

z/OS



Resource Measurement Facility User's Guide

Version 2 Release 1

Note

Before using this information and the product it supports, read the information in “Notices” on page 403.

This edition applies to Version 2 Release 1 of z/OS (5650-ZOS) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 1990, 2015.**

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

Contents

Figures	ix
--------------------------	-----------

Tables	xi
-------------------------	-----------

About this document	xiii
--------------------------------------	-------------

Who should use this document	xiii
How this document is organized	xiii
z/OS information	xiv
How to read syntax diagrams	xv
Symbols	xv
Syntax items	xv
Syntax examples	xvi

How to send your comments to IBM	xix
---	------------

If you have a technical problem	xix
---	-----

Summary of changes	xxi
-------------------------------------	------------

Summary of changes for z/OS RMF User's Guide for Version 2 Release 1, as updated February 2015	xxi
Changes made in z/OS Version 2 Release 1	xxi
z/OS Version 2 Release 1 summary of changes	xxiv

Part 1. Introduction 1

Chapter 1. RMF - your performance management tool. 3

Gathering data.	3
Short-term data collection with Monitor III	3
Snapshot monitoring with Monitor II	4
Long-term data gathering with Monitor I and Monitor III	4
Storing data	5
Reporting data.	5
Short-term interactive performance analysis with Monitor III	5
Snapshot reporting with Monitor II.	6
Long-term overview reporting with the Postprocessor	6
Viewing reports on spreadsheets.	7
Monitoring on the workstation	7
Providing data for other applications	7
Resource monitoring of systems running AIX, Linux, or Windows	8
Accessing data across the sysplex	8
RMF Sysplex Data Server	8
Sysplex data services for SMF data	9
Sysplex data service for Monitor III data	10
Sysplex data gathering service for Monitor II data	10
What you can gather and report	10
Activity monitoring.	10
Delay monitoring	11
Long-term performance analysis with RMF XP	12

Reporting of other SMF data	12
---------------------------------------	----

Part 2. Administration 13

Chapter 2. Setting up RMF 15

Migrating from previous releases	15
Customizing the system environment.	15
Define RMF library authorization	15
Ensure linkage to language environment	16
IPL with the CMB parameter	16
Define an XCF transport class	16
Check the program properties table (PPT)	16
Remove ERBMFRES	17
Global performance data control authority	17
Specifying access definitions.	17
Define RMF user IDs and ensure access to z/OS UNIX System Services.	18
Assign started task procedures to user IDs	18
Considerations for z/OS UNIX level of security	18
Ensure RACF access to the Distributed Data Server (GPMSEVER and GPM4CIM)	19
Ensure READ access for GPM4CIM to the BPX.SMF profile.	19
Configuring PassTicket support for the Distributed Data Server	19
Controlling access to RMF data for the sysplex data services	20
Setting up the RMF control session including Monitor I and Monitor II.	22
Customizing the RMF control session.	22
Specifying priority for RMF	23
Storing gatherer defaults	23
Preallocating Monitor I and Monitor II reporter data sets	23
Setting up the Monitor III gatherer session RMFGAT	24
Defining VSAM data sets.	25
Ensuring common storage tracking	26
Setting up the Distributed Data Server for z/OS	26
DDS options	28
Synchronizing SMF recording intervals	30
Defining SMF record writing	30
Defining SMF synchronization	31
Storing gatherer options	33
Naming parmlib option members	33
Generalizing parmlib members	34
Defining parameters for Monitor I.	35
Defining parameters for Monitor II	37
Defining parameters for Monitor III	39
Considering reporting aspects	39
Setting up the RMF CLISTs	39
Grouping Monitor II and Monitor III Users.	41
Installing workstation functions	41
Installing the RMF Postprocessor XML Toolkit.	41
Installing the RMF Spreadsheet Reporter	41

Installing RMF Performance Monitoring (RMF PM)	41
Installing RMF client/server enabling - RMFCS	41
Installing the z/OSMF Resource Monitoring plug-in	42

Part 3. Operation 43

Chapter 3. Starting and stopping RMF 45

Starting RMF	45
Starting RMF in the sysplex	46
Starting RMF without JES	46
Stopping RMF	47
Controlling the SMF buffer	47

Chapter 4. Starting and controlling Monitor sessions 51

Specifying session options	51
Conflicts between session options	51
Starting RMF sessions	51
Starting a specific Monitor	52
Modifying RMF session options	55
Displaying RMF status	55
Stopping RMF sessions	56
Stopping a specific session	56

Chapter 5. How RMF processes session options 59

When you start an RMF session	59
Displaying a list of options in effect for a session	60
Examples	60
When you modify session options	62
Examples	63

Part 4. Performance management 65

Chapter 6. Performance administration 67

What is performance administration?	67
Defining procedures and parameters	67
Setting performance goals	68
z/OS workload management	68
Service definition	69
Workload management concepts	69
Workloads and service classes	69
Performance goals	70
Importance	70
Service class periods	70
Reporting hierarchy	70
Sysplex view	71
System view	71
Job or resource view	71

Chapter 7. Performance monitoring . . 73

Watching performance goals	73
Observing response times	73
Monitoring throughput	74
Transaction rate	74
Workflow and velocity	75

Observing bottlenecks and exceptions	75
--	----

Chapter 8. Performance analysis . . . 77

What is a performance problem?	77
Reports that provide data for analysis	77

Part 5. Data Gathering Reference 81

Chapter 9. Long-term data gathering with Monitor I 83

Summary of session options	83
Default session options	84
Description of Monitor I data gatherer options	85
CACHE	85
CHAN	85
CPU	86
CRYPTO	86
CYCLE	86
DEVICE	86
ENQ	88
ESS	89
EXITS	89
FCD	90
INTERVAL	90
IOQ	91
MEMBER	92
OPTIONS	93
PAGESP	95
PAGING	95
RECORD	95
REPORT	95
STOP	96
SYNC	96
SYSOUT	97
TRACE	97
VSTOR	103
WKLD	104
VMGUEST	104
Special considerations	104
INTERVAL and CYCLE options	104
STOP, INTERVAL, and REPORT options	105
Device class selection for the DEVICE option	105

Chapter 10. Snapshot data gathering with Monitor II 107

Chapter 11. Short-term data gathering with Monitor III 109

Summary of gatherer session options	109
Default gatherer session options	110
Description of Monitor III data gatherer options	111
CACHE	111
CFDETAIL	111
CYCLE	112
DATASET	112
HFSNAME	113
IOSUB	113
LOCK	113
MASTER	114

MEMBER	114
MINTIME	115
OPD	115
OPTIONS	115
RESOURCE	116
STOP	117
SGSPACE.	117
SYNC	118
SYSOUT	118
VSAMRLS	118
WSTOR	119
ZFS.	120
ZIIPUSE	120
Controlling data set recording	120
Starting data set support	122
Modifying the data set support options.	124
Stopping data set support	124
Data set support for daylight saving time	125

Part 6. Reporting Reference 127

Chapter 12. The online reporting environment 129

Starting the reporters	129
Reference information	130
Quick start	130
Getting help with RMF dialogs	130
Getting help for a report	131
Some words about the tutorial.	131
Message help and stacked messages.	132

Chapter 13. Interactive performance analysis with Monitor III 133

Before you start Monitor III	133
Data set allocation.	133
Transferring Monitor III VSAM data sets to other systems	134
Messages during Monitor III start	135
Sysplex considerations	135
The Monitor III Primary Menu	136
Navigating from the Primary Menu	136
STOP and GO	140
Using the Data Index (DI)	141
Contents of the Data Index.	142
Monitor III session and option commands - Overview.	146
Session commands	146
Option commands.	146
Monitor III session tasks.	147
Backward and forward referencing	147
Cancelling entries on option panels	151
Getting help for RMF commands.	151
Displaying current range data	151
Searching for a field	151
Setting GO mode	152
Activating GRAPHIC mode	152
Printing screens and tabular reports	153
Printing graphic reports	154
Using program function keys	154
Resetting entries on option panels	155

Retrieving the last command	156
Activating TABULAR mode	156
Toggling between tabular and graphic display	156
Displaying user-written reports	156
Changing session options	156
Changing color graphic options	158
Changing language options.	160
Changing report options.	160
Selecting an option set	161
Using cursor-sensitive control	163
Monitor III help facility	163
Monitor III report commands - Overview	163

Chapter 14. Snapshot reporting with Monitor II 167

Monitor II sessions	167
The ISPF session	168
The TSO/E session	171
The background session	172
Structure of Monitor II reports.	172
Display session report fields	172
Different formats of Monitor II reports	174
Different modes of Monitor II reports	174
Monitor II session commands and options - Overview.	175
Display session commands	175
Background-session options	176
Monitor II session tasks	176
Display the menu	177
Display and set options	177
Reset default options	179
Leave options unchanged	179
Display commands	179
Define session options	180
Scroll through report frames	180
Recall the previous data	181
Sort reports	181
Find text string.	181
Repeat Find	182
Set delta mode	182
Create a hardcopy report	182
Refresh a report automatically.	184
Assign PF keys	185
Print a report page	186
Write SMF records.	186
Specify the system to be monitored	186
Stop the session	187
Monitor II report commands - Overview	187
Details of report commands	188
Conflicting session options	198

Chapter 15. Long-term reporting with the Postprocessor 201

Postprocessor reporting	201
Preparing SMF records for postprocessing.	202
Sysplex reporting across time zones	203
Starting the Postprocessor	204
Starting with JCL	205
Starting from ISPF.	206
Defining output data sets	210

How the Postprocessor processes control statements	212
Defining the reporting period	212
Postprocessor completion	213
How to use control statements	213
DATE	214
DELTA	215
DINTV	215
ETOD	218
EXCEPT	218
EXITS	220
EXRPTS	220
OVERVIEW	221
OVW	222
REPORTS	223
RTOD	230
SESSION	230
STOD	231
SUMMARY	231
SYSID	232
SYSOUT	232
SYSRPTS	232
Examples of control statements	234
Overview and exception conditions	238
CPU Activity - SMF record type 70-1	238
Crypto Hardware Activity - SMF record type 70-2	244
Paging Activity - SMF record type 71	248
Workload Activity - SMF record type 72-3	251
Channel Path Activity - SMF record type 73	256
Device Activity - SMF record type 74-1	258
Coupling Facility Activity - SMF record type 74-4	259
Cache Activity - SMF record type 74-5	261
Ficon Director Activity - SMF record type 74-7	267
Enterprise Disk Systems Statistics - SMF record type 74-8	267
PCIE Function Activity - SMF record type 74-9	270
Page Data Set Activity - SMF record type 75	271
Enqueue Activity - SMF record type 77	271
Virtual Storage Activity - SMF record type 78-2	272
I/O Queuing Activity - SMF record type 78-3	272

Chapter 16. Cross platform monitoring with RMF XP. 275

How to set up RMF XP	277
Configuration files and parameters	278
Post-installation steps	280
Security setup	280
How to start GPM4CIM and assign userIDs	281
How to use RMF XP for long-term performance analysis	281
AIX on System p performance data: subtype 1-12	282
Linux on System x performance data: subtype 20-31	282
Linux on System z performance data: subtype 40-53	283
Windows on System x performance data: subtype 60-64	283
How to request SMF record type 104 collection	283

How to authorize GPM4CIM to write SMF record type 104.	284
Diagnostic aids	285

Part 7. Analysis on the workstation 287

Chapter 17. How to work with Postprocessor XML reports 289

Producing and viewing XML reports with the Spreadsheet Reporter	289
Producing and viewing XML reports with the HTTP API of the DDS	289
Viewing XML reports with the RMF Postprocessor XML Toolkit	289
How to install the Postprocessor XML Toolkit	289
How to use the Postprocessor XML Toolkit	290
How to view XML reports in a browser	290

Chapter 18. RMF Spreadsheet Reporter 293

Concepts of performance analysis with the RMF Spreadsheet Reporter	293
Get more out of Postprocessor reports with the Spreadsheet Reporter	294
The Spreadsheet Reporter's resource-oriented concept	296
Installing the Spreadsheet Reporter	298
Prerequisites	298
Installation steps	298
How to use the RMF Spreadsheet Reporter	299
Spreadsheet Reporter resources	300
Spreadsheet Reporter main dialog	302
Menu bar	304
How to work with Resources	304
How to work with Systems	313
How to specify settings for the RMF Spreadsheet Reporter	317
How to specify processing options and report types	317
How to specify reporting periods and duration intervals	322
How to create Working Sets in batch mode	323
Jclgen.bat	323
Collect.bat	324
CreateRptWSet.bat	325
CreateOvwWSet.bat	326
How to migrate Working Sets	326
Using RMF spreadsheet macros	326
Available RMF spreadsheet macros	326
Macros for Report Working Sets	327
Macros for Overview Working Sets	328
Behavior of RMF spreadsheet macros	328
Sets of overview control statements	329
LPAR Overview Report	329
System Overview Report	330
Workload Overview Report	331
Device Overview Report	333
Cache Subsystem Overview Report	334

Channel Overview Report	335
Spreadsheet usage examples	336
Open RMF Overview Spreadsheets usage example	339
Setting the security level for Excel macros	339

Chapter 19. RMF Performance

Monitoring 341

Performance monitoring - Overview.	341
Getting help in RMF PM	342
Putting it all together.	342
Getting started - Installation and setup	342
Prerequisites.	342
Client installation	343
RMF PM - Let's go	343
Sysplex logon dialog	344
Working with sysplexes	345
Sysplex dialog	345
Open and close a sysplex	347
Expand a sysplex	347
Expand a Resource	347
Creating a DataView from a Resource	348
Working with PerfDesks and DataViews	348
Creating a PerfDesk folder	348
Creating a PerfDesk	348
Opening a PerfDesk	348
Expanding the tree	348
Start and stop a PerfDesk	349
Save a PerfDesk	349
Close a PerfDesk	349
Stepping through the history of DataViews	349
Synchronize all DataViews	350
Startup	350
Add a DataView to a PerfDesk	350
Change the name of a PerfDesk	350
Paste DataView to a PerfDesk	350
Exporting and Importing a PerfDesk	351
DataView.	351
Creating a DataView	352
Working with DataViews	352
Sampling dialog	355
Sample intervals	355
Use Common Interval	355
Sample From	355
Sample To	356
Wrap-around buffer size.	356
Set Time-zone	356
Series	356
Series time-stamps.	356
Series definition dialog	356
Plot/Save Series dialog	357
Filter dialog	360
Work Scope dialog	362
Analysis dialog.	363
What is monitored?	363
Sysplexes.	363
Resources	363
Metrics	367
Work Scopes	367
Message Browser	368
Save messages to a file	368

Delete messages	368
Query help for a message	368
Close the Message Browser.	368

Chapter 20. z/OS Management Facility

- Resource Monitoring 369

Chapter 21. RMF Client/Server

enabling (RMFCS) 373

What is RMF Client/Server enabling ?	373
RMFCS monitoring scenarios	375
Scenario I: Message-initiated monitoring	375
Scenario II: Exception-initiated monitoring	377
Scenario III: Command-initiated monitoring	377
Installation and startup of RMFCS components	378
Prerequisites.	378
Installation	378
Customization	378
Startup	381
Setup and usage considerations	382
Multiple PWS connections to multiple systems	382
Prevention of duplicate connections	382
Routing different events to different workstations.	382
GUI session comes up with WFEX	383
Exception handling for all Monitor III report data	383
Exception-initiated monitoring without MPFLSTxx functions	384
The automated approach	384
Components of RMFCS enabling	384
JCL procedures	384
Modules	385
REXX Procedures / CLISTS	385
Listings of RMFCS procedures.	387
REXX procedure ERBCSGUI	387
REXX procedure ERB3RPH3	388
REXX procedure ERBR3WFX	389
REXX procedure ERBR3SYS	393

Part 8. Appendixes 397

Appendix. Accessibility 399

Accessibility features	399
Using assistive technologies	399
Keyboard navigation of the user interface	399
Dotted decimal syntax diagrams	399

Notices 403

Policy for unsupported hardware.	404
Minimum supported hardware	405
Exploitation of the Flash Express feature	405
Programming Interface Information	405
Trademarks	405

Glossary 407

Index 413

Figures

1.	RMF - Your Performance Management Tool	5
2.	RMF Sysplex Data Server Data Flow	9
3.	Console sample output with Monitor I OPTIONS in effect	94
4.	Console Output with OPTIONS in Effect	116
5.	RMF Performance Management Menu	129
6.	Monitor III Primary Menu	136
7.	Monitor III Sysplex Report Selection Menu	137
8.	Monitor III Overview Report Selection Menu	138
9.	Monitor III Job Report Selection Menu	138
10.	Monitor III Resource Report Selection Menu	139
11.	Monitor III Subsystem Report Selection Menu	139
12.	Monitor III User-written Report Selection Menu	140
13.	Header of Monitor III Single-System Reports	140
14.	Data Index	142
15.	Data Index - Condensed Version	143
16.	Data Index Options Panel	145
17.	Option Selection Menu	147
18.	Header of Monitor III Single-System Reports	148
19.	Interactive Chart Utility (ICU) DIRECTORY Panel	154
20.	Session Options Panel.	157
21.	Color Graphic Options - Panel 1	158
22.	Color Graphic Options - Panel 2	159
23.	Language Options Panel.	160
24.	DELAY Report Options Panel	161
25.	Option Set Menu	161
26.	Monitor II Primary Menu	168
27.	Monitor II Address Space Report Selection Menu	169
28.	Monitor II I/O Report Selection Menu	169
29.	Monitor II Resource Report Selection Menu	170
30.	Monitor II Library List and OPT Settings Selection Menu	170
31.	Monitor II User Report Selection Menu	171
32.	Monitor II Display Menu	171
33.	CHANNEL report as example of a Monitor II report header (TSO/E session).	173
34.	ARD, ASD and ASRM Report Options Panel	177
35.	Monitor II TSO/E session - Default Options	178
36.	JCL example for preparing SMF records	204
37.	JCL example for starting the Postprocessor	205
38.	Postprocessor - Setup Panel.	206
39.	Postprocessor Input for SMF Data Sets	208
40.	Postprocessor - Output Data Sets	208
41.	Postprocessor Options panel	209
42.	Relationship between Duration Interval and Reporting Period	217
43.	Cross platform performance monitoring	276
44.	RMF XP startup member in SYS1.PROCLIB(GPM4CIM)	277
45.	GPM4CIM configuration file example for AIX	279
46.	GPM4CIM security concept	281
47.	Log file entries during GPM4CIM processing	285
48.	Sortable data in Postprocessor reports in XML format	291
49.	Spreadsheet - CPU utilization for selected workloads	294
50.	Overview Report and Spreadsheet	296
51.	From SMF Dump Data to a Spreadsheet presentation	297
52.	Spreadsheet Reporter - main dialog	303
53.	New SMF Dump Data	305
54.	Create Report Listing	306
55.	Create Overview Record	308
56.	Create Working Set	310
57.	Available Excel Spreadsheet Macros	312
58.	List of Systems in the view pane	313
59.	Define new System	314
60.	Processing a file containing overview control statements (1)	315
61.	Processing a file containing overview control statements (2)	316
62.	Processing options	318
63.	Report types available in text format	321
64.	Report types available in XML format	322
65.	Reporting periods and duration intervals	323
66.	DASD Activity Report	337
67.	Activity of Top 10 Volumes	338
68.	Spreadsheet Reporter Output — Open RMF Overview Spreadsheets macro	339
69.	RMF PM — main window	343
70.	Sysplex logon dialog	344
71.	Initial DataViews	345
72.	Sysplex Dialog	346
73.	Collecting historical data in RMF PM	350
74.	Sample DataView	351
75.	Filter Dialog	360
76.	Work Scope Dialog.	362
77.	System Status task	370
78.	Resource Monitoring task	371
79.	RMFCS Performance Data View - Example	374
80.	RMFCS Scenario I: Message-Initiated Monitoring	375
81.	RMFCS Scenario II: Exception-Initiated Monitoring	377

Tables

1. Syntax examples	xvi	30. Channel Path Activity - Conditions Based on SMF Record Type 73	257
2. Monitored activities and SMF record types	10	31. Device Activity - Conditions Based on SMF Record Type 74-1	258
3. ddnames for Monitor I and Monitor II Data Sets	24	32. Coupling Facility Activity - Conditions Based on SMF Record Type 74-4	260
4. Summary of Monitor I Session Options	83	33. Cache Activity - Conditions Based on SMF Record Type 74-5	261
5. Monitor I Default Session Options	84	34. Cache Activity - Conditions Based on SMF Record Type 74-5	266
6. Where to specify Monitor I options	94	35. Ficon Director Activity - Conditions Based on SMF Record 74-7	267
7. Monitor III Data Gatherer Session Options	109	36. ESS Link Statistics - Conditions Based on SMF Record 74-8	267
8. Monitor III Default Session Options	110	37. ESS Extent Pool Statistics - Conditions Based on SMF Record 74-8	269
9. Monitor III OPTIONS Command Sources	116	38. ESS Rank Statistics - Conditions Based on SMF Record 74-8	269
10. Field Descriptions for the Data Index	145	39. PCIE Function Activity - Conditions Based on SMF Record 74-9	270
11. Monitor III Session Commands	146	40. Page Data Set Activity - Conditions Based on SMF Record Type 75	271
12. Monitor III Option Commands.	147	41. Enqueue Activity - Conditions Based on SMF Record Type 77	271
13. Program Function Keys Defaults	155	42. Virtual Storage Activity - Conditions Based on SMF Record Type 78-2	272
14. Fields on the Graphic Options Panels	159	43. I/O Queuing Activity - Conditions Based on SMF Record Type 78-3	273
15. Report Commands.	164	44. Supported Reports for spreadsheets	319
16. User-Written Report Commands	166	45. Supported Reports for web browser display	319
17. Monitor II Display Session Areas	173	46. Variables in Jclgen.bat.	323
18. Monitor II Display session commands	175	47. Additional variables in Collect.bat	325
19. Monitor II Background Session Options	176	48. Macros Based on Reports	327
20. Monitor II OPTIONS Command Sources	178	49. Macros Based on Overview Records	328
21. Monitor II Report Commands	187		
22. Postprocessor ddnames - text output	210		
23. Postprocessor ddname - Overview record output	211		
24. Postprocessor ddnames - XML output	212		
25. Postprocessor Control Statement Summary	213		
26. CPU Activity - Conditions Based on SMF Record Type 70-1	239		
27. Crypto Hardware Activity - Conditions Based on SMF Record Type 70-2	244		
28. Paging Activity - Conditions Based on SMF Record Type 71	248		
29. Workload Activity (Goal Mode) - Conditions Based on SMF Record Type 72-3	251		

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.

How this document is organized

This document contains the following parts:

Part 1, “Introduction,” on page 1

This part describes the different components of RMF, and explains how to use them for data gathering, data reporting, and performance management.

Part 2, “Administration,” on page 13

This part gives an overview about the tasks that are required to activate RMF and to tailor all parameters for data gathering according to your requirements.

Part 3, “Operation,” on page 43

Here, you get information about operator tasks for starting, modifying, and stopping the different types of monitoring sessions.

Part 4, “Performance management,” on page 65

The information units contained in this part explain the different tasks that belong to performance monitoring and they describe the various functions of RMF you can use for solving these tasks.

Part 5, “Data Gathering Reference,” on page 81

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 information units.

Part 6, “Reporting Reference,” on page 127

This part deals with the RMF reporting capabilities, and with how to control them. Reports are available to help you with three different tasks:

- Interactive performance analysis, using the Monitor III Reporter Dialog
- Snapshot reporting, using the Monitor II Display Session, with the option of producing reports in printed form
- Long-term overview reporting, using the Postprocessor

In addition, Chapter 16, “Cross platform monitoring with RMF XP,” on page 275 in this part describes how to set up, configure and start RMF XP if you want to monitor the performance of heterogeneous environments running the following operating systems:

- AIX[®] on System p[®]
- Linux on System x[®]
- Linux on System z[®]
- Windows on System x

Part 7, “Analysis on the workstation,” on page 287

In addition to host-based reporting functions in RMF, there are other components available that offer reporting capabilities on the workstation:

- The **RMF Postprocessor XML Toolkit** helps you to browse the Postprocessor reports that are available in XML output format, with your internet browser using the stylesheets provided by RMF.
- The **RMF Spreadsheet Reporter** assists you in converting Postprocessor listings and Overview records into spreadsheets. In addition, it provides sample spreadsheets to help you in presenting and analyzing performance data at a glance.
- The **RMF Performance Monitoring (RMF PM)** provides an interface between your workstation and the z/OS sysplex that gives you the flexibility to create unique scenarios that monitor the performance of your system.
- The **IBM z/OS Management Facility (z/OSMF)** is a web-browser based management console for z/OS. The *z/OSMF Resource Monitoring* plug-in allows cross-sysplex performance monitoring from a single point of control. From the z/OSMF task tree, you can select the following subtasks:
 - The *System Status task* provides an enterprise-wide health check of all z/OS sysplexes.
 - For further analysis, the *Resource Monitoring* task can graphically display RMF Monitor III metrics as well as AIX, Linux, or Windows metrics by means of customizable views.
- The **RMF Client/Server Enabling (RMFCS)** is a concept that makes your performance management independent of a TSO host session. It allows you to establish as many ISPF GUI sessions as you want with any z/OS system in your network that have an APPC or TCP/IP connection configured to your workstation.

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, including the z/OS Information Center, go to the z/OS Internet library (<http://www.ibm.com/systems/z/os/zos/bkserv/>).

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.	
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.	

Table 1. Syntax examples (continued)

Item	Syntax example
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.	

How to send your comments to IBM

We appreciate your input on this publication. Feel free to comment on the clarity, accuracy, and completeness of the information or provide any other feedback that you have.

Use one of the following methods to send your comments:

1. Send an email to mhvrcfs@us.ibm.com.
2. Send an email from the "Contact us" web page for z/OS (<http://www.ibm.com/systems/z/os/zos/webqs.html>).
3. Mail the comments to the following address:
IBM Corporation
Attention: MHVRCFS Reader Comments
Department H6MA, Building 707
2455 South Road
Poughkeepsie, NY 12601-5400
US
4. Fax the comments to us, as follows:
From the United States and Canada: 1+845+432-9405
From all other countries: Your international access code +1+845+432-9405

Include the following information:

- Your name and address.
- Your email address.
- Your telephone or fax number.
- The publication title and order number:
z/OS V2R1.0 RMF User's Guide
SC34-2664-01
- The topic and page number that is related to your comment.
- The text of your comment.

When you send comments to IBM, you grant IBM a nonexclusive right 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

Do not use the feedback methods that are listed for sending comments. Instead, take one of the following actions:

- Contact your IBM service representative.
- Call IBM technical support.
- Visit the IBM Support Portal at z/OS support page (<http://www.ibm.com/systems/z/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 z/OS RMF User's Guide for Version 2 Release 1, as updated February 2015

Changed information

This edition includes the following topics that contain changed information in support of IBM z13:

- “CPU Activity - SMF record type 70-1” on page 238.
- “Crypto Hardware Activity - SMF record type 70-2” on page 244.
- “Paging Activity - SMF record type 71” on page 248.
- “Workload Activity - SMF record type 72-3” on page 251.
- “PCIe Function Activity - SMF record type 74-9” on page 270.
- “Page Data Set Activity - SMF record type 75” on page 271.
- “I/O Queuing Activity - SMF record type 78-3” on page 272.

Changes made in z/OS Version 2 Release 1

This document contains information previously presented in *z/OS RMF User's Guide*, SC33-7990-19, which supports z/OS Version 1 Release 13.

Exploitation of IBM System z Integrated Information Processors

The RMF Monitor III data gatherer RMFGAT has been entitled to partially run on IBM System z Integrated Information Processors (zIIPs). A new Monitor III gatherer option ZIIPUSE/NOZIIPUSE is introduced to determine whether the RMFGAT address space is eligible for zIIP exploitation.

Statistics about CF structures residing in Storage Class Memory

Storage class memory (SCM) usage and statistics information is available for coupling facilities and structures which are allocated with storage class memory.

RMF provides SCM related information in SMF record type 74-4, as well as in the *SCM Structure Summary* and the *Storage Summary* of the *Usage Summary* section in the Postprocessor *Coupling Facility Activity* report.

For structures allocated with SCM, the Monitor III *Coupling Facility Activity* (CFACT) report displays a new Structure Details pop-up window, showing SCM measurements and general structure data.

In addition, new overview conditions are provided for the Postprocessor based on the enhanced SMF record 74-4.

RMF uses the term *storage class memory (SCM)* as a synonym for *Flash Express memory*.

Monitoring PCIe function and zEDC activity

A new Postprocessor *PCIE Activity Report* is available in XML output format and provides measurements about the activity of PCI Express based functions (PCIe functions) and their exploitation of hardware accelerators.

A PCIe function is captured by the report if one of the following hardware feature activities has been measured:

- RDMA (Remote Direct Memory Access) over Converged Enhanced Ethernet
- zEnterprise Data Compression (zEDC) capability using zEDC Express

In addition, RMF provides new overview conditions for the Postprocessor based on a new subtype 9 of SMF record 74.

Support of Group Capacity enhancements and absolute LPAR capacity limits

WLM introduces negative phantom weights for softcapping and uses initial weights to distribute the group capping limit when it becomes necessary to enforce the group limits. RMF adds new fields to SMF record 70-1 and takes the new WLM functionality into account when reporting about capacity groups.

RMF adds support to report on the new absolute LPAR capacity limit that can be defined via the logical partition controls of the Hardware Management Console (HMC). The Postprocessor *Partition Data* report and the Monitor III *CPC Capacity* report display whether either Initial Capping or an absolute LPAR capacity limit was active during a reporting interval.

New RMF Postprocessor overview conditions based on SMF record 70-1 can be used for a more detailed analysis of the hardware capping options.

Enhanced Postprocessor Crypto Hardware Activity report

RMF enhances the Postprocessor *Crypto Hardware Activity* report to provide activity measurements from the Crypto Express4S (CEX4) card configured in one of the three ways:

- Cryptographic CCA coprocessor
- Cryptographic PKCS11 coprocessor
- Cryptographic accelerator

New overview conditions are provided for the Postprocessor, based on the enhanced SMF record 70-2.

Additional Postprocessor reports in XML format

By specifying appropriate ddnames in the job for the Postprocessor output, users can request the following reports in XML output format:

- *Cache Subsystem Activity*
- *Channel Path Activity*
- *Coupling Facility Activity*
- *Enqueue Activity*
- *Hierarchical File System Statistics*
- *I/O Queuing Activity*

- *Page Data Set Activity*
- *PCIE Activity Report*
- *Shared Device Activity*
- *Virtual Storage Activity*
- *XCF Activity*

Cross platform monitoring support for Windows

Beyond the support of the AIX and Linux operating systems, RMF XP has been extended to support Windows systems as monitored endpoints. With the Resource Monitoring plug-in for IBM z/OS Management Facility (z/OSMF), performance metrics from Windows systems can be displayed in the same way and together with metrics from other platforms.

SMF Recording Facility for AIX, Linux and Windows performance data

You can now use RMF XP for long-term performance analysis and capacity planning of your AIX, Linux and Windows systems. For this purpose, you can write performance data collected from the monitored endpoints to the new SMF record type 104.

Controlling the invocation of data reduction exit routines

RMF now controls the names of the data reduction exit routines that are invoked by callers of RMF Monitor II Sysplex Data Gathering service ERB2XDGS or Monitor III Sysplex Data Retrieval service ERB3XDRS.

This documentation describes how to apply the required access control. Especially, if an unauthorized application is making use of parameter <exit_name>, a new RACF resource profile of class FACILITY is required.

Enhancement for collecting VSAM RLS activity data

When specifying the VSAMRLS option to gather data for the Monitor III VSAM RLS Activity report, the limit for active data set masks has been raised from 25 to 50.

Monitoring of pageable large pages activity

RMF provides enhanced performance measurements about memory objects and frames in the following reports:

- In the Postprocessor *Paging Activity* report, the *Memory Objects and Frames* section has been renamed to *Memory Objects and High Virtual Storage Frames* and now contains the following enhanced measurements:
 - additional metrics for high virtual common and shared storage frames
 - metrics for 1 MB frames are now reported in more detail
 - number of auxiliary storage slots for frames from virtual common and shared storage backed on DASD.

In addition, RMF provides new overview conditions for the Postprocessor based on SMF record 71.

- In the Postprocessor *Virtual Storage Activity* report, the information about 1 MB frames in the Private Area Detail section is now separated into the categories *fixed* and *pageable*.
- The Monitor III **Storage Memory Objects** report now provides measurements for 1 MB frames in more detail at system and address space level.

User controlled location of the RMF Master Gatherer system

The new Monitor III data gatherer option MASTER/NOMASTER is introduced to determine whether an individual system is eligible for master gathering.

z/OS Version 2 Release 1 summary of changes

See the following publications for all enhancements to z/OS Version 2 Release 1 (V2R1):

- *z/OS Migration*
- *z/OS Planning for Installation*
- *z/OS Summary of Message and Interface Changes*
- *z/OS Introduction and Release Guide*

Part 1. Introduction

The introduction provides an overview of the capabilities of RMF.

- Data Gathering with Monitor I, Monitor II, and Monitor III
- Reporting with Monitor III, Monitor II, and the Postprocessor
- Creating spreadsheets with the Spreadsheet Reporter
- Resource monitoring of systems running Linux, AIX, or Windows with RMF XP
- Resource monitoring on the workstation with RMF PM

You also find information about the Sysplex Data Server for accessing data across the sysplex.

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
- Client/Server Enabling
- Sysplex Data Server
- Distributed Data Server for z/OS
- Distributed Data Server for AIX, Linux and Windows systems (RMF XP)
- RMF 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

Storing data

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.

Reporting data

All three monitors can create reports, and so does the Postprocessor.

Short-term interactive performance analysis with Monitor III

The Monitor III reporter runs in a TSO/E session under ISPF and provides sysplex or system performance reports by:

- Displaying your current system status in real time mode
- Showing previously collected data that is still available in either in-storage buffers or preallocated VSAM data sets

Monitor III offers a wide spectrum of reports answering questions that arise during the various performance management tasks.

Cursor-sensitive control is one specific highlight of the Monitor III reporter you can use to navigate among different types of reports that all describe the system status at the same point in time from different perspectives. Once you have used it, you

will never want to be without it — it helps you to get the report that points directly to the problems you need to solve.

All reporting is available within one TSO/E session, so there's no need to logon to different systems in the sysplex to get all performance data. All reports are available on one screen.

Snapshot reporting with Monitor II

Monitor II is a snapshot reporting tool for very fast information about how specific address spaces or system resources (processor, DASD volumes, storage) are performing. Monitor II has two modes for reporting on the performance of your system:

- Monitor II display session: You select the ISPF version in the RMF Performance Management menu, or you call the monitor with the TSO/E command RMFMON.
- Monitor II background session: You start a non-interactive session to create a report for printing.

Some reports offer continuous monitoring of single address spaces or DASD devices. You can get an one-line report each time you press ENTER, or you can request a periodically refreshed report.

Long-term overview reporting with the Postprocessor

Typically, you call the Postprocessor in a batch job, although running it in a TSO/E session is possible. You provide a set of options that define the scope of reporting and you get reports of various types with all the data you need for optimum running of your system.

The standard procedure is to allocate SMF data sets or SMF log streams with records from all monitors as input for the Postprocessor. A variation is to get reports on the RMF records that are available in the RMF data buffers of all systems in the sysplex while the Postprocessor is running. This data is automatically made available to the Postprocessor by calling the RMF Sysplex Data Server, a quick path to access performance data without having to go through dumping, sorting, and merging the SMF records.

The Postprocessor offers different types of report:

Interval reports: they show the sysplex performance for each interval for which data has been gathered. Most single-system reports are also available as real-time reports from Monitor I.

Duration reports: the data is summarized over longer periods of time with a maximum value of 100 hours.

Summary, exception, and overview reports: these Postprocessor capabilities let you create the reports you need to manage the performance of your system.

In addition, the Postprocessor can create Overview records which are the base for further spreadsheet processing on the workstation.

The Postprocessor can also generate a set of Postprocessor reports in XML format for display in a web browser.

Viewing reports on spreadsheets

The **Spreadsheet Reporter** is the function in RMF that assists you in converting Postprocessor listings and Overview records into spreadsheets. In addition, it provides sample spreadsheets to help you in presenting and analyzing performance data at a glance.

Monitoring on the workstation

IBM z/OS Management Facility (z/OSMF) is a web-browser based management console for z/OS. The *z/OSMF Resource Monitoring* plug-in allows cross-sysplex performance monitoring from a single point of control. From the z/OSMF task tree, you can select the following subtasks:

- The **System Status task** provides an enterprise-wide health check of all z/OS sysplexes.
- For further analysis, the **Resource Monitoring task** can graphically display RMF Monitor III metrics as well as AIX, Linux, or Windows metrics by means of customizable views.

RMF Performance Monitoring (RMF PM) gives you the capability to construct monitoring scenarios and use them whenever necessary. This is done on the Windows workstation, and the access to the current performance data of your z/OS systems is possible without the need to have a TSO/E session running.

Client/Server Enabling (RMFCS) uses the client/server concept to support performance management for z/OS systems without an active TSO/TCAS subsystem on the host.

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

RMFCS supports event-driven monitoring. That is, predefined events on the MVS™ hosts can be configured to initiate performance monitoring. These events may be either specific system messages, or selected performance data counters that exceed predefined Monitor III exception thresholds.

Providing data for other applications

RMF Distributed Data Server (DDS): Applications that want to access sysplex-wide performance data, can retrieve their input from a single data server on one system in the sysplex, which gathers the data distributed on all systems in the sysplex. Therefore, this is called the Distributed Data Server (DDS).

The DDS offers an HTTP API which can access short-term data from the Monitor III as well as long-term data from the Postprocessor. An application program can send an HTTP request for selected performance data to the DDS.

For more information refer to “Setting up the Distributed Data Server for z/OS” on page 26 and to the *z/OS RMF Programmer’s Guide*.

z/OS Common Information Model (CIM): This z/OS component allows access to RMF performance data from within systems management applications. These applications (called CIM clients) invoke the CIM server, which returns z/OS performance metrics collected by RMF Monitor III.

RMF can also generate indications for monitoring data (for example whether a certain performance metric value is above a threshold), this way enabling CIM clients to support event-based monitoring.

You find information on how RMF supports CIM in the *z/OS RMF Programmer's Guide*.

Resource monitoring of systems running AIX, Linux, or Windows

The z/OS RMF Cross Platform (RMF XP) Distributed Data Server provides CIM-based performance data gatherers to monitor AIX on System p, Linux on System z, Linux on System x and Windows on System x. With RMF XP you can monitor operating systems that can run on an IBM zEnterprise® System, including the zEnterprise Blade Center Extension (zBX).

You can exploit the RMF XP capabilities in the following ways:

- With the *Resource Monitoring* task of the *IBM z/OS Management Facility (z/OSMF)*, performance metrics from connected AIX, Linux or Windows systems can be displayed in the same way and together with z/OS in heterogeneous customer environments.
- Exploiters of RMF XP can send an HTTP request to retrieve performance data from the endpoints running the AIX, Linux or Windows operating systems. Requests can be directed against an instance of the RMF XP core component called GPM4CIM, as soon as it is configured and running. GPM4CIM returns the requested data as a structured XML document.

You find information on how to set up RMF XP in Chapter 16, “Cross platform monitoring with RMF XP,” on page 275. It provides information on how to exploit the HTTP API of the RMF XP DDS if you want to submit requests for AIX, Linux and Windows performance data.

Accessing data across the sysplex

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

- “RMF Sysplex Data Server”
- “Sysplex data services for SMF data” on page 9
- “Sysplex data service for Monitor III data” on page 10
- “Sysplex data gathering service for Monitor II data” on page 10

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 application program may call the RMF Sysplex Data Services from any point in the sysplex.

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 20.

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

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.

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.

What you can gather and report

The type of RMF session you run depends on what you need to know about your system. This section describes which sessions measure and report on each type of activity in the system and the various types of delays. Depending on the type of activity and the system environment, the reports can be either sysplex or single-system reports.

Activity monitoring

The RMF gatherer sessions create either SMF or VSAM data that are available for reporting sessions. The following table

- displays the SMF type of all records that will be written by gatherer sessions
- indicates all Monitor III data stored in VSAM data sets
- shows all report capabilities

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	★	★	★	★
	★			Enclave	★			
	★	79.7	77	Enqueue	★	★	★	★
			74.8	Enterprise Storage Server (ESS)				★
			74.7	FICON director				★
		79.15		IRLM long locks		★		

Table 2. Monitored activities and SMF record types (continued)

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					
	★	79.14	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				★
72.2/4	★	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	★			★

Delay monitoring

In addition to monitoring and reporting system activity, Monitor III reports provide various types of delay information.

Delayed address spaces and groups

For each address space or group of address spaces, Monitor III reports the delay experienced for the report interval and identifies the primary cause for the delay:

- System (all jobs)
- TSO, batch, and started tasks
- ASCH and OMVS address spaces
- Service and report classes and workload groups
- Enclaves

For any service class, report class and workload group, Monitor III reports on response time breakdown, using the GROUP report to display the information.

Delay reasons for address spaces

For each of the above address space groups Monitor III offers information which of the following resources or subsystems caused the delays:

- CICS and IMS subsystem
- Devices
- Enclaves
- Enqueues
- HSM
- JES
- Operator (message, mount, and quiesce)
- Processors
- XCF

Long-term performance analysis with RMF XP

To enable long-term performance analysis of AIX, Linux and Windows systems, you can turn on SMF recording for SMF record type 104. This record type provides one range of subtypes for each supported platform. One specific subtype is used to keep the data for one individual CIM metric category according to the CIM data model on the affected platform.

Subtype 1-12

AIX on System p performance data

Subtype 20-31

Linux on System x performance data

Subtype 40-53

Linux on System z performance data

Subtype 60-64

Windows on System x performance data

For information on the metric categories provided in the subtypes and how to request the collection of SMF record type 104 from the systems of all or selected supported platforms, refer to Chapter 16, “Cross platform monitoring with RMF XP,” on page 275.

Reporting of other SMF data

The Postprocessor provides two reports that are based on SMF data that have been gathered outside of RMF.

WebServer performance reporting

The Postprocessor offers an HTTP Server report to support this important e-business application. The report (based on **SMF record type 103** written by the WebServer) provides usage statistics as well as performance information about the WebServer to assist you in tuning and capacity planning.

Lotus Domino support

The Postprocessor Lotus Domino Server report accepts the **SMF record type 108** written by Lotus Domino and provides feedback on server load as well as the number and type of messages that the server handled.

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.
- Preallocate **reporter data sets** for Monitor I and Monitor II output, to be used instead of the default SYSOUT.
- Tailor the **options** for the Distributed Data Server.
- 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 5, "Data Gathering Reference," on page 81 and Part 6, "Reporting Reference," on page 127.

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.

In addition, this chapter points to installation steps that have to be performed individually by everybody who wants to exploit the **workstation-based functions** that are available with RMF.

Chapter 2. Setting up RMF

After installing RMF, you have to perform certain 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 data sets that you can preallocate, and how to specify them in the start-up procedure
- the JCL procedure for starting the Monitor III gatherer session
- the definition of Monitor III gatherer VSAM data sets
- tailoring of the options for the Distributed Data Server
- synchronization with SMF data recording
- the parmlib members that contain your system's standard gatherer options
- Considering reporting aspects
- the installation of workstation functions

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 Migration*.

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*.

Ensure linkage to language environment

Two components of RMF, the Postprocessor and the Distributed Data Server (GPMSEVER, GPM4CIM), use the services of the Language Environment®. They need access to the data set SYS1.SCEERUN. GPM4CIM additionally needs access to data set SYS1.SCEERUN2. There are two ways of providing this access:

- The recommended way is to include data sets SYS1.SCEERUN and SYS1.SCEERUN2 in the LINKLST of the system on which RMF is running. No further action is then required when starting the separate components.
- If you do not wish to include SYS1.SCEERUN or SYS1.SCEERUN2 in the LINKLST, you must specify these data sets as the STEPLIB of the job step that starts the component.

You can use the following JCL members to do this: for the Postprocessor use ERBSAMPP in SYS1.SAMPLIB and for the Distributed Data Server use GPMSEVER or GPM4CIM in SYS1.PROCLIB.

IPL with the CMB parameter

If you intend to monitor devices other than Tape and DASD, you must IPL with the CMB system parameter and describe the number of extra measurement blocks required. One extra measurement block is required for each extra device number to be monitored. See *z/OS MVS Initialization and Tuning Reference* for more information on this parameter.

Define an XCF transport class

The RMF Sysplex Data Server uses XCF services for its intersystem communication.

It is recommended to keep the number of transport classes small. In most cases, it is more efficient to pool the resources and define the transport class based on message size. For more details, please refer to *Parallel Sysplex® Performance: XCF Performance Considerations* available at <http://www-1.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP100743>.

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 18
- “Assign started task procedures to user IDs” on page 18
- “Considerations for z/OS UNIX level of security” on page 18
- “Ensure RACF access to the Distributed Data Server (GPMSEVERE and GPM4CIM)” on page 19
- “Ensure READ access for GPM4CIM to the BPX.SMF profile” on page 19
- “Configuring PassTicket support for the Distributed Data Server” on page 19
- “Controlling access to RMF data for the sysplex data services” on page 20

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. We recommend to define three user IDs that relate to the three started tasks that RMF provides (of course, you may decide to define only one user ID that is assigned to all three 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) and the RMF Distributed Data Server (GPMSEVER) need to be defined so that they can obtain the correct data and can use the required UNIX system services.

The following example contains RACF commands to define the three 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('/'))
ADDUSER GPMSEVER DFLTGRP(omvsgrp) OMVS(UID(nnn) HOME('/'))
```

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

Assign started task procedures to user IDs

RMF provides three different started tasks. In this step, you define these started tasks to RACF and assign them to the RMF user IDs.

```
RDEFINE STARTED RMF.*      STDATA(USER(RMF)      TRUSTED(YES))
RDEFINE STARTED RMFGAT.*   STDATA(USER(RMFGAT)    TRUSTED(YES))
RDEFINE STARTED GPMSEVER.* STDATA(USER(GPMSEVER)  TRUSTED(YES))
RDEFINE STARTED GPM4CIM.*  STDATA(USER(GPMSEVER)  TRUSTED(YES))
SETROPTS RACLIST(STARTED) REFRESH
```

The Distributed Data Server uses the IWMSRSRG service to register itself for sysplex routing. This service is an authorized service. Therefore, the calling DDS user ID GPMSEVER must either have the attribute TRUSTED or must have explicit READ access to the BPX.WLMSEVER Facility. Otherwise, the DDS cannot propagate hostname and port number for potential exploiters. If you did not mark the GPMSEVER task as TRUSTED(YES), as shown in the example above, you must grant access for this task to the RACF Facility BPX.WLMSEVER as shown below:

```
PERMIT BPX.WLMSEVER CLASS(FACILITY) ID(GPMSEVER) ACCESS(READ)
```

Considerations for z/OS UNIX level of security

If the BPX.DAEMON FACILITY resource is defined, your system has z/OS UNIX security and can exercise more control over your superusers.

Because the RMF distributed data server runs as a daemon, it must have access to the BPX.DAEMON facility, and all programs loaded by GPMSEVER and GPM4CIM must be defined to PROGRAM CONTROL. In addition, access to the BPX.SERVER and BPX.STOR.SWAP facilities must be defined for user ID GPMSEVER.

The minimum definitions for the RMF Distributed Data Server are listed in this example. You can use more generic definitions.

```

PERMIT BPX.DAEMON CLASS(FACILITY) ID(GPMSSERVE) ACCESS(READ)
PERMIT BPX.SERVER CLASS(FACILITY) ID(GPMSSERVE) ACCESS(READ)
PERMIT BPX.STOR.SWAP CLASS(FACILITY) ID(GPMSSERVE) ACCESS(READ)
RDEFINE PROGRAM GPM*      ADDMEM('SYS1.SERBLINK'//NOPADCHK)  UACC(READ)
RDEFINE PROGRAM ERB*      ADDMEM('SYS1.SERBLINK'//NOPADCHK)  UACC(READ)
RDEFINE PROGRAM CEEBINIT  ADDMEM('CEE.SCEERUN'//NOPADCHK)    UACC(READ)
RDEFINE PROGRAM IEEMB878  ADDMEM('SYS1.LINKLIB'//NOPADCHK)    UACC(READ)
RDEFINE PROGRAM CELHV003  ADDMEM('SYS1.SCEERUN2'//NOPADCHK)   UACC(READ)
RDEFINE PROGRAM C128      ADDMEM('SYS1.SCEERUN2'//NOPADCHK)   UACC(READ)
RDEFINE PROGRAM CELHDCPP  ADDMEM('SYS1.SCEERUN2'//NOPADCHK)   UACC(READ)
SETROPTS WHEN(PROGRAM) REFRESH
SETROPTS RACLIST(FACILITY) REFRESH

```

Ensure RACF access to the Distributed Data Server (GPMSSERVE and GPM4CIM)

The RMF Distributed Data Server (GPMSSERVE) uses the RACF application name GPMSSERVE. If the RACF Application class (APPL) is active and the GPMSSERVE application is protected by a profile in this class, a user must have read access to this profile. Otherwise RACF does not allow the user to access the GPMSSERVE application.

Correspondingly, the GPM4CIM component of the DDS provided by RMF XP uses the RACF application name GPM4CIM. If this application is protected by a profile in the same way as described for GPMSSERVE, you also need to provide read access to this application.

The minimum definitions for the DDS for both GPMSSERVE and GPM4CIM are listed in this example. You can also use more generic definitions.

```

RDEFINE APPL GPMSSERVE  UACC(READ)
RDEFINE APPL GPM4CIM    UACC(READ)

```

Ensure READ access for GPM4CIM to the BPX.SMF profile

To write SMF record type 104, the GPM4CIM started task needs at least READ access to the RACF BPX.SMF profile of the FACILITY class, specified with your security authorization facility (SAF) product. See “How to authorize GPM4CIM to write SMF record type 104” on page 284 for a RACF example.

Configuring PassTicket support for the Distributed Data Server

If the RMF Distributed Data Server (DDS) is configured to require authentication (see “Setting up the Distributed Data Server for z/OS” on page 26), instead of a user ID and a password, a user ID and a PassTicket can be supplied.

For more information about PassTickets, see the *z/OS Security Server RACF Security Administrator's Guide* (SA22-7683).

A PassTicket is validated against an application name. The RACF application name of the DDS is GPMSSERVE. Before creating the necessary application profile, the RACF class PTKTDATA must be activated:

```

SETROPTS CLASSACT(PTKTDATA)
SETROPTS RACLIST(PTKTDATA)

```

Define a DDS application profile with an associated encryption key:

```

RDEFINE PTKTDATA GPMSSERVE SSIGNON(KEYMASKED(<key>))

```

where *<key>* is a user-supplied 16-digit value used to generate the PassTicket. You can specify a value of your choice. Valid characters are 0 - 9 and A - F.

The user calling the DDS must have RACF permissions in order to generate PassTickets. Define a profile in the PTKTDATA class controlling access to the PassTicket services and explicitly set the universal access authority to NONE:

```
RDEFINE PTKTDATA IRRPTAUTH.GPMSEVE.* UACC(NONE)
```

The user ID connecting to the DDS needs update permission to the newly created profile:

```
PERMIT IRRPTAUTH.GPMSEVE.* CLASS(PTKTDATA) ID(<user>) ACCESS(UPDATE)
```

where *<user>* is the user ID connecting to the DDS. In a CIM environment, this is the user ID associated to the CIM server started task.

Finally you must activate the changes:

```
SETROPTS RACLIST(PTKTDATA) 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). One application using these services is the RMF Postprocessor, if the SMF records are retrieved directly from the SMF buffers.

Also, if you want to exploit the DDS HTTP API (see the *z/OS RMF Programmer's Guide*), you must grant read access to the ERBSDS.SMFDATA profile for the GPMSEVE user ID, which is assigned to the DDS started task GPMSEVE as described in "Assign started task procedures to user IDs" on page 18.

Another application using the mentioned services is the data gatherer of the Monitor II ILOCK command.

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. For example, a Monitor II reporter session invokes this service when reporting about another system in the sysplex.

ERBSDS.MON3DATA

controls access to Monitor III set-of-samples data by the ERB3XDRS

service. For example, the Distributed Data Server as server address space for users of RMF PM calls this service. If this profile is defined, the TSO user ID of RMF PM users must be authorized. Also, a Monitor III reporter session calls this service when sysplex-wide reports are requested.

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 and reason codes, refer to *z/OS RMF 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 20 and "Controlling the invocation of data reduction exit routines."

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”
- “Specifying priority for RMF” on page 23
- “Storing gatherer defaults” on page 23
- “Preallocating Monitor I and Monitor II reporter data sets” on page 23

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 24).

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 47. 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” on page 53)
- specifying the Monitor I gatherer session options, for example, PARM='MEMBER(10),NOCACHE', (see Chapter 9, “Long-term data gathering with Monitor I,” on page 83)

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 IEFARM 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 33.

Preallocating Monitor I and Monitor II reporter data sets

RMF dynamically allocates all Monitor I and Monitor II message and report data sets to that SYSOUT class that is defined in member ERBRMFxx. However, if you want to route output data to permanent data sets rather than to SYSOUT, you can allocate appropriate data sets in the cataloged procedure.

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.

The message and report data sets that RMF uses for Monitor I and Monitor II sessions, and the ddnames for these data sets are:

Table 3. ddnames for Monitor I and Monitor II Data Sets

ddname	Session	Contains	Allocations	Notes®
MFMESSGE	Monitor I, Monitor II background session	General messages	One allocated each time RMF is started.	To change the SYSOUT class parameter for this data set, you must preallocate the data set.
RMFSCxx	Monitor I, Monitor II background session	Messages pertaining to a particular session	One allocated for each session.	xx is the session identifier (ZZ for Monitor I, III for Monitor III, or the Monitor II session identifier you specified).
MFRnnnnn	Monitor I session	Report output	One ddname and one data set allocated for each interval during the session.	nnnnn is a decimal number from 00001 to 99999; successively generated. For example, if a session has 15 intervals, ddnames are MFR00001 through MFR00015.
MFEnnnnn	Monitor I session	Report output after a recoverable abnormal end	One ddname and one data set allocated for each interval during the session.	RMF uses this data set to re-allocate report data sets after a recoverable ABEND. nnnnn is a decimal number from 00001 to 99999; successively generated. For example, if a session has 15 intervals, ddnames are MFE00001 through MFE00015.
RMFxxnnn	Monitor II background session	Report output	One data set and one ddname allocated for each report.	xx is the session identifier, and nnn is a decimal number from 001 to 999, successively generated. RMF uses only one ddname for each report, regardless of the number of intervals in the session. If you modify session options to stop and then restart a particular report, a new ddname is created when the report is restarted.

Note: If you omit the data set control block (DCB) characteristics for the message and report data sets described above, the characteristics used are:

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

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

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 25
- “Ensuring common storage tracking” on page 26

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 46). You can use a DD DUMMY statement if you do not wish to store the gatherer messages.

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)” on page 141) 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 34 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:

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

“Controlling data set recording” on page 120 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.

Setting up the Distributed Data Server for z/OS

Applications that want to access sysplex-wide performance data, can retrieve their input from a single data server on one system in the sysplex, which gathers the data distributed on all systems in the sysplex. Therefore, this is called the Distributed Data Server (DDS).

The DDS offers an HTTP API which can access short-term data from the Monitor III as well as long-term data from the Postprocessor. An application program can send an HTTP request for selected performance data to the DDS.

Note: Equivalent to the DDS component GPMSEVERE, which gathers performance data from z/OS systems in a sysplex, RMF XP provides a second DDS component called GPM4CIM, which you can use to gather performance data from AIX, Linux and Windows operating systems. For general and setup information about RMF XP refer to Chapter 16, “Cross platform monitoring with RMF XP,” on page 275. The *z/OS RMF Programmer’s Guide* provides information on how to exploit the HTTP API of GPM4CIM if you want to submit requests for AIX, Linux and Windows performance data. The remaining information in the current topic is about GPMSEVERE for z/OS only.

Exploiters of Monitor III performance data provided by the DDS are, among others, z/OS Capacity Provisioning, z/OSMF, or RMF PM. If you want to monitor systems in a sysplex, you must set up a Distributed Data Server (DDS) host session on the system in the sysplex with the highest RMF release. If you want to monitor several sysplexes, each one needs to have an active DDS.

To start the DDS, RMF provides the cataloged procedure stored in SYS1.PROCLIB(GPMSEVERE):

```
//GPMSERVE PROC MEMBER=00
//STEP1 EXEC PGM=GPMDDSRV,REGION=128M,TIME=1440,
// PARM='TRAP(ON)/&MEMBER'
...
//GPMPPJCL DD DISP=SHR,DSN=SYS1.SERBPWSV(GPMPPJCL)
...
```

You can modify it according to your requirements.

Prerequisites for exploiting the Monitor III HTTP API

On those systems where you want to monitor short-term Monitor III data, you need to start the Monitor III gatherer with identical MINTIME and SYNC options (see “Description of Monitor III data gatherer options” on page 111).

Also make sure, that the following prerequisites are met on your z/OS host:

- Unix System Services must be configured in *full function mode*.
- TCP/IP under Unix System Services must be configured and fully initialized.

Prerequisites for exploiting the Postprocessor HTTP API

To get access to Postprocessor data provided by the DDS, the GPMSERVE started task points to a Postprocessor job called GPMPPJCL. A JCL template for this job is stored in SYS1.SERBPWSV(GPMPPJCL).

You must adapt or replace the GPMPPJCL member to suit your installation, ensuring that the DDS is able to run RMF Postprocessor jobs. If you do not want to request Postprocessor data with the DDS, you can omit the GPMPPJCL DD card from the GPMSERVE started task.

By default, the RMF Postprocessor retrieves data from all available intervals in the SMF buffer. You can modify the GPMPPJCL template to retrieve only SMF data from the most current interval by adding a job step similar to the following GETSMF sample job step:

```
//GETSMF EXEC PGM=ERBAPPL,PARM='*/70:78'
//SMFDATA DD DISP=(NEW,PASS),UNIT=SYSDA,SPACE=(CYL,(2,2))
//ERBLIST DD SYSOUT=*
```

The example job step retrieves SMF data from records 70 through 78 for the most current interval, indicated by the question mark used for the start time. If you use ERBAPPL in your GPMPPJCL job, you must add an MFPPIN DD card to the Postprocessor job step RMFP in GPMPPJCL with a backward reference to the SMFDATA DD card of the GETSMF job step, for example:

```
//MFPPIN DD DISP=(OLD,PASS),DSN=*.GETSMF.SMFDATA
```

GPMPPJCL template

```
/*JOBPARM SYSAFF=*
//RMFP EXEC PGM=ERBRMFPP
//MFPPMSGDS DD SYSOUT=*
//XPRPTS DD SYSOUT=*
//XPXSRPTS DD SYSOUT=*
//XPOVWRPT DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//SYSIN DD *
```

Note that the Postprocessor API functionality is only available with JES2 installed. Omit the GPMPPJCL ddname in a JES3 environment.

The complete DDS HTTP API is described in the *z/OS RMF Programmer's Guide*.

DDS options

The preparation of the z/OS component of the DDS host session (GPMSERVE) as server address space for possible exploiters requires the customization of options in a parmlib member that is needed for the GPMSERVE procedure to start the Distributed Data Server.

RMF provides a default parmlib member GPMSRV00, which you may tailor according to your needs. To display the active DDS options, you can use the following command:

```
MODIFY GPMSERVE,OPTIONS
```

Note: A subset of the DDS options is also used by the GPM4CIM component of RMF XP, for example, CACHESLOTS, or HTTP_NOAUTH. For information on the RMF XP setup refer to “How to set up RMF XP” on page 277.

Here is the content of GPMSRV00:

```
/* **** */
/*
/* NAME:          GPMSRV00
/*
/*
/* DESCRIPTION:  PARMLIB MEMBER FOR THE RMF DISTRIBUTED DATA SERVER
/*
/*              HOST ADDRESS SPACE (GPMSERVE)
/*
/*
/* **** */
CACHESLOTS(4)          /* Number of timestamps in CACHE */
DEBUG_LEVEL(0)         /* No informational messages */
SERVERHOST(*)          /* Don't bind to specific IP-Address */
MAXSESSIONS_INET(5)    /* MaxNo RMF PM clients */
SESSION_PORT(8801)     /* TCP/IP port number RMF PM */
TIMEOUT(0)             /* No timeout */
DM_PORT(8802)          /* Port Number for DM requests */
DM_ACCEPHOST(*)        /* Accept from all IP-addresses */
MAXSESSIONS_HTTP(20)   /* MaxNo of concurrent HTTP requests */
HTTP_PORT(8803)        /* Port number for HTTP requests */
HTTP_ALLOW(*)          /* Mask for hosts that are allowed */
HTTP_NOAUTH()          /* No server can access without auth. */
EXCLUDE_REPORTS()      /* Reports to be deactivated */
/* **** */
/* HTTP section via UNIX domain sockets:
/* **** */
MAXSESSIONS_UNIX(1)    /* MaxNo of concurrent HTTP requests */
UNIXSOCKET_PATH(/tmp/gpmserve) /* Pathname for HTTP socket directory */
```

CACHESLOTS

Number of CACHE entries (one for each MINTIME). The valid scope is 3 through 32.

DEBUG_LEVEL

Amount of messages that is sent to the SYSPRINT data set. The valid scope is 0 through 3. DEBUG_LEVEL(0) suppresses all informational messages.

SERVERHOST

TCP/IP address (or hostname) to which the server binds when it opens any listener sockets. You should only use this option, if a host has several TCP/IP addresses (different network adapters) and you want the DDS server to bind its services to one specific TCP/IP address. Make sure, that the value you specify is the valid TCP/IP address (or hostname) of the

host where the DDS server runs. On z/OS, you may use the TSO HOMETEST command to find out the valid TCP/IP addresses.

Example: SERVERHOST(9.164.123.244)

Default: SERVERHOST(*) (any TCP/IP address).

MAXSESSIONS_INET

Maximum number of permitted concurrent RMF PM clients. Additional clients are rejected. The maximum allowed value is 100.

SESSION_PORT

TCP/IP port number for RMF PM clients. It must correspond to the port number, that the clients specify in the SYSPLEX settings.

TIMEOUT

Number of seconds of inactivity, before DDS assumes a timeout condition on the TCP/IP connections for RMF PM clients.

DM_PORT

UDP/IP port number for Tivoli® DM/390 communication.

DM_ACCEPTHOST

TCP/IP address (or hostname) that is allowed to send DM requests. A value of '*' means, that the TCP/IP hosts will not be restricted. You may specify more than one DM_ACCEPTHOST statement.

MAXSESSIONS_HTTP

Maximum number of permitted concurrent HTTP server threads. The maximum allowed value is 100.

HTTP_PORT

TCP/IP port number for HTTP requests.

HTTP_ALLOW

Host names or TCP/IP addresses that can use the HTTP interface. Wildcards * and ? are allowed. You may specify more than one HTTP_ALLOW statement.

Examples:

```
HTTP_ALLOW(*.ibm.com)
HTTP_ALLOW(9.164.*.*)
HTTP_ALLOW(sys?.boeblingen.de.ibm.com)
```

Default: HTTP_ALLOW(*)

Note: If your installation uses a proxy server for http access, you have to specify the hostname of the proxy server to allow access from users which use this proxy server.

HTTP_NOAUTH

Host names or TCP/IP addresses that can use the HTTP interface without authentication (user ID/password). Wildcards * and ? are allowed. You may specify more than one HTTP_NOAUTH statement.

Example: HTTP_NOAUTH(sysa.boeblingen.ibm.com)

Default: HTTP_NOAUTH()

Note:

1. The DDS supports PassTickets (see “Configuring PassTicket support for the Distributed Data Server” on page 19). If you run a CIM server in your z/OS environment, you can use PassTickets to grant this server

access to the DDS. Otherwise, to grant CIM clients access to performance data, you must authorize the CIM server host via HTTP_NOAUTH, since the corresponding CIM RMF monitoring providers use the DDS HTTP API.

2. If your installation uses a proxy server for http access, you have to specify the hostname of the proxy server to allow access from users who use this proxy server.

EXCLUDE_REPORTS

List of Monitor III reports to be deactivated. All reports that are contained in this list will not be provided by the DDS.

You can mix single system and sysplex reports. The listed reports must be separated by a comma. You may specify more than one EXCLUDE_REPORTS statement.

Examples:

```
EXCLUDE_REPORTS(CFOVER,CFSYS)
EXCLUDE_REPORTS(ZFSACT,ZFSSUM)
```

Default: EXCLUDE_REPORTS()

Note: Deactivating reports has impacts on DDS exploiters:

- All metrics based on the deactivated reports will not be provided by the DDS.
- *z/OS Capacity Provisioning* connects to the DDS via the CIM server to obtain performance data from the CPC, SYSINFO, and SYSSUM reports. Therefore, never specify CPC, SYSINFO, or SYSSUM on EXCLUDE_REPORTS() when running *Capacity Provisioning*.

MAXSESSIONS_UNIX

This parameter specifies the maximum number of concurrent HTTP requests to the DDS. The maximum allowed value is 100.

UNIXSOCKET_PATH

This parameter specifies the pathname for the HTTP socket directory. The maximum length of the pathname is twenty characters.

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”
- “Defining SMF synchronization” on page 31

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.
- RMF XP writes record type 104.

Examples:

```
SYS(TYPE(...,72,...))      /* write SMF record type 72 */
SYS(TYPE(...,104(1:12),...)) /* write SMF record 104 with */
                           /* subtypes 1-12 for AIX    */
```

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).

RMF XP

Record type 104 writing requested by setting both of the following two options:

1. Setting the according option TYPE in the active SMFPRMxx parmlib member
2. Option RECORD(YES) in GPM4CIM configuration file specifies record type 104 collection as defined by SMFPRMxx

For detailed information, refer to “How to request SMF record type 104 collection” on page 283.

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, refer to *z/OS MVS Initialization and Tuning Reference*.

- Use the system command SETSMF. For more information, refer to *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 long-term monitoring with RMF XP, you achieve synchronization of SMF record type 104 collection by specifying the INTERVAL parameter in the GPM4CIM configuration file. This parameter defines the length of the monitoring interval and determines the frequency, with which RMF retrieves data from the monitored endpoints.

RMF XP Synchronization for SMF type 104

INTERVAL(length)

Duration of the data monitoring interval in seconds - default:
INTERVAL(300)

For more information see “Configuration files and parameters” on page 278.

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

Storing gatherer options

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

- “Naming parmlib option members”
- “Generalizing parmlib members” on page 34
- “Defining parameters for Monitor I” on page 35
- “Defining parameters for Monitor II” on page 37
- “Defining parameters for Monitor III” on page 39

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 alphanumeric 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 9, “Long-term data gathering with Monitor I,” on page 83, “Details of report commands” on page 188 and Chapter 11, “Short-term data gathering with Monitor III,” on page 109, 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 parmli concatenation in z/OS, it is recommended to define one or more *customer* parmli that can be specified in the LOADnn parmli 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 parmli data set.

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

Generalizing parmli members

In a sysplex environment, each individual system has its own parmli with the corresponding RMF parmli 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, parmli 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 parmli members

Now, you can use the same parmli 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 parmli members, include in the Monitor III parmli 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 parmli 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 36) and an alternative one (“ERBRMF02” on page 37) 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(NOCHDRDR)     /* 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(NOCHDRDR)        /* 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      */
D
PAGING               /* PAGING DATA                      */
NOTRACE              /* NO TRACE REPORT                   */
VSTOR(S)             /* VIRTUAL STORAGE SUMMARY DATA     */
WKLD                 /* WORKLOAD MANAGER DATA            */
D
/*****
/* PART 2: TIMING
/*****
CYCLE(1000)          /* SAMPLE EVERY SECOND (1000 MSEC)   */
NOSTOP               /* ACTIVE UNTIL OPERATOR ISSUES STOP */
SYNC(SMF)            /* USE INTVAL/SYNVAL 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(RVCPUA,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/SYNCCVAL 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") and an alternative one ("ERBRMF03" on page 38) 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”) 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(7))     /* 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(32)             /* 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 */
NOZFS                /* NO ZFS DATA GATHERING */
NOSGSPACE             /* NO STORAGE GROUP SPACE GATHERING */
NOLOCK               /* NO LOCK DATA GATHERING */

```

Considering reporting aspects

- “Setting up the RMF CLISTs”
- “Grouping Monitor II and Monitor III Users” on page 41

Setting up the RMF CLISTs

There are two ways to make the RMF Reporting sessions through the RMF CLIST available to all users in your system: either by SYSPROC concatenation or by a stand-alone CLIST.

SYSPROC concatenation

Concatenate the RMF ISPF dialog library SYS1.SERBCLS to the library associated with file name SYSPROC in your LOGON procedure.

Check the following:

- Make sure all copies of RMF CLISTs from previous releases are deleted from the SYSPROC concatenation. If a CLIST from a previous RMF release is used, the RMF commands will not work.
- Make sure that SYS1.SERBCLS has the same RECFM as the other data sets in the SYSPROC concatenation.
- Make sure the block size for SYS1.SERBCLS is the same or smaller than the block size for the first data set in the SYSPROC concatenation.
- If you changed the name of SYS1.SERBCLS or copied SYS1.SERBCLS into a common dialog library, make sure the new name or common dialog library is associated with SYSPROC.
- If you customized the Monitor III CLISTs by copying the RMF dialog library members into the appropriate common dialog libraries and deleted the ALLOCATE and LIBDEF statements in the RMF CLISTs, make sure that the common dialog libraries are concatenated to the proper file names in your LOGON procedure.

Stand-alone CLIST

Provide a 'stub CLIST' RMFSTART in an established SYSPROC library. The following is a sample stub:

CLIST RMFSTART

```
PROC 0 UTILITY
CONTROL MAIN MSG LIST CONLIST
IF &SYSISPF ^= ACTIVE THEN DO
    ISPSTART CMD(%RMFSTART &UTILITY)
    SET RC = &LASTCC
END
ELSE DO
    ALTLIB ACT APPL(CLIST) DA('SYS1.SERBCLS')
    IF &STR(&UTILITY) = &STR(UTILITY) THEN %ERBRMFU
    ELSE RMF
    SET RC = &LASTCC
    ALTLIB DEACTIVATE APPL(CLIST)
END
```

You can then invoke the RMF Reporter session with the command
%RMFSTART

and you can invoke the Monitor III Report Definition Utility with the command
%RMFSTART UTILITY

This alternative removes the task of copying the RMF CLISTs every time a new RMF release is installed or service is applied to the RMF CLISTs. However, the standard commands to invoke RMF (RMF and ERBRMFU) will not work.

ISPF application ID (APPLID) considerations

When setting up RMF report invocation, consider that RMF uses the following application IDs:

```
ERB
ERBS
ERBU
ERB2
```

Do not use these IDs in a NEWAPPL(id) keyword of an RMF invocation command.

Grouping Monitor II and Monitor III Users

Because both online monitors provide a timer-driven automatic update mode (T command in Monitor II and GO mode in Monitor III), response time for such a user can be meaningless. For example, when Monitor III is running in GO mode with a refresh period of 100 seconds, the TSO response time appears as 100 seconds. Consequently, the response time measured for the service class that the user is in may not be an accurate representation of what happened during the interval.

To avoid this situation, RMF recommends that Monitor II and Monitor III users are put in a separate service class.

Installing workstation functions

In addition to the monitoring functions on the host system, RMF extends its monitoring capabilities by several functions that are available on the workstation. All required programs and procedures are automatically installed in the RMF distribution libraries on your host system through SMP/E. Each user can then perform the installation on the workstation according to the provided descriptions.

Installing the RMF Postprocessor XML Toolkit

Download the RMF Postprocessor XML Toolkit to your workstation.

For instructions on downloading and using the provided functions, refer to “Viewing XML reports with the RMF Postprocessor XML Toolkit” on page 289.

Installing the RMF Spreadsheet Reporter

Download the RMF Spreadsheet Reporter to your workstation.

For instructions on downloading, as well as activation and usage of the Spreadsheet Reporter functions, please refer to chapter Chapter 18, “RMF Spreadsheet Reporter,” on page 293.

Installing RMF Performance Monitoring (RMF PM)

The installation of RMF PM comprises the following steps.

1. Tailoring of parmlib member GPMSRVxx (“Setting up the Distributed Data Server for z/OS” on page 26)
2. Start of the Distributed Data Server (“Starting the Distributed Data Server” on page 53)
3. Installation of RMF PM (“Getting started - Installation and setup” on page 342)

Running RMF PM requires an active TCP/IP connection between the host system and the workstation.

Installing RMF client/server enabling - RMFCS

RMFCS is designed to allow several users to monitor the MVS system individually. Each user who wants to run this function just has to initialize the personal environment by taking the following steps:

1. Customize ISPF C/S session
2. Customize RMFCS procedures
3. Ensure RACF authorization
4. Initialize message-initiated monitoring

5. Initialize exception-initiated monitoring

These steps are described in detail in “Installation and startup of RMFCS components” on page 378.

Installing the z/OSMF Resource Monitoring plug-in

z/OSMF provides the *Resource Monitoring* application as a plug-in. To exploit this application, you need to install and configure z/OSMF accordingly. For instructions on how to achieve this, refer to *IBM z/OS Management Facility Configuration Guide*.

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 5, "Data Gathering Reference," on page 81 and Part 6, "Reporting Reference," on page 127.

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 47)
- specifying the Monitor I gatherer session options (see Chapter 9, “Long-term data gathering with Monitor I,” on page 83)

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” on page 53)

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 59.

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 24).

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 45.

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 51) 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 56. 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 8). 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 22). 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 */  
SMFBUF(RECTYPE(70:79,104(60:64))) /* for Windows on System x */
```

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 find details of the options and commands for all RMF sessions in Part 5, “Data Gathering Reference,” on page 81 and Part 6, “Reporting Reference,” on page 127.

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 22 for more details).
- The RMF **Parmlib** member, or other equivalent data set member containing session options. See “Storing gatherer options” on page 33 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 59.

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 5, “Data Gathering Reference,” on page 81 and Part 6, “Reporting Reference,” on page 127.

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 24. 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 53.

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 59 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 9, "Long-term data gathering with Monitor I," on page 83.

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" on page 188.

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 24). 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 11, “Short-term data gathering with Monitor III,” on page 109.

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 the Distributed Data Server

The Distributed Data Server (DDS) provides the ability to serve multiple clients in a single-server address space. This capability is used, for example, by RMF PM.

To have the Distributed Data Server address space managed by RMF, you can start it automatically by using the DDS option. This option ensures that you always have one active instance of the Distributed Data Server within your sysplex. As soon as the RMF Sysplex Data Server recognizes the DDS option on any system in the sysplex, the Distributed Data Server is started on the RMF master system. The master system is a system with an active Monitor III gatherer, the highest z/OS release and SMFBUF option set. The MASTER parmlib option can be used to define a system to be a master system candidate. If another system becomes the master system, the Distributed Data Server is automatically restarted on this system.

You can specify the DDS option as follows:

PARM='DDS'

to activate the sysplex-wide DDS management with the EXEC statement of the RMF procedure (see “Customizing the RMF control session” on page 22).

S RMF,,,DDS

to activate the sysplex-wide DDS management with the RMF start command.

F RMF,DDS

to activate the sysplex-wide DDS management with the RMF modify command.

You can stop the sysplex-wide DDS management using the **NODDS** option. For example, the command F RMF,NODDS also shuts down the current sysplex-wide DDS instance, regardless on which system the command has been entered.

Example for manual startup of GPMSERVE:: To start the Distributed Data Server manually, assuming that you have stored the corresponding parameters in Parmlib member GPMSRV01, issue the command:

```
S GPMSERVE,MEMBER=01
```

If you use the default member GPMSRV00, you can omit the MEMBER parameter.

To retrieve information about an automatically started Distributed Data Server, you can use the following commands:

```
MODIFY RMFDDS01,OPTIONS shows the currently active GPMSERVE options.
```

```
MODIFY RMFDDS01,DISPLAY shows the currently active connections.
```

Note:

1. If you started the Distributed Data Server address space manually, use GPMSERVE instead of RMFDDS01 as identifier in the commands.
2. You must restart the DDS to ensure correct operation when the local time changed (daylight saving time).

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 find details about RMFCS in Chapter 21, “RMF Client/Server enabling (RMFCS),” on page 373.

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 59.

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:

```
{MODIFY} RMF, {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 47.
- By specifying a time value in the STOP option for a specific session. See Part 5, “Data Gathering Reference,” on page 81.
- By issuing a STOP session command to stop a specific session. All other active sessions continue processing. See “Stopping a specific session.”

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 an 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 55.
- 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. The following is the final option list, including the source for each option.

WKLD -- DEFAULT	DEVICE(DASD) -- DEFAULT
SYNC(SMF) -- DEFAULT	DEVICE(TAPE) -- MEMBER
NOTRACE -- MEMBER	DEVICE(NOUNITR) -- MEMBER
NOEXITS -- MEMBER	CHAN -- DEFAULT
NOENQ -- DEFAULT	PAGING -- DEFAULT
OPTIONS -- MEMBER	CPU -- DEFAULT
STOP(8H) -- DEFAULT	CACHE -- DEFAULT
CYCLE(1000) -- COMMAND	MEMBER(02) -- EXEC
RECORD -- DEFAULT	MEMBER(10) -- COMMAND
SYSOUT(A) -- DEFAULT	IOQ(NOCHDR) -- DEFAULT
NOREPORT -- DEFAULT	IOQ(NOCOMM) -- DEFAULT
NOPAGESP -- MEMBER	IOQ(DASD) -- DEFAULT
DEVICE(NOCHDR) -- DEFAULT	IOQ(NOGRAPH) -- DEFAULT
DEVICE(NONMBR) -- DEFAULT	IOQ(NOTAPE) -- DEFAULT
DEVICE(NOGRAPH) -- DEFAULT	IOQ(NOUNITR) -- DEFAULT
DEVICE(NOSG) -- DEFAULT	VSTOR(S) -- DEFAULT
DEVICE(COMM) -- COMMAND	

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 under “Details of report commands” on page 188. 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. Performance management

RMF offers you a wide variety of views on z/OS system performance. This part of the manual tells you which view will help you most in a particular situation, and what steps you can take to have RMF present you with this view. For help in analyzing the reports produced, see *z/OS RMF Report Analysis*.

There are many situations in which RMF can help you with performance management. You needn't wait until you have a problem in that area; RMF supplies data that you can use to check that things are running smoothly, or to see in good time where improvements may be necessary. The following chapters discuss how RMF helps you with:

- Performance monitoring, seeing that everything is running smoothly
- Performance analysis, getting to the seat of problems
- System tuning, ensuring the best usage of resources
- Capacity planning, ensuring that you have enough resources

Chapter 6. Performance administration

This information unit provides information about planning and preparing for performance management in your sysplex:

- what you should know about z/OS workload management
- how the hierarchical performance view looks

What is performance administration?

Is it a separate task, or is it part of monitoring and analysis?

The answer to these questions depends to a very high degree on the organization and size of your installation. If you have a group of system programmers and performance analysts, then you can assign different performance management tasks to different people. If, on the other hand, you are the only specialist, and responsible for everything, then you might see no need to distinguish the different tasks.

In the following considerations we will concentrate on the task itself without taking into consideration whether it is performed by the same specialist as the monitoring and analysis tasks.

As you see in Part 2, “Administration,” on page 13, we have defined performance administration as the task of setting up everything required for the smoothest possible running of performance measurement and performance management.

Defining procedures and parameters

The operator will start all the data-gathering functions that are performed by the three monitors. Monitor I and Monitor III will probably run continuously, while the Monitor II background session might be started on request only. But, in each case, the START command should be as easy as possible for the operator.

Ease of operation is important also with respect to the automatic start-up procedures in most installations. Here, you do not want to have the operator typing in commands with many parameters or replying to many requests from the application. We therefore recommended you to define all start parameters and options in such a way that the appropriate values are selected by default. Only in exceptional cases should the operator have to override these values to handle a specific situation.

Of course, setting up all gathering options requires a common understanding with the system programmers who work with the data. They have to decide what data is to be gathered. Do they need data for performance monitoring and analysis, or for capacity planning and tuning? Different tasks might need different data, and you have to implement appropriate gathering procedures for each.

When defining the scope of data to be gathered, you also have to specify where to store the data. As you know, the three monitors create two types of data:

- SMF records (Monitor I, Monitor II, Monitor III)
- VSAM records (Monitor III only)

It is part of the administration task to allocate the appropriate data sets to ensure that the performance analysts have access to everything they need; not only to data from today or yesterday, but also to data gathered some time ago.

Setting performance goals

The human view of the performance of a system is often subjective, emotional and difficult to manage. However, the purpose of a system is to meet the business needs of the users.

To match business needs with subjective perception, the concept of *Service Level Agreements* (SLA) was introduced.

The SLA is a contract that objectively describes such measurables as:

- Average transaction response time for network, I/O, CPU, or total
- The distribution of these response times (for example, 90% TSO trivial at less than 0.2 of a second)
- Transaction volumes
- System availability

A *transaction* is a business unit of work and can be a CICS end user interaction or a batch job, for example. Ideally, a transaction is defined from a user's point of view.

The definition and implementation of an SLA might be done in your installation in a more or less formal way, but the more precisely

- the expectations of the users
- the capabilities of the computer shop

have been defined, the easier tracking and monitoring are. This definition is important with regard to the capabilities of performance management in a z/OS system. There, the *Workload Manager* enables you to specify explicit performance goals for your applications, and the reporting capabilities within RMF will allow you to track them directly.

z/OS workload management

z/OS workload management provides a solution for managing workload distribution, workload balancing, and distribution of resources to competing workloads. z/OS workload management is the cooperation of various subsystems (for example, CICS, DB2®, IMS™, JES, APPC, TSO/E, UNIX System Services) with the z/OS workload manager (WLM) component.

Fewer, simpler, and consistent system externals: Workload management provides a way to define z/OS externals and tune z/OS without having to specify low-level parameters. The focus is on setting performance goals for work, and letting the workload manager handle processing to meet the goals.

Externals reflect customer expectations: Workload management provides z/OS performance management externals in a service policy that reflects goals for work, expressed in terms commonly used in service level agreements (SLA). Because the terms are similar to those commonly used in an SLA, you can communicate with users, with business partners, and with z/OS, using the same terminology.

Service definition

Performance administration is the process of defining and adjusting performance goals. Workload management introduces the role of the service level administrator. The service level administrator is responsible for defining the installation's performance goals on the basis of business needs and current performance. This explicit definition of workloads and performance goals is called a *service definition*. The service definition applies to all types of work, including CICS, IMS, TSO/E, UNIX System Services, JES, and APPC/MVS. You can specify goals for all z/OS-managed work, whether online transactions or batch jobs, and the goals apply to the sysplex.

Workload management concepts

The service definition contains all information about the installation needed for workload management processing. There is one service definition for the entire sysplex. The service level administrator sets up *policies* within the service definition to specify the goals for work. He must understand how to organize work, and be able to assign performance objectives to it.

A *service definition* consists of:

- One or more *service policies*, which are a named set of performance goals that an installation tries to meet. You can have different policies to specify goals intended for different times. Service policies are activated by an operator command, or through the ISPF administrative application utility function.
- *Workloads* and *service classes*, which are the categories of work. A workload is a grouping of work in a way that is meaningful for your installation to manage and monitor. It is made up of a group of service classes. You assign performance goals and, optionally, capacity boundaries, to service classes.

In addition, you can define *report classes* which will help you in your reporting with another granularity as being possible with service classes.

The term *workload group* is also used in RMF documents and means the same as *workload*.

- *Resource groups*, which define processing capacity boundaries across the sysplex. You can assign a minimum and maximum number of CPU service units per second to work by assigning a resource group to a service class.
- *Classification rules*, which determine how to assign incoming work to a service class.

Workloads and service classes

To workload management, work is a demand for service, such as a batch job, an APPC, CICS, or IMS transaction, a TSO/E logon, or a TSO/E command. All work running in the installation is divided into workloads. Your installation may already have a concept of workload. A workload is a group of work that is meaningful for an installation to manage and monitor. For example, all the work created by a development group could be a workload, or all the work started by an application, or in a subsystem.

Within a workload, you group work with similar performance characteristics into service classes. You create a service class for a group of work with similar:

- Performance goals
- Resource requirements
- Business importance

You can create a service class for any combination of the above. You assign performance goals to the service classes, such as a response time goal, and you indicate how important it is to your business that the performance goal be achieved.

Performance goals

There are three kinds of goal:

- Response time
- Velocity
- Discretionary

Response time goals indicate how quickly you want your work to be processed. Velocity goals are for kinds of work for which response time goals are not appropriate, such as long-running batch jobs.

Response time

This is the expected amount of time required to complete the work submitted under the service class. Specify either an average response time, or response time with a percentile. A percentile is the percentage of work in that performance period that should complete within the response time.

You must specify the goal for system response time, not “end-to-end” response time. That is, workload management does not control all aspects of system performance, so response-time scope is confined to the time SRM has control of the work.

Velocity

This is a measure of how fast work should run when ready, without being delayed for resources. Velocity is a percentage from 1 to 99. The formula for velocity is:

$$\text{Velocity} = \frac{\text{\# Using Samples}}{\text{\# Using Samples} + \text{\# Delay Samples}} * 100$$

Please, refer to “Workflow and velocity” on page 75 for details and the difference to the term *Workflow*.

Discretionary

Workload management defined goal. Associate this goal with work for which you do not have a specific performance goal. Work with a discretionary goal is run when resources are available.

Importance

Importance is a reflection of how important it is to achieve the service-class goal. Workload management uses importance only when work is not meeting its goal. It is required for all goal types except discretionary. Importance applies on a performance-period level and you can change importance from period to period. There are five levels of importance: 1 to 5, 1 being the highest importance.

Service class periods

A service class with a goal and an importance is called a service class period.

Reporting hierarchy

There are different scopes for the reports available in RMF, this can be explained with the reporting hierarchy that RMF offers.

Sysplex view

Comprehensive reports with summary and overview data about the sysplex. You can get these reports either interactively from Monitor III, or as Postprocessor reports based on measurement data gathered by Monitor I or Monitor III.

System view

Reports that provide information for one selected system. This may be a stand-alone system or a member of a sysplex. Single-system reports offer a more detailed level of performance information.

All reports that are not explicitly related to one specific address space or system resource are called system reports. In contrast, job or resource reports concentrate on one specific component of your system.

Job or resource view

This is the deepest level of detail. It concentrates on single jobs (or, more precisely, address spaces) or single-system resources, and helps you analyze a performance problem that is indicated by a sysplex or system report.

Chapter 7. Performance monitoring

This information unit describes what to do on a daily basis to keep your finger on the pulse of the system. Thus you may avoid being surprised by performance degradation caused by gradually-changing factors. The task of performance monitoring involves:

- watching performance goals
- observing response times
- monitoring throughput
- observing bottlenecks and exceptions

Watching performance goals

You have set performance goals for your workload within a sysplex by means of WLM, and you should observe if and how well these goals are being met. You can do this for your whole workload at one glance, independent of the single system your work is actually running on.

RMF offers two reports that provide information about performance goals, as defined in the active performance policy, and the corresponding actual values.

Monitor III - Sysplex Summary Report

Use this report as the entry point for this kind of monitoring. It shows all active workloads with their performance values, including the goals for each service class period.

The performance status line offers a very easy way of monitoring the performance of your sysplex. It is displayed when continuous monitoring is active (in GO mode), and shows you the performance history of your system for the past two hours at a glance.

Postprocessor - Workload Activity Report

Use this report if you want to evaluate the attainment of performance goals for time intervals in the past. This sysplex report shows performance goals and actual values at different levels of detail (from policy summaries to service class period details). You can select the type of information that best meets your installation's requirements.

Observing response times

There might be times where you are interested in monitoring response times for single users and groups of users. This may be the case if you have defined service level agreements based on response times, or if you get complaints from users about slow response times on the system. For this task, you will call the

- *Monitor III - Sysplex Summary Report*: This report displays the average response time for each service and report class period.

If you would like more information about one specific service class period, use the cursor-sensitive control of the Monitor III to navigate directly to the

- **Monitor III - Response Time Distribution Report:** With this report, you get detailed data for the service class period you are interested in:
 - If you have defined a response time goal for a service class period, you see a response time distribution graphic which is the lowest level of detail that is available.
 - In any case, you see the response time for each system from which this period is getting service.

If you want to concentrate on a single system, you get response time data from the

- **Monitor III - System Information Report:** This report shows the average response time for all workloads and service classes.

You get more detailed data for one group with the

- **Monitor III - Group Response Time Report:** The average response time is displayed as time that is split up into *using* and *delayed* time frames, so you can see how much time this group of address spaces was using the system resources, and how much time was spent waiting for resources.

If you need response time data for longer periods of time, then you will work with Monitor I data. Get the information by creating a

- **Postprocessor - Workload Activity Report:** This sysplex report provides response times for all service class periods and — if you have defined response time goals — response time distribution information. You select the level of detail by the corresponding report options.

Monitoring throughput

There are various indicators that show throughput values.

Transaction rate

Throughput definition is:

$$\text{Throughput} = \frac{\text{\# transactions or jobs}}{\text{time}}$$

You get this information in various reports.

- **Monitor II Address Space SRM Data Report:** Data is available for all address spaces: you get the transaction count and the total session time for each address space.
 - **Monitor III Reports:** You can get the transaction rate on sysplex level with the
 - Sysplex Summary report.
You get summary statistics for all workloads, service classes, and service class periods.
 - Response Time Distribution report and Work Manager Delay report
The transaction rate is shown for one selected service class period.
- Throughput data on system level are available with the
- System Information report.
The values are summarized by high-level groups (system, TSO, UNIX System Services ...) and for all workloads and service classes.
 - Group Response Time report.
This report shows detailed values for one selected service class period.

- **Postprocessor Reports:**
 - Workload Activity report.
On sysplex level, you get transaction rates for all workloads and service classes.
 - Exception report.
You can define exception criteria for transaction rates of specific workloads or service classes.

Workflow and velocity

Another way to characterize the throughput could be to take the workflow or velocity data that is shown in several reports:

- *Workflow* is a term created by Monitor III for *reporting* purposes.
- *Velocity* is a term created by workload management for *managing* purposes.

For both terms, the definition has the same formula:

$$\text{Value} = \frac{\text{\# Using Samples}}{\text{\# Using Samples} + \text{\# Delay Samples}} * 100$$

What is the difference between workflow and velocity?

Velocity (as a managing indicator) considers only the processor, the storage, and DASD devices — these are the resources which are under control of workload management.

Workflow (as a reporting indicator) reflects all system components (for example, tape activities or delays caused by mounts or HSM).

Monitor III Reports: Workflow data are shown primarily in the

- Workflow/Exceptions report
- System Information report
- Group Response Time report
- Sysplex Summary report: here you find velocity data

Postprocessor Reports: The

- The Workload Activity report

shows the velocity values.

Observing bottlenecks and exceptions

There are two approaches to monitoring performance:

- You can check the performance of your system by observing indicators like *performance goals*, *performance index*, *workflow*, or *response times*, most of which have already been discussed.
- Or you can use a variety of Monitor III reports to look for *exceptions* or *delays* that might be the source of a performance problem.

To define your performance exceptions directly, use the

- **Monitor III - Workflow/Exceptions Report:** There are many types of exceptions you can specify: CPU utilization, response times, number of users, storage activities and many others.

You define thresholds and corresponding colors, and Monitor III indicates when a threshold has been reached.

Now, you can start investigating the reason, which hopefully will enable you to solve the problem either immediately, or with the next IPL, or with planning for a more powerful processor in the near future.

If you would prefer information about exceptions for a longer time range, you will call the

- ***Postprocessor - Exception and Overview Reports:*** For these single-system reports, you can define many types of exceptions or thresholds, on the basis of CPU, I/O, workload or paging data gathered by Monitor I. The reports list all relevant data and allow you to create the detailed interval reports.

The other method is to look directly for delays — situations in which jobs are waiting for resources (for example, processor, devices, storage).

- ***Monitor III - System Information Report:*** This report gives you an overview of all applications in your system at different levels (system, TSO, batch, and so on) or grouped by workloads or service classes. The information includes speed and delay indicators.

If you like to create your own performance reports that should contain the data you are interested in, you can do this with RMF PM.

- ***RMF PM - PerfDesks:*** RMF PM takes its input data from Monitor III. The data is suitable for monitoring and analyzing performance in real time and in the recent past. It provides a selected subset of the information provided by the Monitor III gatherer: general performance data, performance data for jobs, and workload-related performance data.

You can collect real-time data, combine data from different collection types, or even from different applications, and group resources together. Once you have created these scenarios, you can save them in your own panels, called *PerfDesks*.

Chapter 8. Performance analysis

Monitoring your systems is an ongoing process. Analyzing performance problems is a task that needs to be performed only from time to time — but the full capability of RMF will help you here, too.

This information unit:

- discusses some ideas about performance problems
- points to reports that can help you analyze problems and evaluate possible solutions

What is a performance problem?

There are many views on what constitutes a performance problem. Most of them revolve around unacceptably slow response times or high resource usage, which we can collectively refer to as "pain." The need for performance investigation and analysis is indicated by, for example:

- Slow or erratic response time:
 - Service level objectives being exceeded
 - Users complaining about slow response
 - Unexpected changes in response times or resource utilizations
- Other indicators showing stress;
 - Monitor III Workflow/Exceptions
 - System resource indicators (for example, paging rates, DASD response)
 - Expected throughput on the system not being attained

Finally you need to decide whether a given situation is a problem worth pursuing or not. This decision is based on your own experience, knowledge of your system, and sometimes politics. We simply assume for the following discussions that you are trying to relieve some sort of numerically quantifiable performance problems in your system.

Generally, a performance problem is the result of some workload not getting the resources it needs to complete in time. Or the resource is obtained but is not fast enough to provide the desired response time.

The most frequent cause of performance problems is having several address spaces compete for the same resource. These could be a hardware resource (processor, device, storage) or serially usable software resource (catalog, VTOC). While one address space is using the resource, the other address spaces are delayed. Therefore, one key aspect of Monitor III is to make visible who is using what, and who is delayed.

Reports that provide data for analysis

Several Monitor III reports provide this information on different levels of detail.

Monitor III Reports

- Delay report

This report shows the address spaces that have the highest delays in your system.

- Job report

An address space can be delayed for many reasons. Several variations of this report (for example, PROC, DEV, JES, HSM, OPER) provide detailed information for your analyzing process.

- Resource reports

You can also analyze selected resources and see who is using and who is waiting for these resources.

- Storage reports

There are several types of storage reports that provide very detailed information about storage consumption (paging, migration, frames available, ...) and utilization of common storage (CSA, SQA, ECSA, and ESQA).

- Work Manager Delay report

This sysplex report provides information for your CICS and IMS subsystems and shows several types of delay that might be the source of a current performance problem.

Postprocessor Reports

The reason for a performance problem can also be the overloading of resources in your systems, for example, of the processors, channels or devices. Here, you will find the best overview in the Postprocessor reports that are based on Monitor I data:

Some of the reports you might use are the:

- CPU Activity report
- Coupling Facility Activity report
- Channel Activity report
- Cache Subsystem Activity report
- Device Activity report
- Paging Activity report
- Virtual Storage Activity report

These are long-term reports showing intervals, typically of 15 or 30 minutes, that you have defined with your gathering options.

Monitor II Reports

If you are interested in some snapshot data to analyze the current status of your system, you can get them from Monitor II.

You can get information about the utilization of the different resources from these reports:

- Central Storage/Processor/SRM Activity report
- Channel Path Activity report
- Device Activity report
- Paging Activity report

Other Monitor II Reports

If you see overloading of resources because of:

- temporary or permanent workload peaks
- single applications that dominate others

you might be interested in evaluating your performance inhibitors.

If you need a snapshot of the current system status, you can get this with other Monitor II reports.

Here you will find data about resource consumption (CPU time, I/O rates, storage utilization) for all address spaces:

- Address Space Resource Data report
- Address Space SRM Data report

If you want to concentrate on one specific address space, you can tailor the report accordingly as a jobname report to monitor only one selected job in your system.

Postprocessor Workload Activity Report

The other possibility is to get resource data for longer time frames either as interval (for example, 30 minutes) or duration (up to 100 hours) report with the Postprocessor Workload Activity Report.

This report provides resource data for different levels of detail. In addition to system control indicators such as service classes and workloads, you might also specify selected report classes to get the data reported according to your requirements and needs.

Postprocessor Trace Activity Report

In specific situations, it might be necessary to get more detailed data. In this case, start Monitor I with some trace options.

You can get many types of data about the utilization of different system components and various activities in the system that can help you in solving a complex performance problem.

Spreadsheet Reporter Macros

Based on Postprocessor data, you can perform the analysis of your system on the workstation using the Spreadsheet Reporter. You select the time range and scope of the data that is important to understand your system, you create records and download them to the workstation. There, you can use several spreadsheet macros which will provide you summary and detail reports for your key system components (processor, storage, DASD and Cache subsystem) and for your important workloads.

Depending on the spreadsheet application (either Excel or Lotus® 1-2-3®), a different set of macros is available, for example:

- System Overview Report
- Workload Overview Report
- I/O Subsystem Report
- Cache Subsystem Overview Report

- Coupling Facility Trend Report

Part 5. 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 9. 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 52)
- 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 22)
- the RMF Monitor I parmlib member ERBRMF00 (see “ERBRMF00” on page 36)

RMF merges the input to a final set of options for the session. See Chapter 5, “How RMF processes session options,” on page 59 for details.

Summary of session options

Monitor I creates SMF records type 70 — 78, you find an overview in “Activity monitoring” on page 10.

Table 4 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 4. Summary of Monitor I Session Options

Option	Description	Details on
Activity Options		
<u>CACHE</u> / <u>NOCACHE</u>	Cache activity	“CACHE” on page 85
<u>CHAN</u> / <u>NOCHAN</u>	Channel path activity	“CHAN” on page 85
<u>CPU</u> / <u>NOCPU</u>	Processor activity	“CPU” on page 86
<u>CRYPTO</u> / <u>NOCRYPTO</u>	Cryptographic hardware activity	“CRYPTO” on page 86
<u>DEVICE</u> (type)/ <u>NODEVICE</u>	Device activity	“DEVICE” on page 86
{(SUMMARY)} ENQ{(DETAIL[,majorname[,minorname]])}/ <u>NOENQ</u>	Enqueue contention activity	“ENQ” on page 88
<u>ESS</u> [(opt_list)]/ <u>NOESS</u>	Enterprise Disk System statistics	“ESS” on page 89
<u>FCD</u> / <u>NOFCD</u>	FICON® director activity	“FCD” on page 90
<u>IOQ</u> (opt_list)/ <u>NOIOQ</u>	I/O queuing activity	“IOQ” on page 91
<u>PAGESP</u> / <u>NOPAGESP</u>	Page data set activity	“PAGESP” on page 95
<u>PAGING</u> / <u>NOPAGING</u>	System paging activity	“PAGING” on page 95
<u>TRACE</u> (variable[,opt_list)]/ <u>NOTRACE</u>	Trace variables for the Trace Activity report	“TRACE” on page 97
{(S)} <u>VSTOR</u> {(D [,jobname1,jobname2,...])}/ <u>NOVSTOR</u>	Virtual storage activity	“VSTOR” on page 103
<u>WKLD</u> / <u>NOWKLD</u>	Workload activity	“WKLD” on page 104

Table 4. Summary of Monitor I Session Options (continued)

Option	Description	Details on
VMGUEST/NOVMGUEST	Processor activity	"VMGUEST" on page 104
Time-frame Options		
{(1000)} CYCLE{(nnn)}	The length of the cycle at the end of which RMF makes sampling observations	"CYCLE" on page 86
{(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 90
[M] STOP (value[H])/NOSTOP	Desired duration of the Monitor I session, in minutes (M), or hours (H)	"STOP" on page 96
{(SMF)} {(RMF,mm)} SYNC {(RMF,mmM)}/NOSYNC	Interval synchronization with the SMF or the RMF interval synchronization on the minute	"SYNC" on page 96
Reporting Options		
{OPTIONS}/{NOOPTIONS} {OPTN } {NOOPTN }	Option list for the session to be printed at the operator console	"OPTIONS" on page 93
RECORD/NORECORD	Specifies whether measured data is to be written to SMF records	"RECORD" on page 95
{(REALTIME)} REPORT{(DEFER)}/NOREPORT	Specifies production of printed interval reports of measured data	"REPORT" on page 95
SYSOUT(class)	SYSOUT class to which the formatted printed reports are directed	"SYSOUT" on page 97
Environment Options		
EXITS/NOEXITS	User exit routines to be executed during session processing to gather or report additional data	"EXITS" on page 89
MEMBER(list)	Parmlib member containing Monitor I session options	"MEMBER" on page 92

Default session options

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

Table 5. 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
DEVICE(DASD)	Measures DASD activity (not other classes of device)
NOESS	Does not measure Enterprise Disk System statistics
NOFCD	Does not measure FICON director activity
IOQ(DASD)	Measures I/O queuing activity on logical control units for DASD
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
CYCLE(1000)	Takes sample measurements once a second (1000 milliseconds)
INTERVAL(30M)	Combines data every 30 minutes (value is ignored for SYNC(SMF))

Table 5. Monitor I Default Session Options (continued)

Default Option	Description
STOP(8H)	Ends the session 8 hours after it was started
SYNC(SMF)	Synchronizes the reporting interval with 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)
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.

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 34, 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 104 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:

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

NMBR requests specific device numbers, where **aaaa**, **bbbb**, and **zzzz** each represent hexadecimal device numbers. You can omit leading zeros. You can specify any combination of a single device number, in the format **aaaa**, a list of device 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, when specified, cancels any existing list of device numbers.

- Any of the following classes:

CHRDR/NOCHRD

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.


Note: For more information on non-DASD/TAPE measurement, refer to the CMB parameter in IEASYS in *z/OS MVS Initialization and Tuning Reference*.

- If you request the following:

```
DEVICE(NODASD, NMBR(0288,0291), 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 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 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 30 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 22 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 device numbers, where aaaa, bbbb, and zzzz each represent hexadecimal device numbers. You can omit leading zeros. You can specify any combination of a single device number, in the format aaaa, a list of device 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 alphanumeric 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 33. 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 22.

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 59.

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 94 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 6 on page 94 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 6. Where to specify Monitor I options

Source	Where Option is specified
-- COMMAND	On a START or MODIFY command.
-- 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 30 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

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

MCVMGAGE
expanded storage migration age

MCVSBTF
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.

OMDGOEB
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
RMCASSC	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
<u>ALL</u>	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 10. 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 Chapter 14, “Snapshot reporting with Monitor II,” on page 167.

Chapter 11. 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 33 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 111.
- The **parm** field of the START session command that starts the session. See “Starting a specific Monitor” on page 52.

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 55.
- The response to the OPTIONS option.

Table 7 gives a summary of the Monitor III gatherer session options. The referenced pages describe the options in detail.

Table 7. Monitor III Data Gatherer Session Options

Option	Effect	Details on
CACHE(suboption...)	Defines cache data gathering.	“CACHE” on page 111
CFDETAIL	Defines level of detail for data gathering for the coupling facility.	“CFDETAIL” on page 111
CYCLE(nnnn)	Sets the length of the cycle at the end of which RMF samples data.	“CYCLE” on page 112
DATASET(suboption...)	Controls data set recording of sampled data.	“DATASET” on page 112
HFSNAME(suboption...)	Controls data set recording for z/OS UNIX file systems.	“HFSNAME” on page 113
IOSUB	Controls data set recording of I/O-subsystem and channel-path activity.	“IOSUB” on page 113
LOCK	Defines data gathering for lock reporting (spin locks and suspend locks).	“LOCK” on page 113
MASTER	Makes the system eligible/uneligible to be the RMF Master Gatherer system.	“MASTER” on page 114
MEMBER(list)	Specifies Parmlib members containing session options.	“MEMBER” on page 114
MINTIME(nnn)	Specifies the interval at which data samples are summarized.	“MINTIME” on page 115
OPD	Defines data gathering for OMVS process data.	“OPD” on page 115
OPTIONS	Controls display of the current options at the start of a session.	“OPTIONS” on page 115
RESOURCE(...)	Specifies the job entry subsystem (JES) to be used.	“RESOURCE” on page 116

Table 7. Monitor III Data Gatherer Session Options (continued)

Option	Effect	Details on
SGSPACE(suboption...)	Defines data gathering for storage group space and disk space monitoring.	"SGSPACE" on page 117
STOP(value)	Sets the duration of the data gatherer interval.	"STOP" on page 117
SYNC	Synchronizes MINTIME within the hour.	"SYNC" on page 118
SYSOUT(class)	Specifies the SYSOUT class for gatherer messages.	"SYSOUT" on page 118
VSAMRLS(suboption...)	Controls data gathering for VSAM RLS activity.	"SYSOUT" on page 118
WSTOR	Sets the size of the RMF local storage buffer.	"WSTOR" on page 119
ZFS	Defines data gathering for monitoring zFS activity.	"ZFS" on page 120
ZIIPUSE	Specifies whether the Monitor III data gatherer is entitled to execute partially on IBM System z Integrated Information Processors (zIIPs).	"ZIIPUSE" on page 120

Default gatherer session options

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

Table 8. Monitor III Default Session Options

Default Option	Description
CACHE	Defines cache data gathering.
CFDETAIL	Defines partial data gathering for the coupling facility.
CYCLE(1000)	Takes data samples once a second (1000 milliseconds).
DATASET(STOP,NOSWITCH)	No data set recording will be done.
IOSUB	Defines data gathering for the I/O subsystem and for channels.
MASTER	Makes the system eligible to be the RMF Master Gatherer system.
NOLOCK	No data gathering for lock reporting (spin locks and suspend locks).
MINTIME(100)	Builds a set of samples every 100 seconds.
OPD	Defines data gathering for OMVS process data.
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.
RESOURCE(*JES2,JES2)	Assumes that JES2 is installed on the system.
NOSTOP	The session does not stop automatically after a predefined time; you must use a STOP command.
VSAMRLS	Activity data is gathered for VSAM RLS by storage class.
SYNC(0M)	MINTIME is synchronized on the hour.
WSTOR(32)	Sets the RMF local storage buffer to 32 megabytes.
NOZFS	No data gathering about zFS activity.
ZIIPUSE	Specifies whether the Monitor III data gatherer is entitled to execute partially on IBM System 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 via the DATASET option) into VSAM data sets.

SMF records

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

You find detailed information about all record types in “Activity monitoring” on page 10.

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 34, 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 114).

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 120.

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 through 3 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 55.

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 IEF RDER DD statement in the RMF cataloged procedure. See “Setting up the RMF control session including Monitor I and Monitor II” on page 22.

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 30 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 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 120.)

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 9. 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.
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(7)) -- 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(7)) -- DEFAULT
ERB305I   III : WSTOR(32) -- 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 : NOZFS -- DEFAULT

```

Figure 4. Console Output with OPTIONS in Effect

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).

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.

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.


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 30 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 32 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 data gathering about zFS file system activity.


ZIIPUSE




Specifies whether the Monitor III data gatherer is entitled to execute partially on IBM System 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 7 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 119) 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(7)) -- 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(32) -- DEFAULT
```

For more information, see Chapter 5, “How RMF processes session options,” on page 59.

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” on page 134.

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 55.

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 7 to 5 megabytes.

The following command modifies the options:

```
MODIFY RMF,MODIFY III,DS(SWITCH),DS(DEL(RMF.DS04)),DS(WHOLD(5))
```

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 five 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, please see "Data set allocation" on page 133.

Part 6. Reporting Reference

This part deals with the RMF reporting capabilities, and how to control them. Reports are available to help you with three different tasks:

- Interactive performance analysis, using the Monitor III Reporter Dialog
- Snapshot reporting, using the Monitor II Display Session, with the option of producing reports in printed form
- Long-term overview reporting, using the Postprocessor

In addition, Chapter 16, “Cross platform monitoring with RMF XP,” on page 275 describes how to set up, configure and start RMF XP if you want to monitor the performance of heterogeneous environments running the following operating systems:

- AIX on System p
- Linux on System x
- Linux on System z
- Windows on System x

Chapter 12. The online reporting environment

A common ISPF interface gives you access to Monitor II, Monitor III and the Postprocessor.

Online help is provided for Monitor II and Monitor III dialogs and the Postprocessor.

This information unit informs you about:

- the Performance Management menu
- online help
- the tutorial
- message help

Starting the reporters

The RMF Performance Management menu offers easy access to the reporting capabilities of the Monitor II and Monitor III display sessions and the Postprocessor. Just enter the TSO/E command

- RMF (or RMFJPN for the Kanji version)

You will see the following panel:

RMF - Performance Management

z/OS V2R1 RMF

Selection ==>

Enter selection number or command on selection line.

1 Postprocessor

2 Monitor II

3 Monitor III

U USER

R RMF SR

P RMF PM

N News

Postprocessor reports for Monitor I, II, and III

Snapshot reporting with Monitor II

Interactive performance analysis with Monitor III

User-written applications (add your own ...)

Performance analysis with the Spreadsheet Reporter

RMF PM Java Edition

What's new in z/OS V2R1 RMF

(PP)

(M2)

(M3)

(US)

T TUTORIAL

X EXIT

RMF Home Page: <http://www.ibm.com/systems/z/os/zos/features/rmf/>

5650-ZOS Copyright IBM Corp. 1994, 2013.
Licensed Materials - Property of IBM

Figure 5. RMF Performance Management Menu

From here, you can access the RMF Reporter you want by entering on the selection line:

- The selection number
- The abbreviation shown in parentheses to the right of the choice

Select U or US to access any user-written applications that you have defined.

Enter T to see a tutorial menu, from which you can select the RMF component you want to know more about.

Enter X to leave this panel without starting any reporter.

Reference information

In addition to the selections for invoking a specific function, there are some selections that provide information either about the current release of RMF or about functions you can perform on your workstation. If you want to use these workstation functions, at first you have to install them.

Performance analysis with the Spreadsheet Reporter

The Spreadsheet Reporter allows you to convert RMF data to spreadsheet format and provides a practical approach how to use spreadsheet macros for converted reports and Overview records.

You find all details in Chapter 18, “RMF Spreadsheet Reporter,” on page 293.

RMF PM Java Edition

RMF Performance Monitoring (RMF PM) allows you to monitor the performance of your z/OS host from a workstation through a TCP/IP interface to one or more z/OS sysplexes. You logon to any sysplex, and you can monitor the resources in the corresponding sysplex.

You find all details in Chapter 19, “RMF Performance Monitoring,” on page 341.

What's new in z/OS V2R1 RMF

Here, you find a comprehensive overview about all new functions and enhancements in the current release of RMF.

RMF in the Internet

Did you ever visit the RMF homepage in the Internet? Our address is:
<http://www.ibm.com/systems/z/os/zos/features/rmf/>

Here, you get the most current information about RMF - try it.

Quick start

You can bypass the Primary menu if you want to get directly to the reporter you need. To do this, enter the RMF command with the appropriate option:

- RMF PP to call the Postprocessor
- RMF MON2 to call Monitor II
- RMF MON3 to call Monitor III
- RMF UTIL to call the Monitor III Utility (see the *z/OS RMF Programmer's Guide*)

Getting help with RMF dialogs

Online help for RMF reporting sessions includes a tutorial, help, and message help panels. For more information about a report or a panel, press PF1. To use the RMF tutorial, either enter the T command on the Primary menu, or enter =T (using the ISPF “jump” facility) from the command line on any panel.

Getting help for a report

Press the Help key while viewing any RMF panel to see the Extended Help for that panel. The Extended Help provides access to all information related to the panel.

For tabular reports, an example of the report is shown at the top of the Extended Help. Field Help is available for all of the highlighted column headers shown in this example.

Note: In some cases, the help for several fields has been grouped together (for example, there is only one help for all fields in the report header). The highlighted line below the example indicates which column headers in the last line have been combined into a single help topic.

To see help for one of the highlighted fields in an example of a report, tab to it and press the Help key (using the Tab key will show you which fields have separate help topics).

The non-highlighted areas in the example of the report represent sample data. There is no additional help available for these areas.

Help for fields on option panels and graphic reports (Monitor III only) is provided through a list.

What do the highlighted areas mean?

The help panels for RMF contains two types of highlighted phrases. One type is called **emphasized text** and the other type is called a **reference phrase**.

- Emphasized text is highlighted merely to provide emphasis, and you cannot tab to it.
- A reference phrase is a highlighted phrase that you can tab to. If you do so, and then press the Help key, you will be presented with more information related to the phrase.

Some words about the tutorial

The tutorial provides an overview of the latest RMF features, and also acts as a reference tool for system programmers, service administrators, performance analysts and operators who use RMF.

The tutorial consists of separate sections for Monitor III, Monitor II and the Postprocessor, respectively.

The Monitor III tutorial has been expanded to include several short scenarios that illustrate how to use some of the most common Monitor III reports.

Because of its task-oriented structure, you should be able to use this tutorial as an educational tool, by going through all of the information from start to finish, and also as a reference tool to find specific information.

Wherever possible, this tutorial takes advantage of the detailed help that is available for each report. It does this by providing an overview of a task with links into the existing help information.

Message help and stacked messages

To access the message help, press PF1 after the RMF message appears on the panel. When multiple messages occur at the same time, RMF displays the first message, and stacks the others. When you press PF1, RMF displays the help panel for the first message. Below the help text, "Additional messages have occurred" appears. Press ENTER on the message help panel to display the stacked messages.

Chapter 13. Interactive performance analysis with Monitor III

This information unit describes how to find your way through the ISPF panels that are your window on the data that Monitor III provides.

To start a Monitor III session, just enter the TSO/E command RMF and select “Monitor III” from the “RMF - Performance Management” panel that comes up.

This information unit provides the following information:

- what to do with the Monitor III Primary menu
- how to get to the other screens you need
- which commands and reports are available

Monitor III gives you a single point of control for monitoring resource usage within a sysplex. You can specify the appropriate system ID for the view you want in any system report.

Before you start Monitor III

Enabling RMF:

RMF is an optional feature 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. The Monitor III session cannot be started, and you will receive the message:

ERB911I RMF is not enabled to run on this system

Data set allocation

During a Monitor III Reporter session, you can display either data gathered by a running Monitor III Gatherer session, or data recorded on VSAM data sets during an earlier gatherer session on any system. If you intend to display data from VSAM data sets, you must allocate them before you start the Monitor III Reporter session:

```
ALLOC FI(RMFDS00) DA(vsam_dsname) SHR
```

If you allocate more than one data set, and you can allocate data sets from different systems (for example, all members of your sysplex), then the DDNAMEs must be in ascending sequence without gaps. For example, if you need to allocate three data sets, the DDNAMEs would be RMFDS00, RMFDS01, and RMFDS02.

If you used names that were not contiguous, for example RMFDS00, RMFDS01, and RMFDS03, RMF would disregard those names following the gap that are not contiguous (in our example, RMFDS03).

Note:

1. If you are allocating data sets from a sysplex, it is of key importance that you allocate **all data sets** of the sysplex to enable complete reporting.
2. You can allocate only VSAM data sets which do not belong to an active Monitor III Gatherer session.

Sysplex Allocation

If you have a sysplex with four members, and you have a naming convention that the VSAM data sets of each member have the name SYS1.ERB.&SYSNAME.VSAM (see “Generalizing parmlib members” on page 34), then you would use this allocation:

```
ALLOC FI(RMFDS00) DA('SYS1.ERB.SYSTEMA.VSAM') SHR
ALLOC FI(RMFDS01) DA('SYS1.ERB.SYSEMB.VSAM') SHR
ALLOC FI(RMFDS02) DA('SYS1.ERB.SYSEMC.VSAM') SHR
ALLOC FI(RMFDS03) DA('SYS1.ERB.SYSEMD.VSAM') SHR
```

This example implies that SYSTEMA is the MVS system name of the first member.

For more information, see “Using the Data Index (DI)” on page 141.

Transferring Monitor III VSAM data sets to other systems

A Monitor III Reporter session that uses preallocated data sets does not require the Monitor III data gatherer to be running on the same system. You can therefore display on one system data that RMF has gathered on another system. This allows you, for example, to run Monitor III Reporter sessions on one system, and send the data sets from other locations to be analyzed there. Once transmitted, the data sets can be preallocated and then analyzed during a reporter session in the usual manner.

Note: This transfer of Monitor III data from one system to another is not required for real-time monitoring in the sysplex. If you want to access current data from any system in the sysplex during a reporter session, the data is made available through the sysplex data server automatically.

Sending data sets to a different system

When you have collected data in several VSAM data sets, use the CLIST **ERBV2S**, which is supplied with RMF, to unload them to a sequential data set for transport.

ERBV2S Syntax

```
ERBV2S vsam_dsn seq_dsn [TRACKS(num_tracks)]
```

Where:

vsam_dsn

The name of the Monitor III VSAM data set.

seq_dsn

The name of the sequential data set to be created.

If this parameter is specified as *, ERBV2S creates a data set name according to the following rules:

- The suffix **SEQ** is appended to the input data set name
- The first qualifier is replaced by the user's dsname prefix

num_tracks

The size of the primary extent of the sequential output data set. The default is 250 tracks. Any unused space is released after REPRO.

Examples

To unload the data from VSAM data set RMF.MONIII.DS1 to sequential data set RMF.MONIII.DS1.UNLOAD, enter:

```
ERBV2S 'RMF.MONIII.DS1' 'RMF.MONIII.DS1.UNLOAD'
```

To unload the same data to sequential data set userid.MONIII.DS1.SEQ, enter:

```
ERBV2S 'RMF.MONIII.DS1' *
```

Use TRANSMIT to send the resulting sequential file to another system for analysis.

Receiving data sets at the analyzing system

When you have received the data sets, prepare them for display by running the CLIST **ERBS2V**, which is supplied with RMF. ERBS2V allocates a Monitor III VSAM data set and REPROs the input sequential data set to that VSAM data set.

ERBS2V Syntax

```
ERBS2V seq_dsn vsam_dsn [VSAMVOL(volume)] [TRACKS(num_tracks)]
```

Where:

seq_dsn

The name of the sequential input data set that contains unloaded Monitor III VSAM data.

vsam_dsn

The name of the Monitor III VSAM data set to be created.

volume

The name of the volume on which the VSAM data set is to be allocated. If this parameter is omitted, the VSAM data set is allocated on the same volume as the input sequential data set *seq_dsn*.

num_tracks

The size of the primary extent of the VSAM output data set. If this parameter is omitted, the allocated space of the sequential input data set will be used.

Example

To load the sequential data set RMF.MONIII.DS1.UNLOAD into the VSAM data set named RMF.M3.DS1 on volume DATA10, enter:

```
ERBS2V 'RMF.MIII.DS1.UNLOAD' 'RMF.M3.DS1' VSAMVOL(DATA10)
```

Messages during Monitor III start

There are two special cases in which you might see a message on your terminal after calling Monitor III:

- ADM0873 I IF AVAILABLE, PLEASE SELECT PCLK, OTHERWISE, PRESS 'ENTER'

This message indicates that your 3270 terminal either has no graphic capability, or that you run on a multisession terminal (for example 3279) in a session that has not been defined in the VTAM control unit as graphic session. As result, Monitor III can create tabular reports only.

- IEC130I ADMPC DD STATEMENT MISSING

This message might appear in a 3270 emulator session on your workstation. You can ignore it, and Monitor III will create graphic reports.

Sysplex considerations

You might have systems in your sysplex with different releases of RMF installed. To avoid problems when reporting Monitor III data, always use an RMF reporter version that is at least equal to or higher than the highest RMF gatherer version used to collect the data to be reported.

The Monitor III Primary Menu

To start a Monitor III session, just enter the TSO/E command RMF and select “Monitor III” from the “RMF - Performance Management” panel that comes up. The panel that RMF displays in response to your selection is the Monitor III Primary Menu:

```
RMF Monitor III Primary Menu                                z/OS V2R1 RMF
Selection ==>
Enter selection number or command on selection line.

  S SYSPLEX          Sysplex reports and Data Index          (SP)
  1 OVERVIEW         WFEX, SYSINFO, and Detail reports        (OV)
  2 JOBS             All information about job delays          (JS)
  3 RESOURCE         Processor, Device, Enqueue, and Storage   (RS)
  4 SUBS             Subsystem information for HSM, JES, and XCF (SUB)

  U USER            User-written reports (add your own ...)   (US)

                                O OPTIONS    T TUTORIAL    X EXIT

                    5650-ZOS Copyright IBM Corp. 1986, 2013.
                    Licensed Materials - Property of IBM

F1=HELP  F2=SPLIT  F3=END   F4=RETURN  F5=RFIND  F6=TOGGLE
F7=UP    F8=DOWN   F9=SWAP  F10=BREF  F11=FREF  F12=RETRIEVE
```

Figure 6. Monitor III Primary Menu

Navigating from the Primary Menu

On this panel, you can tell RMF

- What you want reported
- How you want it reported
- Whether for single or multiple systems

Select **what** you want reported in one of two ways:

1. Make a selection from the list that starts at the top of the panel:
 - The single number or letter on the left
 - The command shown in upper case beside it
 - The abbreviated command shown on the right in parentheses

RMF then presents you the appropriate selection menu that allows you to select the individual report you want. Make your selection here in the same way.

2. Select an individual report directly by entering the appropriate command on the selection line. The available commands, with abbreviations and a short description of what the resulting report will contain, are listed in Table 15 on page 164. These commands are called *report commands*.

Use the parameters of the report commands to narrow down the reports to essentials.

The first choice, SYSPLEX, leads you to the sysplex reports, and choices 1 to 4 lead you to single-system reports. You can specify the system you want on the panel that displays the individual single-system reports by overtyping the current system.

In addition, option U or USER leads to a menu with user-written reports. There you find three IBM-supplied examples that are created by help of the Monitor III Utility (see the *z/OS RMF Programmer's Guide*). Each installation can use this menu to offer all installation-specific own reports.

You can also influence **how** RMF presents the reports in two ways:

1. By entering 0 or OPTIONS to reach the Option Selection menu. This guides you in specifying the report options for the session.
2. By using *option commands* to call up the data-entry panels for the options you want to specify. You will find these commands listed in "Option commands" on page 146.

As you can see from the bottom line of the panel you can also:

- Call up a tutorial about Monitor III by entering T or TUTORIAL
- End the session by entering X or EXIT

SYSPLEX

The SYSPLEX command displays the Sysplex Report Selection menu. Use this menu to select one of the sysplex reports, or the Data Index.

All sysplex reports provide a sysplex view of your system. Whenever you invoke one of these reports, the data from all systems belonging to the sysplex is retrieved and transferred to the reporting system via the RMF Sysplex Data Server.

RMF Sysplex Report Selection Menu

Selection ==>

Enter selection number or command for desired report.

Sysplex Reports

1 SYSSUM	Sysplex performance summary	(SUM)
2 SYSRTD	Response time distribution	(RTD)
3 SYSWKM	Work Manager delays	(WKM)
4 SYSENG	Sysplex-wide Enqueue delays	(ES)
5 CFOVER	Coupling Facility overview	(CO)
6 CFSYS	Coupling Facility systems	(CS)
7 CFACT	Coupling Facility activity	(CA)
8 CACHSUM	Cache summary	(CAS)
9 CACHDET	Cache detail	(CAD)
10 RLSSC	VSAM RLS activity by storage class	(RLS)
11 RLSDS	VSAM RLS activity by data set	(RLD)
12 RLRLRU	VSAM LRU overview	(RLL)

Data Index

D DSINDEX	Data index	(DI)
-----------	------------	------

Figure 7. Monitor III Sysplex Report Selection Menu

OVERVIEW

The OVERVIEW command displays the Overview Report Selection menu. Use this menu to select Workflow/exceptions, system information, and various detail reports.

```

RMF Overview Report Selection Menu

Selection ==>

Enter selection number or command for desired report.

Basic Reports
  1 WFEX      Workflow/Exceptions      (WE)
  2 SYSINFO   System information        (SI)
  3 CPC       CPC capacity

Detail Reports
  4 DELAY     Delays                   (DLY)
  5 GROUP     Group response time breakdown (RT)
  6 ENCLAVE   Enclave resource consumption and delays (ENCL)
  7 OPD       OMVS process data
  8 ZFSSUM    zFS summary              (ZFSS)
  9 ZFSACT    zFS activity              (ZFSA)
 10 SPACEG    Storage space            (SPG)
 11 SPACED    Disk space                (SPD)
 12 LOCKSP    Spin locks                (LSP)
 13 LOCKSU    Suspend locks            (LSU)

```

Figure 8. Monitor III Overview Report Selection Menu

JOBS

This command displays the Job Report Selection menu, which shows available reports about job delays. Use this menu to choose the specific job you want to analyze and the type of delay you want reported.

To get a list of active job names, use cursor-sensitive control on the **Jobname** field to invoke the Job Report Options panel.

```

RMF Job Report Selection Menu

Selection ==>

Enter selection number or command and jobname for desired job report.

Jobname ==> _____

 1 DEVJ      Delay caused by devices      (DVJ)
1A DSNJ      ..Data set level            (DSJ)
 2 ENQJ      Delay caused by ENQ          (EJ)
 3 HSMJ      Delay caused by HSM          (HJ)
 4 JESJ      Delay caused by JES          (JJ)
 5 JOB       Delay caused by primary reason (DELAYJ)
 6 MNTJ      Delay caused by volume mount (MTJ)
 7 MSGJ      Delay caused by operator reply (MSJ)
 8 PROCJ     Delay caused by processor    (PJ)
 9 QSCJ      Delay caused by QUIESCE via RESET command (QJ)
10 STORJ     Delay caused by storage      (SJ)
11 XCFJ      Delay caused by XCF          (XJ)

These reports can also be selected by placing the cursor on the
corresponding delay reason column of the DELAY or JOB reports and
pressing ENTER or by using the commands from any panel.

```

Figure 9. Monitor III Job Report Selection Menu

Job-oriented reports show delay components for jobs, such as resource delays, subsystem delays, operator, and device delays.

RESOURCE

The RESOURCE command displays the Resource Report Selection Menu. Use this menu to select reports on processors, devices, enqueue and storage. Use this menu to choose what resource you want to see delays or storage problems for.

RMF Resource Report Selection Menu

Selection ==>

Enter selection number or command for desired report.

Processor	1 PROC	Processor delays	(PD)
	1A PROCU	Processor usage	(PU)
Device	2 DEV	Device delays	(DD)
	3 DEVR	Device resource	(DR)
	3A DSND	..Data set level by DSN	(DSN)
	3B DSNV	..Data set level by volume	(DSV)
Enqueue	4 ENQ	Enqueue delays	(ED)
	5 ENQR	Enqueue resource	(ER)
Storage	6 STOR	Storage delays for each job	(SD)
	7 STORF	Storage usage by frames	(SF)
	7A STORM	Storage usage by memory objects	(SM)
	8 STORR	Storage usage for each resource	(SR)
	9 STORS	Storage summary for each group	(SS)
	10 STORC	Common storage summary	(SC)
	11 STORCR	Common storage remaining	(SCR)
I/O Subsystem	12 CHANNEL	Channel path activity	(CH)
	13 IOQUEUE	I/O queuing activity	(IQ)

Figure 10. Monitor III Resource Report Selection Menu

SUBS

The SUBS command displays the Subsystem Report Selection menu. Use this menu to select HSM, JES, and XCF Delay reports.

RMF Subsystem Report Selection Menu

Selection ==>

Enter selection number or command for desired subsystem report.

1 HSM	Hierarchical Storage Manager delays	(HD)
2 JES	Job Entry Subsystem delays	(JD)
3 XCF	Cross System Coupling Facility delays	(XD)

Figure 11. Monitor III Subsystem Report Selection Menu

USER

The USER command displays the *User-written Report Selection Menu*. Use this menu to select your user-written reports or those examples that are provided with Monitor III.

RMF User-written Report Selection Menu	
Selection ==>	
Enter selection number or command for desired report.	
1 MSI	Migration SYSINFO including Execution Velocity
2 DSD	Detailed Storage Delays
3 RG	Resource Group Data
Device Reports	
DA DEVN	Device Activity
DT DEVT	Device Trend
	Device => _____
System Reports	
ST SYSTREND	System and Workload Trend
	Workload => _____

Figure 12. Monitor III User-written Report Selection Menu

STOP and GO

Monitor III Reporter sessions can run in two modes: STOP and GO. You can specify the mode in commands or session options.

STOP mode - This is the default mode

When you start Monitor III, the first report presents either the current time interval or, if you are reporting on preallocated data sets, the newest data.

When navigating among the various reports, you always cover the same range. This enables you to see your sysplex or system data from different viewpoints that belong together. You can modify the time or the range either by using the BREF and FREF commands, or by overtyping the time, date, or range fields in the header of the report panel:

RMF V2R1 TITLE		Line 1 of 30
Command ==>		Scroll ==> HALF
Samples: nnn	System: syst Date: mm/dd/yy Time: hh.mm.ss	Range: 100 Sec

Figure 13. Header of Monitor III Single-System Reports

GO mode - You start this by command or option

GO mode is available only when reporting on current data in the sysplex. It is not possible with preallocated data sets. Use it to monitor your system continuously. By specifying a Refresh value in the session options, you define the frequency at which the requested report will be updated. Ideally, this interval should be the same as the gathering interval defined in the MINTIME gatherer option.

Note: We recommend a separate service class for TSO users who run permanently or frequently in GO mode, to avoid falsifying the average TSO response time. When you run the Monitor III Reporter in GO mode, each display of the updated report is considered as a TSO transaction. If the range is, for example, 100 seconds,

the response time for each of these transactions is counted as 100 seconds. This has a significant impact on the overall TSO response-time report, especially on systems with a small number of TSO users.

For more information on STOP and GO modes, refer to “Setting GO mode” on page 152.

Using the Data Index (DI)

The Data Index (DI) provides information about the data that is currently available for your Reporter session:

- Current® data from all active gatherers in the sysplex
- Preallocated data sets from previous gatherer sessions

To display the index, select it on the Primary menu or enter DI or DS on any command line.

You can also see if data is missing, or could not be retrieved for one of the following reasons:

- No data is available on the system
- The system does not respond
- The gatherer for the system is not active
- RMF is not active on a system
- The preallocated data set is empty or has an error

Thus the Data Index provides a compact overview of information about all systems belonging to the sysplex, regardless of whether RMF is active or not.

Contents of the Data Index

```

RMF V2R1  Data Index - RMFPLEX1                               Line 1 of 22
Command ==>                                                    Scroll ==> HALF

Samples: 118      System: MVS2  Date: 11/27/13  Time: 10.12.00  Range: 120  Sec

      ----Begin/End----
System --Date-- --Time-- -DDNAME- -----Data Set Name-----

MVS1  11/27/13  10.03.20
      10.12.00          * * *      In-storage buffer      * * *
MVS1          SYS00002 RMF.MONITOR3.DATASET1.MVS3
      * * *      Data from system MVS3      * * *
MVS1          SYS00001 RMF.MONITOR3.DATASET2.MVS3
      * * *      Data from system MVS3      * * *

MVS2  11/27/13  09.11.00 SYS00002 RMF.MONITOR3.DATASET1.MVS2
      09.14.00
MVS2  11/27/13  10.03.00 SYS00003 RMF.MONITOR3.DATASET2.MVS2
      10.12.00          * * *      Currently active      * * *
MVS2  11/27/13  10.03.00
      10.12.00          * * *      In-storage buffer      * * *

MVS3  11/27/13  09.11.00 SYS00002 RMF.MONITOR3.DATASET1.MVS3
      09.14.00
MVS3  11/27/13  10.03.00 SYS00003 RMF.MONITOR3.DATASET2.MVS3
      10.12.00          * * *      Currently active      * * *
MVS3  11/27/13  10.03.00
      10.12.00          * * *      In-storage buffer      * * *
TEST
      * * *      No response      * * *

```

Figure 14. Data Index

For each active Monitor III data gatherer in the sysplex, the Data Index lists:

- All data sets written by the gatherer
- The RMF in-storage buffer

For a Reporter session with preallocated data sets, the index lists these data sets.

Reducing information on the report

The screen allows you to display all data sets that are available throughout the whole sysplex. As this may be a long list, you can use the **DDNAMES/DSNAMES** option on the Report Options panel to reduce the data-set level information per system. If this option is used to exclude the data set names from the index, the layout changes, and the screen looks as shown in the following figure.

Command ==>		RMF V2R1 Data Index - RMFPLEX1		Line 1 of 4	
				Scroll ==> HALF	
Samples: 118	System: RMFE	Date: 11/27/13	Time: 10.12.00	Range: 120	Sec
-----Begin-----		-----End-----			
System	--Date--	--Time--	--Date--	--Time--	
MVS1	11/27/13	10.03.20	11/27/13	10.12.00	
MVS2	11/27/13	09.11.00	11/27/13	10.12.00	
MVS3	11/27/13	09.11.00	11/27/13	10.12.00	
TEST			* * *	No response	* * *

Figure 15. Data Index - Condensed Version

This screen displays information about data that is available throughout the sysplex. It shows at a glance for which time ranges data is available on each system, or if no data is available at all, or could not be retrieved because of special conditions.

Data sources

Two situations should be distinguished:

- Preallocated data sets

In this situation, the reporter retrieves data only from the preallocated data sets to the local reporter session, independent of any gatherers that are running on the various systems. It is possible to preallocate data sets created on different systems. The Data Index shows all data available in all the data sets, with the respective system-ID.

- Gatherer Session - no preallocated data sets

Here, the Data Index shows the data available through the gatherers running in the sysplex. For each gatherer, this may be the in-storage-buffer and, if data-set support is active, the data sets on which the gatherer is recording.

Rows with data that are available on the local system are displayed in turquoise. All other rows are displayed in dark blue.

Messages

The following messages can be shown in special cases:

*** Currently active ***

The currently active data set for the Monitor III data gatherer session (appears only on the Data Index for a reporter session without preallocated data sets)

*** In-storage buffer***

The local storage buffer entry of the Monitor III data gatherer

*** Empty ***

Data set with no usable data. For a session without preallocated data sets, data set recording might not be active and RMF cannot find the LRECL or CI SIZE for the data sets. For a session with preallocated data sets, the data set might be empty or contain other than sampled data gathered during a Monitor III data gatherer session.

*** No data available ***

There is no data available for the system listed in the System: field on this line.

***** No response *****

A system that is part of the sysplex, according to the XCF system name list, does not reply to the request for data

***** Gatherer not active *****

RMF is active on a system, but the Monitor III gatherer is not started

***** RMF not active in xxxxxxxx *****

The RMF address space is not active on system xxxxxxxx. Therefore, no data can be reported for this system.

The eight-character MVS system name xxxxxxxx is defined in the SYS1.PARMLIB(IEASYSxx) parameter SYSNAME.

The four-character SMF system ID, defined in the SYS1.PARMLIB(SMFPRMxx) parameter SID(yyyy) cannot be determined, and is set to '????'.

The following messages occur when the data gatherer tried to use the data set.

***** Not Found *****

Uncataloged data set specified on the DATASET option of the Monitor III data gatherer session (the data set is unusable)

***** Invalid RECSIZE *****

Data set specified with an invalid record size (the data set is unusable)

***** Invalid CISIZE *****

Data set specified with an invalid control interval size (the data set is unusable)

***** Open Error RC=xx reason=xxx *****

Error in opening the data set (the data set is unusable)

***** Close Error RC=xx reason=xxx *****

Error in closing the data set (the data set is unusable)

***** VSAM error RC=xx reason=xxx *****

Error in reading the VSAM data set (the data set is unusable)

***** DYNALLOC RC=xx IRC=xxxx ERC=xxxx *****

Dynamic allocation error (the data set is unusable)

***** UNALLOC RC=xx IRC=xxxx ERC=xxxx *****

Data set unallocated (the data set is unusable)

***** Sample time exceeds current time *****

Data set with a sample time that is later than the current system time. The system time has probably been incorrectly set. (This message does not appear on the screen with preallocated data sets.)

***** Data from sysplex xxxxxxxx *****

For either preallocated data sets or gatherer data sets, a data set that is from a sysplex other than the one selected has been encountered. Only one sysplex can be represented by the data on the Data Index. No other reports can be shown as long as this error persists.

***** Data from system xxxx *****

The reporter cannot report data from gatherer data sets from another system. The gatherer marks the data sets as unusable if more than one system has written to a data set. The reporter cannot access the data in data sets that are marked unusable.

The reporter also cannot report data from different sysplexes in one session.

Field descriptions

Table 10. Field Descriptions for the Data Index

Field Heading	Meaning
System (on detailed and condensed version)	The four character SMF system identifier.
Begin/End Date Time (on detailed version)	These are the beginning and ending dates/times for the data in the usable and not empty data sets or the in-storage buffers.
Begin Date Time (on condensed version)	The begin date and time for which data is available on the respective system.
End Date Time (on condensed version)	The end date and time for which data is available on the respective system.

Note: If you are using old data, the sysplex ID and other fields may be blank.

Cursor-sensitive control

Cursor-sensitive control on the *System* field switches to the selected system. This means that data from the requested system is retrieved, if available, and the Data Index is redisplayed, with the selected system shown in the header System field, and the corresponding lines of the report shown in turquoise.

Data Index options

RMF Data Index Options

Command ==>

Change or verify parameters. Press END to save and end.

DDNAMES/DSNAMES ==> YES

Sort Order ==> ASCEND

Include DDNAMES / DSNAMES information (YES NO)

Sort data set names (ASCEND DESCEND)

Figure 16. Data Index Options Panel

The Data Index has two options: the *DDNAMES/DSNAMES* and the *Sort Order* options.

DDNAMES/DSNAMES

Allows switching between a panel listing data set names, as shown in Figure 14 on page 142, and a panel giving a condensed list of systems belonging to the sysplex, as shown in Figure 15 on page 143.

Sort Order

Specifies the sort order of the displayed data sets.

The sort criteria are at first the System ID and within each system the end date/time of the available data.

If more rows than one with the same system ID exist, the usable data sets that are not empty are listed first, then the empty data sets, and finally the unusable data sets. The usable data sets that are not empty are sorted by the end time of the stored data.

Monitor III session and option commands - Overview

This chapter provides an overview of the following types of Monitor III commands:

- “Session commands”: these commands help you to work within a session.
- “Option commands”: these commands help you to define the appearance of a session.

Session commands

Table 11. Monitor III Session Commands

Task	Command	Parameters	Result
“Backward and forward referencing” on page 147	BREF FREF	DATE = TIME = RANGE = SYSTEM =	Changes date, time, range, and system
“Cancelling entries on option panels” on page 151	CANCEL		Restores options to state at panel entrance (except Job Report Options panel)
“Getting help for RMF commands” on page 151	COMMANDS (COM, CMD)		Displays RMF help menu for commands
“Displaying current range data” on page 151	CURRENT (CU)		Retrieves current data for display
“Searching for a field” on page 151	FIND (F, FI)		Searches for character string on report panels
“Setting GO mode” on page 152	GO		Switches to GO mode processing
“Activating GRAPHIC mode” on page 152	GRAPHIC (GR)	ON OFF	Switches to graphic mode Switches to tabular mode
“Printing screens and tabular reports” on page 153	HARDCOPY (HC)	ON OFF REPORT SCREEN	Prints all displayed reports Prints no reports Prints tabular copy of reports Prints a copy of the screen image
“Printing graphic reports” on page 154	ICU		Sends report data to ICU and starts an ICU session
“Using program function keys” on page 154	PFK		Displays list of PF keys
“Resetting entries on option panels” on page 155	RESET		Resets options (excluding JOBNAME options) to default values shipped with RMF
“Retrieving the last command” on page 156	RETRIEVE		Displays last command entered on the command line
“Searching for a field” on page 151	RFIND		Repeats the FIND command
“Activating TABULAR mode” on page 156	TABULAR (TAB)	ON OFF	Switches to tabular mode Switches to graphic mode
“Toggling between tabular and graphic display” on page 156	TOGGLE (TOG)		Switches between tabular and graphic display

Option commands

To change the options of an RMF Monitor III Reporter session, select 0 on the Primary menu, or enter the command OPTIONS on the command line of any panel.

In response, RMF displays the Option Selection menu:

RMF Option Selection Menu

Selection ==>

Select one of the following options or enter command. Press ENTER.

1 SESSION	Set Session Options	(S0)
2 COLOR	Set Graphics Colors and/or Patterns	(C0)
3 LANGUAGE	Set Language and Date/Time Options	(L0)
4 ROPTIONS	Select report options for	
	REPORT ==> _____	(R0)
5 OPTSET	Change or Select Option Set	(0S)

Figure 17. Option Selection Menu

To leave the panel without making a selection, enter END on the selection line.

Table 12. Monitor III Option Commands

Task	Command	Result
"Changing session options" on page 156	SESSION	Specifies options that are valid for all reports displayed during this session.
"Changing color graphic options" on page 158	COLOR	Defines the colors you like to have in the reports.
"Changing language options" on page 160	LANGUAGE	Defines language-specific display of date and time.
"Changing report options" on page 160	ROPTIONS	Sets or modifies options for a specific report. Therefore, if you make this selection, you must enter a report name in the field REPORT ==> Report names and their valid abbreviations are listed in Table 15 on page 164.
"Selecting an option set" on page 161	OPTSET	Builds a set of options and stores it for later use. If you build several different sets of options, you can select the appropriate one for a given session.

Monitor III session tasks

This topic provides information about the tasks that you can perform during Monitor III sessions, like for example, printing screens and tabular reports, changing session options or using cursor-sensitive control.

Backward and forward referencing

In STOP mode, you can obtain reports on any data in the data gatherer's in-storage buffer or, with data set support, data contained in user-defined data sets through the use of the BREF (backward referencing) and FREF (forward referencing) commands. You can also do this by using the Date, Time, System and Range fields on the report panels.

Issue this command from a report screen. If you issue this command on a non-report screen, RMF displays the last report viewed for the current Date and Time, and for the Range specified on the Session Options panel. If you have not viewed any reports during the session, RMF displays the Workflow/Exceptions (WFEX) report.

Depending on the parameters specified, and on whether or not you are using data-set support, you can display data from either:

- The data gatherer's in-storage buffer on any or all of the systems in a sysplex
 - The data gatherer data sets on any or all systems in the sysplex
- or
- Preallocated data sets

You can use the DI report to list, by system ID, the beginning date/time and end date/time for samples stored on each data set used during data set recording.

Please keep in mind that the easiest way to specify all these values is to enter them directly into the report header line, as in Figure 18, rather than on BREF or FREF commands:


RMF V2R1 TITLE		Line 1 of 30
Command ==>		Scroll ==> HALF
Samples: nnn	System: syst	Date: mm/dd/yy Time: hh.mm.ss Range: 100 Sec

Figure 18. Header of Monitor III Single-System Reports


The syntax of the BREF and FREF commands is:




D:



T:



R:



S:



The parameters are all optional, and indicate the following:

DATE

Specifies the month, day, and year of the data you want. If you omit this parameter, RMF uses the date displayed on the screen. Leading zeroes can be

omitted. The sequence you use for the month, day and year on the BREF/FREF commands must be the same as the sequence specified on the language options panel. RMF supports a sliding window which covers the time frame:

Current Year - 50 ↔ Current Year + 49

This sliding window will be used to define the correct value of the century.

TIME

Specifies the hour, minute, and second of the data you want to retrieve first. If you omit this parameter, RMF uses the begin or end time of the report currently displayed on the screen. The conditions under which RMF uses the begin or end time are described later in this section. Leading zeroes can be omitted. Seconds or hours can be omitted if they are zeroes. For example, specify TIME=9.5 or TIME=9:5, rather than TIME=09.05.00 or TIME=09:05.00. You can use T as an abbreviation for TIME.

RANGE

Specifies the time range over which you want RMF to summarize and present the sampled data. Valid time range values are 0 to 9999 seconds or 0 to 166 minutes. If you specify a value without M or S, RMF uses seconds. If you omit the RANGE parameter, RMF uses the RANGE value currently on the screen. You can use R as an abbreviation for RANGE.

SYSTEM

Allows you to report on any single system in the sysplex. For systemname specify the name of the system you want to report on. All following single-system reports show data from the specified system, until you specify another system.

Note:

1. If the data defined by the DATE, TIME, and RANGE parameters is not available in the data gatherer's storage or, if you have specified data sets, in either the data gatherer's storage or user-defined data sets, RMF issues a message to indicate which data is available.
2. If you have specified data sets during a Monitor III data gatherer session, time gaps in the recorded sets of samples might have occurred during data set recording. If, during the reporter session, RMF detects gaps for the requested RANGE time, the following occurs:
 - If all of the data defined by the DATE, TIME, and RANGE parameters is not available because of a time gap, RMF issues messages describing the BEGIN/END time of the gap.
 - If part of the data defined by the DATE, TIME, and RANGE parameters is not available because of one or more time gaps, RMF issues a message to indicate the BEGIN/END time of the first gap. RMF displays the available data, but because some reported values like TCB + SRB time depend on the actual time of the sampling, the results can be misleading.
3. If the TIME specified is not exactly at the beginning of a MINTIME interval, or the RANGE is not a multiple of MINTIME, RMF might present more data than you request. RMF always presents the data that includes the TIME and RANGE values you specify except if the begin or end time of an interval lies within a time gap.
4. You should be aware that a large RANGE value increases the local storage area and CPU time needed by the data reporter.

The BREF and FREF commands perform the same function when you specify a DATE and/or TIME value (with or without a RANGE value). Both commands allow you to pinpoint the time at which you want to start viewing data collected either prior to or subsequent to entering STOP mode.

The BREF and FREF commands perform different functions when one of the following conditions occurs.

- RANGE is the only parameter specified
- No parameters are specified.

Using BREF/FREF with the RANGE value

If RANGE is the only parameter you specify, the FREF and BREF commands use the TIME value currently displayed on the top of the screen. The FREF command uses the TIME value as the beginning time of the new report and adds the RANGE value you specify to obtain the report interval. The BREF command uses the end time of the current report interval (TIME plus RANGE value displayed at the top of the screen) and subtracts the RANGE value you specify to obtain the beginning time of the new report interval.

Using BREF and FREF by specifying only a RANGE value allows you to include in the report interval data from the current report interval indicated by TIME and RANGE at the top of the screen. With BREF, you can access data in a previous interval as well as the current interval; with FREF, you can access data in a later interval as well as the current interval. For example, if the current RANGE and MINTIME values are 100 seconds, and the TIME on the top of the display screen is 9:00:00, then RMF displays a report containing data from 9:00:00 to the TIME + RANGE value at the top of the screen, which would be 9:01:40. To view data from a previous interval, as well as the currently displayed interval (9:00:00 to 9:01:40), specify BREF R=200. RMF presents a report containing data from 8:58:20 to 9:01:40. (8:58:20 being the TIME + RANGE value at the top of the screen minus 200.) To include one more preceding interval, specify BREF R=300 to present data from 8:56:40 to 9:01:40. If you want to display data only from the current interval again (9:00:00 to 9:01:40), shorten the range parameters on the command to 100 seconds (BREF R=100).

Using the FREF command you can display reports containing data from subsequent intervals. For example, specify FREF R=200 to display a report containing data from 9:00:00 to 9:03:20.

If you specify RANGE=0 with the date and/or time, you can pinpoint the time at which you want to start viewing data. RANGE=0 causes RMF to adjust the range to the smallest possible value, which is the MINTIME value you specified in the data gatherer options. If you want to begin viewing the report at the TIME value on the display screen and you specify BREF or FREF, you must specify RANGE = 0.

Using BREF/FREF without parameters

If you do not specify any parameters, the FREF command uses the TIME value on the display screen and *adds* the RANGE value (on the screen) to calculate the begin time of the data RMF retrieves. The BREF command uses the TIME value on the display screen and *subtracts* the RANGE value to calculate the beginning time of the data RMF retrieves.

Once you pinpoint the time that you want to start viewing data collected by the data gatherer, you can issue additional FREF or BREF commands to move

backward and forward in time. You can also use the PF10 or PF11 keys, which have default settings of BREF and FREF, respectively.

Cancelling entries on option panels

If you have made changes and wish to return to the values that were originally on the panel when you first entered the screen, enter on the command line:

►►—CANCEL—◄◄

CANCEL cancels all changes you have made except for Date, Time, and Range fields.

Note: CANCEL does not work on the Job Report Options panel.

Getting help for RMF commands

The COMMANDS command displays the RMF help menu for commands, where you can access a definition of the command you want more information on:

►►—COMMANDS—◄◄

Displaying current range data

To display a report with data from the current time for the length of the current range value, enter on the command line:

►►—CURRENT—◄◄

“Current range value” means the range value specified on the Session Options panel. This may be different from the range you saw last, if the range has been altered explicitly during the session by:

- A BREF or FREF command with an explicit RANGE option
- A BREF or FREF command using the range displayed on the screen
- Entering a range in the report panel input field

Issue this command from a report screen. If you issue it on a non-report screen, RMF displays the last report viewed for the current date and time, and for the range specified on the Session Options panel. If you have not viewed any reports during the session, RMF displays the Workflow/Exceptions (WFEX) report.

Note: The CURRENT command does not work when you are reporting from preallocated data sets.

Searching for a field

To search for a field on a scrollable report, enter on the command line:

►►—FIND—*string*—◄◄

where *string* is a character string that can be enclosed in single quotes, but cannot contain any blanks.

To find the next occurrence of that string enter:

►►—RFIND—◄◄

RFIND is usually assigned to PF5.

When you issue a FIND command on a **tabular report**, RMF searches *from the cursor position* down, displays the line where the character string was found as the top row, and positions the cursor at the beginning of the character string.

When you issue a FIND command on a **graphic report**, RMF actually searches the tabular version of the report. That is, you can use FIND and RFIND successfully on a character string that does not appear in the graphic report, but appears in the tabular report. The search takes place *from the top line* down. RMF displays the bar corresponding to the line of the tabular report in which it found the character string as the top graphic bar, and positions the cursor on the command line.

Setting GO mode

To switch from STOP mode to GO mode, enter on the command line:

►►—GO—◄◄

When you enter GO, RMF resets the Range from the value on the Session Options panel.

These are some rules to keep in mind while using the GO command:

- You can not enter any commands on the command line while in GO mode.
- If you enter GO on a panel that is not a report, the last displayed report will be displayed in GO mode or, if no report has been previously displayed during the session, the Workflow/Exceptions report.
- You cannot enter GO mode during a reporter session with preallocated data sets. If you have specified MODE(GO) on the Session Options panel, RMF ignores the GO option.

STOP mode is the default. To ensure the default mode for your system is current with the RMF default, enter the RESET command from the Session Options panel. RESET reestablishes the RMF default settings. When a new option set is created for a new user, the mode is automatically set to STOP.

To switch from GO mode to STOP mode, press the ATTN key or the PA1 key. When using a programmable workstation, typically you will get these keys with the right mouse-click. This action freezes the current report so you can page through it. While in STOP mode, the data gatherer continues to collect data and place it in local storage. With data set recording, the data gatherer continues to copy data from local storage to the data sets.

Activating GRAPHIC mode

If you are in TABULAR mode when you start a session, use the GRAPHIC command to switch modes. Enter the GRAPHIC command on the command line of any report:



RMF activates GRAPHIC mode, and if your terminal supports graphics, and your installation has the Graphical Data Display Manager (GDDM) and the Presentation Graphics Feature (PGF) program products, both Version 2 Release 1 or later, you can display graphic reports. The default for GRAPHIC is ON. To return to tabular report display, specify GRAPHIC OFF on the command line.

Printing screens and tabular reports

Enter the HARDCOPY command on the command line to print a screen or a report. This command has the syntax:



The parameters, which are optional, have the following effect:

ON Prints all reports requested during the session, and is equivalent to specifying HARDCOPY on the Session Options panel.

OFF

Ends the hardcopy mode.

To print a single report or screen when you specify HARDCOPY OFF on the Session Options panel, enter HARDCOPY SCREEN or HARDCOPY REPORT.

SCREEN

Prints the displayed screen.

HARDCOPY SCREEN will print any report-screen image.

REPORT

Prints the whole report (because a report can be longer than one screen).

The command causes RMF to print all frames of the report whether they are displayed or not.

RMF writes all reports requested during the session to the output data set you specified on the Session Options panel, or to SYSOUT if an output data set is not specified. The output data set must have the DCB parameters:

RECFM(VBA),LRECL(137)

Hardcopy prints only tabular reports; if you specify HARDCOPY ON and access any graphic reports during a session, RMF prints the tabular version of the report.

If you enter HARDCOPY without parameters on the command line, the default is ON, which prints the tabular version of all reports you access during the session.

Note: You should use the ISPF PRINT command only in tabular mode. If used in graphic mode, unpredictable results will occur. For more information about ISPF commands, see *z/OS ISPF User's Guide Vol I*.

Printing graphic reports

To print RMF graphic reports, use the Interactive Chart Utility (ICU). Issue the ICU command from the command line of a graphic report:

►►—ICU—◄◄

The ICU command creates a graphics data file (GDF) of the current screen, starts an ICU session, and displays the initial empty DIRECTORY panel.

To display all of the GDFs created, type **L** in the Commands column and **GDF** in the Type column of the line marked *******. Figure 19 shows a sample DIRECTORY panel.

DIRECTORY					
ADM1042 I 3 ITEM(S) LISTED					
Commands	Name	Type	Library	Date and Time	Description
			DDNAME No.	Last Written	
***		GDF	ADMGDF	09 NOV 2013 10:24 AM	
001	DELAY	GDF	ADMGDF	0 09 NOV 2013 10:24 AM	RMF
002	WFEX	GDF	ADMGDF	0 08 NOV 2013 1:43 PM	RMF
003	WFEX1	GDF	ADMGDF	0 03 NOV 2013 9:47 AM	RMF
Commands: D (Delete) P (Pick Name) C (Copy From) / (Scroll Here)					
SH (Show GDF File) PR (Print GDF File)					
*** Line only: L (List) T (Copy To)					
PF: 1=Help 6=Show Description 7=Up 8=Down 9=Exit					

Figure 19. Interactive Chart Utility (ICU) DIRECTORY Panel

On the ICU DIRECTORY panel, use the ICU commands to view, print, and process the GDFs. For more information on the ICU commands, use the HELP (PF1) key. To return to the RMF session, use the EXIT (PF9) key.

When you enter the ICU command, RMF saves the displayed screen of the graphic report as a member in the data set:

userid.RMFZR21.ADMGDF(report name)

The member remains in this data set until you delete it. RMF saves and re-uses this data set every time you start a Monitor III Reporter session. Because RMF uses the report name as the member name, the next time you enter ICU on the same report, the member is overwritten. To avoid this, you can either:

- Copy the GDF member into a new member with a different name.
On the ICU DIRECTORY panel, you can use the “Copy from” and “Copy to” commands.
- Rename the member before entering the ICU command again.

Using program function keys

Issue the PFK command to display the program function keys:

►►—PFK—◄◄

You can also use the ISPF KEYS command.

Table 13 shows the default PF key settings. The settings for PF keys 13 to 24 are identical to the settings for PF keys 1 to 12.

Table 13. Program Function Keys Defaults

PF Key	Default Setting
PF1	HELP
PF2	SPLIT
PF3	END
PF4	RETURN
PF5	RFIND
PF6	TOGGLE
PF7	UP
PF8	DOWN
PF9	SWAP
PF10	BREF
PF11	FREF
PF12	RETRIEVE

Using PF keys to build commands

When you press a program function key, RMF builds a command by using the command string defined for the PF key and adding any text in the input line. For example, if you specify T=10.05 on the input line and press PF10, RMF builds the command BREF T=10.05.

Changing PF key settings

To change the settings for any of the 24 PF keys, use the ISPF KEYS command to access the ISPF PFK screen. There, change the setting next to the PF key, and press ENTER. Changes remain in effect until you alter them again.

Note: PF key changes are not stored in RMF option sets. There is only one set of PF key definitions associated with your RMF session.

Resetting entries on option panels

To ensure the RMF default settings for option panels are in effect, enter RESET on the command line or the respective panel.

►►—RESET—◄◄

RESET reestablishes RMF's default settings.

Note: Because there is no default value for jobname, the RESET command is not valid on the job report options panel.

The CANCEL command changes the value back to what it was when you entered the panel.

For Workflow/Exception (WFEX) and GROUP report options, if you specified YES at “Customization” on the Session Options panel, RESET invokes automatic customization and re-establishes defaults.

Retrieving the last command

Use the ISPF RETRIEVE command to recall the last command you entered.

►►—RETRIEVE—◄◄

Activating TABULAR mode

Issue the TABULAR command on the command line to display tabular reports:

►►—TABULAR—☐ ON
 ☐ OFF—◄◄

TABULAR ON is the default. To return to a graphic display, you can specify TABULAR OFF on the command line of any panel.

Toggling between tabular and graphic display

To switch between tabular and graphic displays, press PF6 or enter the TOGGLE command on the command line of any report on a terminal that supports graphics.

►►—TOGGLE—◄◄

TOGGLE causes RMF to change the display format between graphic and tabular, maintaining the same scrolling position on the screen.

Displaying user-written reports

Monitor III includes a user exit for both the data gatherer and the data reporter session. Use the Report Format Definition Utility to create unique user reports. Specify the report selection on the user-report menu. See the *z/OS RMF Programmer's Guide* for more information on user-written reports.

Changing session options

The Session Options panel lets you specify options that apply to more than one report. To display it, select "Set Session Options" from the Option Selection menu, or enter the command SESSION on the command line of any panel.

RMF Session Options		
Command ==>		
Current option set: SYS1POL on MVS1		
Change or verify parameters. Press END to save and end.		
Mode	==> STOP	Initial mode (STOP GO)
First Screen	==> PRIMARY	Initial screen selection (ex: PRIMARY)
Refresh	==> 100	Refresh period (in seconds)
Range	==> 100S	Time range 10-9999 sec (ex: 100S, 100)
		1-166 min (ex: 2M)
Time Limit	==> NONE	Time limit or NONE
		1-999 min (ex: 10M)
		1-128 hours (ex: 1H)
Hardcopy	==> OFF	Hardcopy mode (ON OFF) (ex: ON)
SYSOUT	==> A	Class for printed output (ex: A)
Output Data Set	==>	
		Data set for hardcopy (Overrides SYSOUT)
Report Format	==> TABULAR	(GRAPHIC TABULAR) (ex: GRAPHIC)
Customization	==> YES	Automatically tailor WFEX report (YES NO)
Input Data Set	==> 'SYS1.PARMLIB'	
		Data set for customization (COMPAT mode only)

Figure 20. Session Options Panel

Figure 20 shows the RMF default session options. The values saved on this panel become part of the current option set, and apply to all displayed reports whenever that option set is in effect. For more information about options sets, see “Selecting an option set” on page 161.

From the Session Options panel, you can:

- Select the display mode (STOP or GO)
- Select the panel you want to appear when you start an RMF session
- Set the refresh period for the reports (GO mode only)
- Set the time range over which you want data reported
- Set the time limit for reports (GO mode only)
- Turn hardcopy mode on or off
- Specify the SYSOUT class
- Specify an output data set for hardcopy reports. This overrides the SYSOUT specification. The data set must already exist. See “Printing screens and tabular reports” on page 153.
- Choose graphic or tabular display for Monitor III reports
- Choose automatic customization for the WFEX report
- Specify the Parmlib from which customization information is to be taken

For more information about the parameters on the Session Options panel, use the HELP (PF1) command.

To leave the panel and save the changes, use the END (PF 3) command. If RMF detects errors, it displays the Session Options panel again with an appropriate error message. If all entries are correct, the changes take effect immediately and remain valid for subsequent sessions, until they are changed again or you choose another option set.

The options “Mode” and “First Screen” are exceptions. They take effect when you start the next RMF session.

To leave the panel without making any changes, enter CANCEL on the command line. If you have not typed anything in, F3 or END has the same effect.

Changing color graphic options

This two-part panel allows you to specify colors and patterns for the graphic displays of Monitor III reports. You can use this panel only if:

- GDDM and PGF are installed on your host, and
- Your terminal supports graphics

To display the first part of the panel, select 2 on the Option Selection menu, or enter the command COLOR on the command line of any panel. To access the second part of the panel press the DOWN key, and to return to the first part, press the UP key.

Table 14 on page 159 describes the fields on the color graphic options panel.

Command ==>>>>RMF Color Graphic OptionsScreen 1 of 2

Change or verify the color and pattern assignments.
Press DOWN to view the second screen. Press END to save and end.

ID	Name	Color	Pattern	ID	Name	Color	Pattern
1	CMD Line	==>> 7		9	ENQ DLY	==>> 2	==>> 14
2	Headings	==>> 1		10	HSM DLY	==>> 2	==>> 5
3	Title	==>> 7		11	JES DLY	==>> 2	==>> 13
4	CON	==>> 5	==>> 3	12	LOCL DLY	==>> 2	==>> 12
5	DSC	==>> 5	==>> 7	13	OUTR DLY	==>> 2	==>> 14
6	PND	==>> 5	==>> 5	14	PROC DLY	==>> 2	==>> 10
7	COMM DLY	==>> 2	==>> 9	15	STOR DLY	==>> 2	==>> 9
8	DEV DLY	==>> 2	==>> 12	16	SWAP DLY	==>> 2	==>> 13

1234567Colors




Figure 21. Color Graphic Options - Panel 1

RMF Color Graphic Options Screen 2 of 2

Command ==>

Change or verify the color and pattern assignments.
Press UP to view the first screen. Press END to save and end.

ID	Name	Color	Pattern	ID	Name	Color	Pattern
17	VIO DLY	==> 2	==> 11	21	EApp1%_	==> 2	==> 0
18	DEV USG	==> 4	==> 3	22	==> _____	==> -	==> -
19	PROC USG	==> 4	==> 0	23	==> _____	==> -	==> -
20	USING	==> 4	==> 4	24	==> _____	==> -	==> -
28	SUBS DLY	==> 2	==> 13	25	==> _____	==> -	==> -
29	OPER DLY	==> 2	==> 5	26	==> IOSQTIME	==> 2	==> 0
30	XCF DLY	==> 2	==> 11	27	==> SRVCTIME	==> 6	==> 0
31	OTHR DLY	==> 2	==> 11	32	==> CONT RT	==> 2	==> 12

1 2 3 4 5 6 7 Colors



Figure 22. Color Graphic Options - Panel 2

Table 14. Fields on the Graphic Options Panels

Field Heading	Meaning
ID	Specifies the ID number that represents the item's name, color and pattern assignments.
Name	Specifies the name of the report item that the colors and patterns represent.
Color	Specifies a number (1-7) that represents the color for the graphic bar that depicts the data for Name.
Pattern	Specifies a number (0-16) that represents the pattern for the graphic bar that depicts data for Name.

Items on the Color Graphic Options panel can represent the command line, headings, titles, and the graphic bars that contain and display data on the graphic reports. You can change, for example, the color of the command line (CMD Line), or the pattern for the device delay bar (DEV DLY). You enter changes directly on the panels by choosing colors or choosing patterns.

Choosing Colors

You can choose among 7 colors. The numbers corresponding to the colors (1-7) appear at the bottom of the screen. Enter the desired color number under the Color column of the items you want to change.

Choosing Patterns

You can choose among 17 patterns. The numbers corresponding to the patterns (0-16) appear at the bottom of the screen. Enter the desired pattern number under the Pattern column of the items you want to change.

If you want to add an item to be reported, enter it under the Name column on one of the blank lines on the second color graphic options panel and assign it a color and pattern. Entering the CANCEL and RESET commands changes the values on both panels, regardless of which one it was entered on. See “Cancelling entries on option panels” on page 151.

To save changes on the color graphic panels and exit, use the END key. Changes become part of your current option set and are saved across sessions.

Changing language options

The Language Options panel allows you to specify the following for all Monitor III output and report options:

- Format of the date
- Character used to separate the date
- Character used to separate the time
- Character used as a decimal point for output values.

Note: On **input**, the decimal point is always ‘.’.

To display the language options panel, select option 3 on the Option menu or enter the LANGUAGE command on the command line of any panel.

RMF Language Options		
Command ==>		
Change or verify parameters. Press END to save and end.		
Date Format	==> MDY	Order for input and output Month (M), Day (D), and Year (Y)
Date Separator	==> /	Date separator for output (/ or . or -)
Time Separator	==> .	Time separator for output (. or : or -)
Decimal Point	==> .	Decimal point in output (. or ,)

Figure 23. Language Options Panel

For more information about the parameters, use the HELP (PF1) command.

Changing report options

The Report Options panels allow you to change the options for all RMF reports. You can customize reports to allow for different jobs, resource names, and workflow exceptions to appear in the report displays. In addition, you can specify service classes, report classes and workload groups.

To obtain the Report Options panel for a report, specify the ROPTIONS command on the command line of the report you wish to change. Figure 24 on page 161 shows a Report Options panel for the DELAY report. For a complete description of these panels for each Monitor III report, see *z/OS RMF Report Analysis*.

You can also select the Report Options panel for a report from the Option Selection menu. Enter the full name of the report (or any valid abbreviation) on the REPORT line of the ROPTIONS selection, and select 4 on the command line of the Option Selection menu.

Many Report Options panels offer **wild-card support**. To select groups or jobs with similar names, you can use an asterisk (*) as a wild card in the last position of the name. You will find details in *z/OS RMF Report Analysis*.

RMF saves all of the values entered on a Report Options panel in your current option set. The options take effect immediately.

```

RMF Delay Report Options: DELAY
Line 1 of 1
Command ==> Scroll ==> HALF

Change or verify parameters. To exit press END.
All changes (except for Summary and Criterion specification) will apply to
DELAY, DEV, ENQ, HSM, JES, PROC, STOR, STORC, STORF, and XCF.

Class      ==> ALL      Classes: ALL TSO BATCH Started task ASCH OMVS
Service class ==> *ALL    *ALL or one of available service classes below
Summary    ==> NO      Class summary lines on DELAY report (YES NO)
Criterion  ==> 0        Minimum delay to include job in DELAY report

Jobs       ==> NO      View job selection/exclusion panel next (YES NO)

Available Service classes
APPPRIME  NRPRIME  OMVS      TSOPRIME  SYSTEM  SYSSTC

```

Figure 24. DELAY Report Options Panel

Selecting an option set

An option set contains all the options that you can define on the option panels:

- Session
- Color graphic
- Report
- Language

The Option Set menu lets you build or select different sets of options to control an RMF display session. To display the menu, enter option O on the Option menu or enter the OPTSET command on the command line of any panel.

The menu allows you to add or delete option sets. All option sets appear in alphabetical order on the panel; however, only one option set can be active or current for an RMF session, and you cannot delete an active option set. If an option set is not current, RMF saves it by name and description. The recommended option set can be deleted only if automatic customization is not active (see “When you use automatic customization” on page 162).

RMF is shipped with a default option set called INITIAL, which appears on the Option Set menu:

```

RMF Option Set Selection Menu
Line 1 of 2
Command ==> Scroll ==> HALF

Enter a code in the action column or fill in a new option set. Press END.
Action codes: Delete (D)  Select (S)

Current Option Set:    STANDARD on SYS5
Recommended Option Set: STANDARD on SYS5

Action  Name      System  Description
-
-
-  STANDARD  SYS5    Generated from option set INITIAL 04/14/13
-  INITIAL    initial RMF options

```

Figure 25. Option Set Menu

When you use automatic customization

If you use automatic workflow/exceptions (WFEX) customization, RMF creates or selects option sets for you. If you specified Customization YES on the Session Options panel, RMF automatically selects the option set listed under Recommended Option Set and makes it current. When automatic customization selects the current option set, all options, not only the WFEX report options, are switched.

With automatic customization, every time data is retrieved from the data gatherer, RMF checks that the options set name and the system ID of the data match the option set name and the system ID of the current option set.

If the option set name and the system ID match, processing continues under the current option set. If the option set name and the system ID do not match, RMF does one of the following:

- If an option set exists whose name and system ID match the option set name and system id of the data from the data gatherer, RMF selects that option set and makes it current.
- If no option set exists with a matching name, RMF creates a new option set and makes it current. For option set name, RMF uses the name of the active service policy. RMF sets the options for all reports, except the WFEX and GROUP reports, from the previous option set that was in effect.

Note: The automatic customization can be performed only if you have access authority to the Parmlib data set. Otherwise, you will get an error message and Monitor III will continue with its default options.

Creating a new option set

To create a new option set, enter a name and a description on the input lines on the option set selection menu, and press enter. RMF initializes the new option set with the values of the current option set. An entry in the Description field is optional.

Making an option set current

To make an option set current, place an S in the Select column next to the option set name. You can create a new option set and make it current at the same time by placing an S next to the option set name you specify on the input line and then pressing ENTER. The option set you select becomes the current option set.

If automatic customization is active, and you select an option set other than the recommended option set, customization is de-activated. To re-activate automatic customization, you must make the recommended option set current.

Deleting an option set

To delete an option set, enter D in the Select column next to the name of the option set and press ENTER. RMF displays a warning panel to confirm the delete. However, you cannot delete the current option set. If automatic customization is active you cannot delete the recommended option set. If customization is not active, you can delete the recommended option set.

Changing an option set

If you want to change options in an option set, you must first make the option set current; then change the session, color graphic, report, and language options, using the option panels. RMF records the changes that you make on these panels during the session in the current option set.

Using cursor-sensitive control

Cursor-sensitive control lets you place the cursor on a field in a tabular report, press the ENTER key, and see another report containing any additional information about the field. You can move from one RMF report to the other without returning to the primary menu or entering specific commands.

RMF keeps track of your path. Pressing the END (PF3) key returns you to the previous report until you reach the point at which you started.

Note: If you press the RETURN (PF4) key, or use the jump function, or, in a sysplex environment, switch from one system to another, RMF displays the Primary menu and you lose all return paths.

If you issue any RMF command while using cursor-sensitive control, or use cursor sensitivity to select a new system ID, RMF will erase the return path up to the point at which you did so.

Cursor-sensitive control is active on:

- Most fields on all tabular reports except STORCR
- The Jobname field of the Job Report Selection menu
- The Report Type field of the Option Selection menu
- All system lines in the Data Index

Cursor-sensitive control is not active on:

- Most selection and option panels
- Graphic reports
- The STORCR report panel
- RMF reports that you modify

Monitor III help facility

For the Monitor III Reporter dialog, an online help structure is available, in addition to the relevant part of the RMF Tutorial. You can get help for any panel by pressing PF1.

For more details on scope and handling of online help, see “Getting help with RMF dialogs” on page 130.

Monitor III report commands - Overview

Table 15 on page 164 lists all report commands with their parameters and abbreviations. The “How to request this report” section in the *z/OS RMF Report Analysis* for each report shows an example of the command and parameters.

You can enter the commands on any command line.

The **Parameters** column in Table 15 on page 164 indicates what parameters, if any, you can specify on the respective commands:

cfname

A coupling facility name

job_class

One of the following names of a job class:

ALL (A)

ASCH (AS)

BATCH (B)
 OMVS (O)
 STC (S)
 TSO (T)

Note:

1. This parameter is optional. If it is not specified, ALL is used by default.
2. In addition, ENC (or E) can be specified as class for the DELAY report.

dsname

A data set name

jobname

A job name

period A service or report class period

resource

A resource name

service_class

A service class name

s/r-class

A service or report class name

ssid A cache subsystem identifier

sstype The name of a subsystem that schedules enclaves

storage_class

A storage class name

volser A serial number of a volume

wlm The name of a workload group, a service class, or a report class

Table 15. Report Commands

Command	Parameters	Displays	Abbreviation
CACHDET	ssid	Cache detail report	CAD
CACHSUM		Cache summary report	CAS
CFACT	cfname	coupling facility activity report	CA
CFOVER	cfname	coupling facility overview report	CO
CFSYS	cfname	coupling facility system report	CS
CHANNEL		Channel path activity report	CHAN, CH
CPC		CPC capacity report	
DELAY	job_class, service_class	Delays report for all jobs or specified job groups	DEL, DLY, DL
DELAYJ	jobname	Job report variation for specified job reflecting primary delay reason	DLJ, DJ, DELJ, DLYJ, JOB, JO
DEV	job_class, service_class	Device delays report for all jobs or specified job groups	DD, DVD
DEVJ	jobname	Device delays variation of job report for specified jobname	DDJ, DVJ
DEVR	volser	Device delays report for all or specified resources	DR, DVR
DSINDEX		Data index information	DS, DI
DSND	dsname	Data set delays report for all or specified data sets	DSN

Table 15. Report Commands (continued)

Command	Parameters	Displays	Abbreviation
DSNJ	jobname	Data set delays - Job report for specified jobname	DSJ
DSNV	volser	Data set delays - Volume report for specified volume	DSV
ENCLAVE	sstype	Enclave activity report	ENCL
ENQ	job_class, service_class	Enqueue delays report for all jobs or specified job groups	ED
ENQJ	jobname	Enqueue delays variation of job report for specified jobname	EJ
ENQR	resource	Enqueue delays for all or specified resources	ER
GROUP	s/r-class, period	Group response time breakdown	GP, GRP, GD, RT, GRT
HSM	job_class, service_class	HSM delays report for all jobs or specified job groups	HD
HSMJ	jobname	HSM delays variation of job report for specified jobname	HJ
IOQUEUE		I/O queuing activity report	IOQ, IQ
JES	job_class, service_class	JES delays report for all jobs or specified job groups	JD
JESJ	jobname	JES delays variation of job report for specified jobname	JJ
JOB	jobname	Job report variation for specified job reflecting primary delay reason	JO, DELAYJ, DLYJ, DELJ, DLJ, DJ
LOCKSP	HELD SPIN <u>BOTH</u>	Spin Lock Report about held spin locks and/or address spaces spinning due to a request for a spin lock	LSP
LOCKSU	LOCAL GLOBAL <u>BOTH</u>	Suspend Lock Report about local and/or global suspend locks	LSU
MNTJ	jobname	Operator delays variation for mount request of job report for specified jobname	MTJ
MSGJ	jobname	Operator delays variation for message request of job report for specified jobname	MSJ
OPD		OMVS process data	
PROC	job_class, service_class	Processor delays report for all jobs or specified job groups	PD
PROCJ	jobname	Processor delays variation of job report for specified job	PJ
PROCU	job_class, service_class	Processor usage of a job per processor type (standard or special purpose processors)	PU
QSCJ	jobname	Operator delays variation for quiesce command of job report for specified jobname	QJ
RLSDS	dsname	VSAM RLS activity by data set	RLD
RLSLRU		VSAM LRU overview	RLL
RLSSC	storage_class	VSAM RLS activity by storage class	RLS
SPACED		Disk space report	SPD
SPACEG		Storage space report	SPG
STOR	job_class, service_class	Storage delays report for all jobs or specified job group	SD
STORC	job_class, service_class	Common storage report	SC
STORCR		Common storage remaining at end of job report	SCR
STORF	job_class, service_class	Detailed information on frame counts for all jobs or specified job group	SF

Table 15. Report Commands (continued)

Command	Parameters	Displays	Abbreviation
STORM	job_class, service_class	Detailed information about the use of memory objects within the system	SM
STORJ	jobname	Storage delays variation of job report for specified job	SJ
STORR		Storage space and paging activity report for all system volumes	SR
STORS	wlm	Summarized storage information by workload group, service or report class	SS
SYSENQ		Sysplex enqueue delays report	ES
SYSINFO	wlm	System information, total and by user groups	SY, SYS, SI
SYSRTD	s/r-class, period	Response time distribution report	RTD
SYSSUM	wlm	Sysplex summary	SUM
SYSWKM	s/r-class, period	Work manager delays report for subsystems	WKM
WFEX		Workflow/exceptions screen	WE, WF
XCF	job_class, service_class	Cross-system coupling facility delays report	XD
XCFJ	jobname	XCF delays variation of the job report for specified jobname	XJ
ZFSACT		zFS file system activity	ZFSA
ZFSSUM		zFS file system summary	ZFSS

Table 16 contains commands for the examples of user-written reports that were delivered with RMF.

Table 16. User-Written Report Commands

Command	Parameters	Displays	Abbreviation
DEVN		Device activity	DA
DEVT		Device trend	DT
DSD		Detailed storage delays	
RG		Resource group data	
SYSTREND		System trend	ST

Chapter 14. Snapshot reporting with Monitor II

This information unit guides you in using the Monitor II sessions:

- ISPF sessions
- TSO/E sessions
- background sessions
- session commands and report commands

Monitor II sessions

There are three types of Monitor II report sessions:

- ISPF session (page “The ISPF session” on page 168)

Start this session with the command:

RMF

This leads to the RMF Primary menu, where you select **2** to start the Monitor II ISPF session.

- TSO/E session (page “The TSO/E session” on page 171)

Start this session with the TSO/E command:

RMFMON

- Background session (page “The background session” on page 172)

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

ISPF sessions and TSO/E sessions are also referred to as *Display sessions*, in contrast to *background session*.

You can obtain a **printout** of a Monitor II report:

- during a display session
- during or at the end of a background session

You can get the same reports in all sessions using a similar syntax:

- Display session

You call the reports using *commands* that conform to TSO/E syntax rules:

Example: ASD T,A

In an ISPF session, you can also select the reports from a menu.

- Background session

You call the reports using *options* that conform to option syntax rules:

Example: ASD(T,A)

Note: Starting with z/OS V1.2 RMF, there is no longer a local 3270 display session. However, you can have access to Monitor II reports without an active TSO/TCAS subsystem by means of the RMF Client/Server Enabling (RMFCS); see Chapter 21, “RMF Client/Server enabling (RMFCS),” on page 373.

The ISPF session

Enablement: RMF is an optional feature 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. The Monitor II session cannot be started, and you will receive the message:

ERBA000I RMF is not enabled to run on this system

When you select *Monitor II* on the *RMF Primary Menu*, you get the *Monitor II Primary Menu*. You can go from here to the category of report that you want to display, or you can choose the tutorial or exit from Monitor II. You can also enter Monitor II report commands on the selection line.

RMF Monitor II Primary Menu

z/OS V2R1 RMF

Selection ==>

Enter selection number or command on selection line.

1 Address spaces

2 I/O Subsystem

3 Resource

L Library Lists

U User

Address space reports

I/O Queuing, Device, Channel, and HFS reports

Enqueue, Storage, SRM, and other resource reports

Program library and OPT information

User-written reports (add your own ...)

T TUTORIAL

X EXIT

5650-ZOS Copyright IBM Corp. 1994, 2013.

Licensed Materials - Property of IBM

Figure 26. Monitor II Primary Menu

The selection **U** displays the *RMF Monitor II User Reports* selection menu. This option is only meaningful if you have written some reports of your own and included them into the selection menu. For information on how to do this, see the *z/OS RMF Programmer's Guide*.

The Monitor II tutorial is available from the "RMF Primary Menu" as well as the "Monitor II Primary Menu". For information about how to use the Monitor II tutorial, see "Getting help with RMF dialogs" on page 130.

When selected, each category of Monitor II presents a more detailed selection panel showing the individual reports. The categories are:

- Address space reports
- I/O queuing, device, channel, and HFS reports
- Enqueue, storage, SRM, and other resources reports
- Program library and OPT information - allows you to check whether the status of the program libraries or the settings of the OPT parameters are correct for your current environment
- User-written reports.

Address-space reports

This panel lets you choose what you want to know about address-space activity. The reports offered in the lower part of the panel present information by job name,

so if you select one of them, you must enter the appropriate jobname in the “Options” panel that corresponds to the chosen report.

Here is what the panel looks like:

```
RMF Monitor II Address Space Report Selection Menu
Selection ==>

Enter selection number or command on selection line.

1 ARD          Address space resource data
2 ASD          Address space state data
3 ASRM         Address space SRM data

4 ARDJ         Address space resource data by jobname
5 ASDJ         Address space state data by jobname
6 ASRMJ        Address space SRM data by jobname
```

Figure 27. Monitor II Address Space Report Selection Menu

Instead of making a selection, you can enter any Monitor II report command at the selection prompt.

The next panel to appear is the report panel you have chosen.

I/O Queuing, Device, Channel, and HFS reports

From this panel, you can choose whether you want information about channel path activity, about I/O queuing activity, about device activity, or about hierarchical file systems of the UNIX System Services.

Here is what it looks like:

```
RMF Monitor II I/O Report Selection Menu
Selection ==>

Enter selection number or command on selection line.

1 CHANNEL      Channel path activity
2 IOQUEUE      I/O queuing activity

3 DEV          I/O device activity
4 DEVV         I/O device activity by volume or number

5 HFS          Hierarchical file system statistics
```

Figure 28. Monitor II I/O Report Selection Menu

Instead of making a selection, you can enter any Monitor II report command at the selection prompt.

Enqueue, storage, and SRM reports

For your better orientation, the choices in this panel have been divided into:

- Enqueue activity reports, at the top of the panel
- Storage and System Resource Management-related reports, at the bottom

```

RMF Monitor II Resource Report Selection Menu
Selection ==>

Enter selection number or command on selection line.

Enqueue

1 SENQ          Enqueue activity
2 SENQR         Enqueue reserve activity

Storage and SRM

3 PGSP          Page data set activity
4 SPAG          System paging activity
5 SRCS          Central storage / processor / SRM

Other Resources

8 SDS           RMF Sysplex Data Server activity
9 ILOCK         IRLM Long Lock detection

```

Figure 29. Monitor II Resource Report Selection Menu

Instead of making a selection, you can enter any Monitor II report command at the selection prompt.

Program library and OPT information

This panel appears when you select option **L** in the Monitor II Primary Menu. The available types of library lists and the active OPT member are presented for selection.

```

RMF Monitor II Library List and OPT Settings Selection Menu

Enter selection number or command on selection line.

1 Link list      LNKLSTxx - Link Library list          (LLI)
2 LPA list       LPALSTxx - LPA Library List           (LLI LPA)
3 APF list       IEAAPFxx - Authorized Program List     (LLI APF)

4 OPT            IEAOPTxx - OPT Settings                (OPT)

Selection ==>

```

Figure 30. Monitor II Library List and OPT Settings Selection Menu

The Monitor II commands that are executed for each selection are shown in parentheses to the right of the selections. You can enter one of these at the selection prompt, if you prefer.

For the library lists, entering the command rather than making a selection, allows you to alter the scope of the resulting list. You can then specify the operand **A**, which causes device-type and serial-number information to be included in the list. This information is suppressed by default, because the necessary processing is time-consuming and only justifiable when you really want it. You will find a detailed description of the LLI command in “LLI” on page 194.

User Reports

This panel appears in response to the choice **U** in the Primary menu. The names of your reports are presented for selection, if you formatted the panel as described in

the z/OS RMF Programmer's Guide:

```
RMF Monitor II User Reports
Selection ==>

Enter selection number or command on selection line.

1 USER          User report 1
```

Figure 31. Monitor II User Report Selection Menu

Here, too, you can enter any Monitor II report commands after the selection prompt, instead of making a selection.

The TSO/E session

Enablement:

RMF is an optional feature 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. The Monitor II reporter session will not start, and you will receive the message:

ERB111I RMF IS NOT ENABLED TO RUN ON THIS SYSTEM

When you start Monitor II using the command RMFMON, Figure 32 is the panel you see first.

RMF DISPLAY MENU		
NAME	PFK#	DESCRIPTION
ARD	1	ADDRESS SPACE RESOURCE DATA
ASD	2	ADDRESS SPACE STATE DATA
ASRM	3	ADDRESS SPACE SRM DATA
CHANNEL	4	CHANNEL PATH DISPLAY
DEV	6	SYSTEM DEVICE DATA
PGSP	7	SYSTEM PAGING SPACE DATA
SENQ	8	SYSTEM ENQUEUE CONTENTION
SENQR	9	SYSTEM ENQUEUE RESERVE
SPAG	10	PAGING DATA
SRCS	11	CENTRAL STORAGE / CPU / SRM DATA
ARDJ		RESOURCE DATA FOR SPECIFIC JOBNAME
ASDJ		STATE DATA FOR SPECIFIC JOBNAME
ASRMJ		SRM DATA FOR SPECIFIC JOBNAME
DEVV		SYSTEM DEVICE DATA FOR A SPECIFIC VOL/NUMBER
IOQUEUE		I/O QUEUING ACTIVITY DATA
SDS		RMF SYSPLEX DATA SERVER
LLI		PROGRAM LIBRARY INFORMATION
ILOCK		IRLM LONG LOCK DETECTION
OPT		SETTINGS
USER		USER PICTURE

Figure 32. Monitor II Display Menu

You can enter session commands, or you can select a specific report by entering the report command name shown in the NAME column, or by pressing the corresponding PF key, shown in the PFK# column.

Issue all commands from the **input area**. This is where the cursor appears when you begin a session.

The background session

For a background session, the definition of all session and report options is done either with the appropriate Parmlib member (default member ERBRMF01), or with additional options that the operator can specify in a START or MODIFY command. For details, refer to “Starting a specific Monitor” on page 52.

Structure of Monitor II reports

This chapter provides some general information about Monitor II reports.

- Monitor II reports can have different formats:
 - Table reports
 - Row reports
- Monitor II can also have different report modes:
 - Total mode
 - Delta mode
- ISPF sessions and TSO/E sessions have different report headers.

Furthermore, the different ways to display and modify report options depending on the session type are explained in “Display and set options” on page 177.

Display session report fields

A Monitor II report header looks different, depending on whether you use the ISPF interface or the TSO/E interface.

If you are using an ISPF session

Each report consists of:

- a header line identifying the report
- a status line for CPU, UIC, PR and System. For a description of these fields, refer to Table 17 on page 173.

This line also contains the current setting of the report mode (Delta/Total).

- a variable number of report data lines.

If you are using the TSO/E session

Each report consists of:

- a title line
- two lines of heading information
- a variable number of report data lines.




Figure 33. CHANNEL report as example of a Monitor II report header (TSO/E session)

Table 17. Monitor II Display Session Areas

Area	Function
Report title	The type of measurement data
F	Indicates more pages
Input Area	Issue all commands from here. Separate commands from the MIG=xxx by at least one blank space, or use all 32 spaces, otherwise you get a syntax error.
Status/Message Area	This area contains: CPU Current average processor utilization. This information depends on the activity of Monitor I. If Monitor I CPU gathering is active, the header line shows two views separated by a slash (/): - The MVS view of the CPU utilization - The LPAR view of the CPU utilization If Monitor I CPU gathering is not active, the header line shows: - The SRM view of the CPU utilization. '***' due to missing CPU measurement data for the LPAR view UIC The current system unreferenced interval count. Values greater than 9999 are displayed as nnK to indicate a multiple of 1000. The maximum value is 65K. PR The rate of page-ins per second excluding swap-ins, VIO (virtual input/output), reclaims, and hiperspaces. System The SMF identifier associated with this system.
Report Name Area	The report name.
Mode Area	The current setting for the delta mode (either D for delta or T for total) and hardcopy mode (either H for hardcopy, or blank)
Header Area	Consists of two lines of column headings that identify the data fields included in the report.
Output Area	Contains the report data.

When you begin a session, the cursor appears in the **input area**. During the session you issue all display commands from this area. Other areas indicated in the figure are described in Table 17 on page 173.

Different formats of Monitor II reports

Monitor II offers two types of reports:

- **Table Reports** - Example: ASD Report

Table reports have a variable number of data lines.

- **Row Reports** - Example: ASDJ Report

Row reports have only one line of data. When you request a row report repeatedly, each request adds one line of data to the display. You can use the repetitive requests to build a table of information.

Note: The current line might not be displayed on the screen if you have selected the ISPF option PFSHOW ON or if you are in split-screen mode. You can get the line either by issuing the command PFSHOW OFF or by appropriate scrolling.

Different modes of Monitor II reports

Monitor II offers two modes for the session reports:

- **Total mode**

A total mode report shows the cumulative total since the beginning of the Monitor I interval.

- **Delta mode**

A delta mode report shows the change in the activity since the previous request for the report.

Delta Mode Report

A delta report reflects changes in the activity shown in any report type. Its reporting interval is the time between two consecutive Monitor II requests.

To enter delta mode, type D on the command line, and press ENTER. This establishes the base for reporting, but **does not** request a report.

To request the first delta report, press ENTER again. The reporting interval is the time between the last total report and this first delta report, and the data reflects the change in activity within this interval. If no reports of this type have yet been requested in the current Monitor II session, the first delta report shows null values ('--') in the measurement columns.

To request further delta reports, just press ENTER each time. In each subsequent report in delta mode, the data reflects the change in activity since the previous report.

If a Monitor II interval expires between two consecutive Monitor II requests, no data is reported, and RMF prompts you to press ENTER.

To return to total mode, enter the command D OFF on the command line.

Monitor II session commands and options - Overview

Table 18 provides an overview of the Monitor II display session commands.

Display session commands

Table 18. Monitor II Display session commands

Task	ISPF Command	Parm	TSO/E Command	Parm	Result
"Display the menu" on page 177	RETURN (PF4)		M		Returns to the Primary menu.
"Display and set options" on page 177	RO		MM		ISPF Displays the Report Options panel for the current report TSO/E Displays the report option defaults for both the gatherer and reporter, and the current PF key assignments.
"Reset default options" on page 179	RESET				ISPF On Report Option panels, resets all optional values to the defaults specified in the menu ERBFMENU
"Leave options unchanged" on page 179	CANCEL				ISPF On Report Option panels, ends the option dialog without making any changes
"Display commands" on page 179	COMMANDS				ISPF Shows you all the available commands
"Scroll through report frames" on page 180	PF8/PF7		F		ISPF Forward/backward scrolling TSO/E Forward scrolling
"Recall the previous data" on page 181	Rrep	opts	Rrep	opts	Recalls the previous report, where rep is the report name and opts are any options for the report
"Sort reports" on page 181	SORT (PF6)	A D			ISPF Sorts the report by the column in which the cursor is located. Ascending order Descending order
"Find text string" on page 181	FIND	string			ISPF Searches for a text string in a report
"Repeat Find" on page 182	RFIND (PF5)				ISPF Repeats a previously entered FIND command
"Set delta mode" on page 182	D	ON OFF	D	ON OFF	Sets the DELTA mode Sets the TOTAL mode
"Create a hardcopy report" on page 182	H	ON OFF	H	ON OFF	Prints all displayed reports Prints no reports
"Refresh a report automatically" on page 184	GO	n	T	m,n	ISPF Causes an automatic refresh of the report data every n seconds TSO/E Updates a report automatically, where m is the number of times you want to update the report, and n is the number of seconds between updates. Specify this command after requesting a report.

Table 18. Monitor II Display session commands (continued)

Task	ISPF Command	Parm	TSO/E Command	Parm	Result
"Print a report page" on page 186	PRINT		P		ISPF Writes the currently displayed screen to the ISPF list data set TSO/E Writes the currently displayed report to the preallocated report data set
"Specify the system to be monitored" on page 186	SYS[TEM]	smf_id	SYS[TEM]	smf_id	Identifies system to be monitored
"Assign PF keys" on page 185	KEYS		#rep	opts	ISPF Standard ISPF key assignment TSO/E Assigns PF keys where rep is the report name and opts are any options for the report. You must press the PF key you want assigned to that report after entering the command.
"Stop the session" on page 187	=X		END		Stops the session In TSO/E, you can also use the Z, QUIT, QQ, X, or STOP command to stop the session.

Background-session options

Table 19. Monitor II Background Session Options

Task	Background Session Option	Result
"Refresh a report automatically" on page 184	<u>SINTV</u> { (30S) } { (value[S]) }	Specifies number of seconds in each measurement interval.
"Stop the session" on page 187	<u>STOP</u> {value {M}} {value {H}}/NOSTOP	Desired duration of the Monitor II session, in minutes (M), or hours (H).
"Set delta mode" on page 182	<u>DELTA</u> / <u>NODELTA</u>	Specifies whether RMF should report total values or values that reflect changes since the previous measurement.
"Write SMF records" on page 186	<u>RECORD</u> /NORECORD	Specifies whether measured data is to be written to SMF records.
"Create a hardcopy report" on page 182	<u>REPORT</u> { (REALTIME) } / (NOREPORT) { (DEFER) }	Specifies production of printed interval reports of measured data.
"Create a hardcopy report" on page 182	SYSOUT(class)	SYSOUT class to which the formatted printed reports are directed.
"Define session options" on page 180	MEMBER (list)	Parmlib member, or list of members, containing Monitor II background session options.
"Display and set options" on page 177	{OPTIONS} / {NOOPTIONS} {OPTN } {NOOPTN }	Print an options list at the operator console at the start of the session.

Monitor II session tasks

This topic provides information about the tasks that you can perform during Monitor II sessions, like for example, defining session options, sorting reports, writing SMF records or creating hardcopy reports.

Display the menu

For an ISPF session

Press PF4 or enter the RETURN command to return to the Primary Menu (see Figure 26 on page 168).

For a TSO/E session

To display the menu of available reports, issue the menu command:

M

Figure 32 on page 171 shows the menu panel. The menu lists each report name, its PF key assignment, and a description.

From the display menu, you can display the first report in the menu with defaults by pressing ENTER.

Note: If you assigned a different PF key to the first report in the display menu, and you press ENTER while the report field is blank, RMF displays the original report assigned to the PF key.

Display and set options

For an ISPF session

In an ISPF session, you select the report you want rather than specifying it in a command. When you enter a command for certain report types, you can specify options as part of the commands.

The options are remembered from one session to the next. The options used for the first session are the RMF defaults, but you can alter these in the Report Options panel for the respective report.

Call up the appropriate Report Options panel by entering the RO command at the command prompt of the report panel.

Figure 34 presents an example of an Report Options panel:

RMF Monitor II - Address Space Options

Command ===>

Change or verify parameters. The input entered on this panel applies to ARD, ASD, and ASRM. To exit press END.

Class	===> T	Specify one of the following workloads: A=All, B=Batch/STC, T=TSO, AS=ASCH, O=OMVS
Inactive	===> NO	Specify YES to include inactive address spaces.

Figure 34. ARD, ASD and ASRM Report Options Panel

Type the options you want, and press ENTER. If an option is invalid for the report, RMF will issue a message telling you this, and leave the option panel on the screen for you to correct your input. If no message is issued, the values you entered are valid, and you can enter the END command or press PF3 to have them accepted.

For a TSO/E session

You get an overview about the default options for all commands by entering the command:

The standard definition of the default options is shown in this figure:

RMF DISPLAY MENU			
NAME	PFK#	DEFAULT OPERANDS FOR GATHERER AND REPORTER	
ARD	1	A,A,	--- A,I,
ASD	2	A,A,	--- A,I,
ASRM	3	A,A,	--- A,I,
CHANNEL	4		
DEV	6	DASD	
PGSP	7	PAGE	
SENQ	8	D --- S	
SENQR	9	ALLVSR	
SPAG	10		
SRCS	11		
ARDJ			
ASDJ			
ASRMJ			
DEVV			
SDS			
IOQUEUE		DASD	
LLI			
ILOCK			
OPT			
USER			

Figure 35. Monitor II TSO/E session - Default Options

The data gatherer and reporter options are separated by three dashes (---). If both the options default are the same, only one set of options is displayed. If there are no default options for the data gatherer, the reporter options are displayed preceded by three dashes (---).

For a background session

To display the current options during start of a background session, either specify **OPTIONS**

in the Parmlib member (for example ERBRMF01), or in the **START** command. Then, 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.

Table 20 shows each possible option followed by its input source.

Table 20. Monitor II **OPTIONS** Command Sources

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

Table 20. Monitor II *OPTIONS* Command Sources (continued)

Source	Where Option is Specified
REPLY	The option was changed from the operator console in reply to message ERB306I.

The following is an example of the console output produced when *OPTIONS* is in effect.

```

ERB103I LS : OPTIONS IN EFFECT
ERB103I LS : NOCHANNEL -- DEFAULT
ERB103I LS : NOPGSP -- DEFAULT
ERB103I LS : NODEVV -- DEFAULT
ERB103I LS : NODEV -- DEFAULT
ERB103I LS : NOSENQR -- DEFAULT
ERB103I LS : NOSENQ -- DEFAULT
ERB103I LS : NOASRMJ -- DEFAULT
ERB103I LS : NOASRM -- DEFAULT
ERB103I LS : NOARD -- DEFAULT
ERB103I LS : NOSRCS -- DEFAULT
ERB103I LS : NOSPAG -- DEFAULT
ERB103I LS : NOARDJ -- DEFAULT
ERB103I LS : NOASDJ -- DEFAULT
ERB103I LS : NOIOQUEUE -- DEFAULT
ERB103I LS : SYSOUT(A) -- MEMBER
ERB103I LS : OPTIONS -- MEMBER
ERB103I LS : REPORT(DEFER) -- MEMBER
ERB103I LS : RECORD -- MEMBER
ERB103I LS : STOP(30M) -- MEMBER
ERB103I LS : SINTV(30S) -- MEMBER
ERB103I LS : NODELTA -- MEMBER
ERB103I LS : NOUSER -- MEMBER
ERB103I LS : ASD -- MEMBER

```

Reset default options

For an ISPF session only

On the command line of any Report Options panels, you can enter the command:

```
RESET
```

This causes RMF to reset all the optional values available for the corresponding report to the those specified in the picture macro. These options take effect for the rest of the session.

Leave options unchanged

For an ISPF session only

On the command line of any Report Options panel, you can enter the command:

```
CANCEL
```

This causes RMF to continue the session without any changes to the options. You can use this command when you have inadvertently misspelled an option, or have decided not to specify one that you have typed in after all.

Display commands

For an ISPF session only

To display all the commands that are available, enter the command:

```
COMMANDS
```

Define session options

For a background session only

You can define whether other Parmlib members with Monitor II options should be used for the session, either using the START command or as part of the options in the ERBRMF01 Parmlib member.



This specifies the Parmlib member(s) — you can specify up to five members — that contain Monitor II background options for the session, where (— xx —) contains from one to five members, separated by commas. Each member must be a two-character alphameric value, which RMF adds to the ERBRMF prefix to form the member name. For the Monitor II background session, the default is 01 indicating Parmlib member ERBRMF01. If you want to use your own Parmlib members, make sure you specified your Parmlib data set on the IEF RDER DD statement in the RMF cataloged procedure. See “Setting up the RMF control session including Monitor I and Monitor II” on page 22.

For more information on Parmlib members, including the contents of the Monitor II member ERBRMF01, see “Storing gatherer options” on page 33.

Scroll through report frames

For an ISPF session

Scrolling through ISPF session panels is achieved in the usual ISPF manner by using PF7 (Backward) and PF8 (Forward). The indication Line x of y at the top right corner of the panel tells you where you are in the report, and how many lines there are. The prompt SCROLL ==> shows you the current scroll amount, and you can change the scroll amount by altering the value, as in other ISPF panels.

For a TSO/E session

To scroll through a multi-frame table report (a report that has more than the maximum number of lines for your device), use the frame command:

F

When RMF displays the first frame of a multi-frame report, a frame command (F) automatically appears in the input area. After inspecting the data in the current frame, press ENTER to see the next frame. Continue the process until you have seen all of the data that you require. If you decide at any point that you do not need to see all of the frames in a report, blank out the frame command or issue a new command. When RMF displays the last frame in the report, the input area is blank.

If you enter F when there are no subsequent frames, RMF ignores the command.

For example, if you are using a terminal with 24 output lines, an F appears in the input area and the first 21 lines of data appear in the output area. The F indicates that you are viewing a multi-frame report and should enter the frame command (F) to view the next frame of output data. Because the input area already contains an F, you can view the next frame by pressing ENTER. The F continues to appear

in the input area until all frames of data have been viewed. When the last frame is displayed, the end of the report is indicated by a blank input area.

Recall the previous data

For ISPF and TSO/E sessions

To cause the most recently displayed set of data (either a full table for a table report or a single line for a row report) to be displayed again, use the recall command. The syntax is:

Rrep [options]

where **rep** is the report command and **options** are the options for the report. Do not leave any blanks between R and the report command name.

Example: You have requested the ASD report for all address spaces with the command:

ASD

Now, you would recall the report for TSO/E address spaces only:

RASD T

Sort reports

For an ISPF session only

On the command line of most report panels, you can enter the command:



Before you press ENTER, place the cursor in one of the columns of the report. When you press ENTER, Monitor II will sort the lines of the report by the contents of this column.

This handling is easier if you use PF6, which is defined as SORT command.

You can sort the report in ascending or descending order. If you do not specify the sort order, then columns with numerical values will be sorted in descending order, and columns with character values will be sorted in ascending order.

Find text string

For an ISPF session only

To find a character string in the report, you can enter the command:

FIND textstring

If the string contains blanks, you must enclose it in quotes.

Repeat Find

For an ISPF session only

You can repeat a previous FIND command using:

RFIND (PF5)

Set delta mode

For ISPF and TSO/E sessions

To set the delta mode for the session, use the DELTA command. The syntax is:



When delta mode is in effect, certain fields in some reports, such as the processor (CPU) time in the ARD report, reflect values that show the change since the previous invocation of the report. The first request for the report will show the value RMF detects at that time; all subsequent invocations of the report will show only the change since the previous report.

Delta mode is set off when the session begins. You must enter D ON or D to set delta mode on for the session. Later, if you want to turn off delta mode, enter D OFF. All report fields that can reflect either session or delta values will then reflect session totals rather than changes.

The mode area indicates the current setting of the delta mode for the session (either Delta/Total for an ISPF session, or D/T for a TSO/E session).

For a background session

To set the delta mode for the session, use the DELTA command:



Note: DELTA/NODELTA is a *reporting* option and has no impact on SMF recording in a background session.

Create a hardcopy report

For ISPF and TSO/E sessions

To create a report from the current session, you set the hardcopy mode on, using the hardcopy command:



When hardcopy mode is in effect, all data in all reports requested during the session is written to a preallocated output data set.

Data Set Allocation for ISPF and TSO/E Sessions

- Allocate the output data set before you start the Monitor II session:

```
ALLOC F(RMFDMSO) DS(data.set.name) SHR
```

where **data.set.name** is the name of the data set to be used for the hardcopy output. Any existing contents of the data set are overwritten.

- To add output from this session to existing data, use the command:

```
ALLOC F(RMFDMSO) DS(data.set.name) MOD
```


Without any data set allocation, the output will be written to SYSOUT class A.

A single output data set is created for all print command (TSO/E session only) and hardcopy command output generated during a single session. Because reports requested might be multi-frame reports and you might choose not to scroll through all of the data during the display session, the data written to the output data set when hardcopy mode is in effect can be more extensive than that displayed on the screen.

Hardcopy mode is turned off when the session begins. You must enter H ON or H to set hardcopy mode on. The command takes effect with the next report you enter.

For a background session

You define with the option



whether or not printed interval reports of the measured data are to be produced. When reports are to be produced (REPORT specified), the REALTIME or DEFER option indicates when the reports are to be printed.

When you omit the option, the default is REPORT(DEFER). If you specify REPORT, you must specify either REALTIME or DEFER; otherwise you get a syntax error.

REALTIME indicates that the reports are to be printed at the end of the session, and when you modify session options for one of the following reasons:

- To end a request for a particular report
- To end a request for all reports
- To replace REPORT(REALTIME) with REPORT(DEFER)

Example: For example, assume that the options ASRM, SPAG, and REPORT(REALTIME) are in effect for an active session. If you end the request for the system paging report by replacing SPAG with NOSPAGE, any accumulated paging reports will be printed. If you change REPORT(REALTIME) to REPORT(DEFER) or NOREPORT, all accumulated reports will be printed.

DEFER indicates that the reports are to be printed after you stop RMF.

You can allocate data sets for the reports in the start-up JCL for the background session, or you can route the output to a SYSOUT class.

Data Set Allocation

You define one or more JCL statements

```
//RMFxxxx DD DSNAME=data.set.name, DISP=disp
```

xx is the session identifier, and **nnn** is a decimal number from 001 to 999.

Without this pre-allocation the output will be routed to SYSOUT. You can define the output class using the option



Class A is the default. You cannot modify the SYSOUT option during session processing.

Refresh a report automatically

For an ISPF session

To refresh a report automatically, enter on any report panel the command:




where **n** is a decimal integer, and **4** is the default. This causes the report to be refreshed automatically every **n** seconds.

To stop automatic refresh of the report, press the ATTN or PA1 key.

For a TSO/E session

To update reports automatically, use the timed update command. The syntax is:



where **m** is the number of times you want to refresh the report and **n** is the number of seconds between refreshes. You can specify a maximum value of 99 for **m** and a maximum value of 999 for **n**. Defaults are:

- 10 for **m**
- 4 for **n**

That is, RMF updates the report automatically 10 times at 4-second intervals.

To stop the timed update, press the attention (PA1) key.

When you issue a timed update command (T), RMF displays the length of the time interval, and the number of intervals remaining in the **input area**.

For a background session

To define the length of a measurement interval, specify:



value

specifies the number of seconds in each measurement interval. The range is from 1 to 3600 seconds. The default is 30 seconds. The S is not required, but you can include it as a reminder that the specification is in seconds.

When you specify a small SINTV value, RMF overhead increases and excessive swapping can result. In such a case, you could make the RMF address space non-swappable to minimize this overhead.

Note: When you change either the SINTV or the STOP options during the session, the duration of the session could be affected. See “Conflicting session options” on page 198.

Assign PF keys

For an ISPF session

You can use the standard ISPF capability of assigning functions to PF keys by calling the ISPF command KEYS. Most PF keys in Monitor II have the standard setting as in all other ISPF applications. There are two exceptions:

PF5 RFIND command

PF6 SORT command

If you want to make your own settings, you might use PF10 and PF11 (or PF13 — PF24), they have no predefined meaning in an ISPF Monitor II session. All PF key definitions remain valid across sessions.

For a TSO/E session

You can assign PF keys to arbitrary reports for the duration of a Monitor II display session.

RMF assigns the first 12 program function (PF) keys as shown in Figure 32 on page 171. To override these assignments, enter a pound sign (#) in the input area, followed by a report name, and any report options. Then press the PF key you want to assign to the report. The assignment remains for the duration of the Monitor II display session. RMF displays the default options menu after each # command so that you can confirm your PF key assignments.

Example: To associate PF 10 with the device report for devices with numbers 0150 through 0350 and device 0370 you specify:

```
#DEV NUMBER(0150:0350,0370)
```

and press PF 10. To request the report, press PF 10. For the duration of your session, PF 10 is associated with this report.

For a TSO/E session, the attention key (PA1) can be used in the same way as for any TSO/E command, as long as the timed update command is not active.

Print a report page

For an ISPF session

To print one page of a report, enter the ISPF command:

PRINT

This writes a copy of the currently-displayed screen contents to the ISPF list data set, which you can print or display after the session.

For a TSO/E session

To print the data currently displayed on the screen, use the print command:

P

See “Create a hardcopy report” on page 182 for more information about the output data set.

General remark

Because both commands cause only the current screen image to be printed, you would have to use repetitive scroll commands and print commands to print all of the data in a multi-frame table report.

Use this command when you want to print a single report and the session is not in hardcopy mode. In hardcopy mode, the entire report is automatically printed, and you would not need to use the print command. You find details in “Create a hardcopy report” on page 182.

Write SMF records

For a background session only

You can define whether SMF records should be written during a background session. You do this using the option:

►►

RECORD
NORECORD

 ◀◀

Specify the system to be monitored

For all display sessions

You can specify which individual system in a sysplex you want a report to refer to. This may be the system you are using to run your Monitor II session, or another system. You do this using the option:

►► SYSTEM—*smf_id* ◀◀

where *smf_id* is the identifier of the system for which you want the reports to be generated.

For an ISPF session

You can use the SYSTEM command, or, alternatively, you can overwrite the value of the SYSTEM field in the header of the report panel with the identifier of the desired system.

Stop the session

For an ISPF session

To stop the session, enter the ISPF skip command on the command or selection line of any panel:

=X

For a TSO/E session

To end the session, enter:



For a background session

You can specify the duration of a background session in minutes (M) or hours (H) with the option



The range is from one minute to one week (168 hours or 10,080 minutes). The default value is ten minutes. If you specify a value outside the range, RMF substitutes the default value. If neither M nor H is specified, M (minutes) is assumed. NOSTOP means the session can be ended only by a STOP command.

Note:

1. You can stop a session at any time with the session STOP command regardless of the value specified for this option.
2. The STOP option applies only to a particular Monitor II background session. RMF remains active until the operator issues a STOP system command.
3. When you change either the SINTV or the STOP options during the session, the duration of the session could be affected. See “Conflicting session options” on page 198.

Monitor II report commands - Overview

Table 21 gives an overview of the available report commands.

Table 21. Monitor II Report Commands

Display Session Syntax	Background Session Syntax	Report
ARD [class,status]	ARD [(class,status)]/NOARD	Address space resource data reporting. See “ARD” on page 189.
ARDJ jobname	ARDJ (jobname)/NOARDJ	Address space resource data reporting for a particular job. See “ARD” on page 189.

Table 21. Monitor II Report Commands (continued)

Display Session Syntax	Background Session Syntax	Report
ASD [class,status]	<u>ASD</u> [(class,status)]/ <u>NOASD</u>	Address space state data reporting. See “ASD” on page 190.
ASDJ jobname	ASDJ(jobname)/ <u>NOASDJ</u>	Address space state data reporting for a particular job. See “ASDJ” on page 190.
ASRM[class,status]	ASRM[(class,status)]/ <u>NOASRM</u>	Address space SRM data reporting. See “ASRM” on page 190.
ASRMJ jobname	ASRMJ(jobname)/ <u>NOASRMJ</u>	Address space SRM data reporting for a particular job. See “ASRMJ” on page 190.
CHANNEL	CHANNEL/ <u>NOCHANNEL</u>	Channel path activity data reporting. See “CHANNEL” on page 191.
DEV [type]	DEV [(type)]/ <u>NODEV</u>	Table reporting on I/O device activity. See “DEV” on page 191.
DEVV {VOLSER(xxxxxx)} {NUMBER(yyyy) }	DEVV{(VOLSER(xxxxxx))}/ <u>NODEVV</u> {(NUMBER(yyyy))}	Row reporting on a specific direct access device. See “DEVV” on page 192.
HFS [hfsname]	--	Table reporting on UNIX hierarchical file system statistics. See “HFS” on page 193.
ILOCK [ALL]	--	IRLM Long Lock detection. See “ILOCK” on page 193.
IOQUEUE [type]	IOQUEUE[(type)]/ <u>NOIOQUEUE</u>	I/O request queuing reporting. See “IOQUEUE” on page 194.
LLI {LNK}{,A} {LPA} {APF}	--	Program library information listing. See “LLI” on page 194.
OPT	--	OPT Settings report. See “OPT” on page 195.
PGSP{PAGE}	PGSP {(PAGE)}/ <u>NOPGSP</u>	Page data set activity reporting. See “PGSP” on page 195.
SDS	--	RMF Sysplex Data Server activity reporting. See “SDS” on page 195.
SENQ {S } {D } {A,sysname } {E,sysname } {majorname[,minorname]}	{(S)} {(D)} SENQ {(A,sysname)} {(E,sysname)} {(majorname[,minorname])} <u>/NOSENQ</u>	Enqueue contention activity reporting. See “SENQ” on page 196.
SENQR {ALLVSR} {volser }	SENQR {(ALLVSR)}/ <u>NOSENQR</u> {(volser)}	Reserve activity reporting. See “SENQR” on page 197.
SPAG	SPAG/ <u>NOSPAG</u>	System paging activity reporting. See “SPAG” on page 198.
SRCS	SRCS/ <u>NOSRCS</u>	Central storage/processor/SRM activity reporting. See “SRCS” on page 198.
USER	USER/ <u>NOUSER</u>	Specifies whether or not a user-specified activity is to be measured. See “USER” on page 198.

Details of report commands

This section describes the Monitor II report commands in alphabetical order. Program defaults are underscored where appropriate.

The same report commands are available for Monitor II display and background sessions, and the command syntax in the both types of sessions is similar, but what you can specify varies from one type of session to the other. For background details refer to Table 21 on page 187.

Display session

Specify commands either in the command or selection line of the ISPF panel or in the input area of a TSO/E panel.

Background session

Specify background options in either one or both of the following:

- The **parm** field of the session command START that you issue to start the session. See “Starting a specific Monitor” on page 52 for the START session command syntax.
- The Monitor II background session Parmlib member ERBRMF01. See the explanation of the background session MEMBER option (“Define session options” on page 180). The contents of ERBRMF01 is described in “Storing gatherer options” on page 33.

RMF uses a program default for any option you do not specify.

Chapter 5, “How RMF processes session options,” on page 59 describes how RMF merges the options from these three sources. You can modify the options during a session as described in “Modifying RMF session options” on page 55.

Note: Some report options (ARDJ, ASDJ and ASRMJ) have “jobname” as a suboption. You can specify only one of these options per session. If you want to monitor several jobs in parallel, you have to start several background sessions.

ARD



Specifies address space resource data reporting, where **class**, and **status** specify the following selection criteria for the address spaces to be included:

class

<u>A</u>	All address spaces
AS	ASCH address spaces
B	batch, started task, and mount task address spaces
O	OMVS address spaces
T	TSO/E address spaces

status

A	All address spaces
<u>I</u>	Active address spaces; that is, those address spaces that are currently executing, non-swappable, or swapped out and eligible for swap-in

The operand fields are positional; if you omit any one, you must replace it with a comma. RMF uses the default value for any omitted operand.

Monitor II background sessions each set separate defaults for their data gathering routine and the reporting routine. For the data gathering routine, the value for **class** is A, and the value for **status** is A. For the data reporting routine, the value

for **class** is A, and the value for **status** is I. This means that data is gathered for all address spaces, but reports are generated only for active ones.

You can change the menu defaults to fit the needs of your particular installation as described in the *z/OS RMF Programmer's Guide*.

Example: For a display session or background session:

- To report on all address spaces, enter
ARD ,A or ARD(,A)
- To report on all OMVS address spaces that are currently active, enter
ARD 0 or ARD(0)

ARDJ

►► ARDJ—*jobname*—————►►

Specifies address space resource data reporting for a particular job. Specify the job you want to measure under **jobname**.

ASD

►► ASD

A
class

 ,

I
status

 —————►►

Specifies address space state data reporting. For a description of the options, see the ARD command.

ASDJ

►► ASDJ—*jobname*—————►►

Specifies address space state data reporting for a particular job. Specify the job you want to measure in **jobname**.

ASRM

►► ASRM

A
class

 ,

I
status

 —————►►

Specifies address space SRM data reporting. For a description of the options, see the ARD command.

ASRMJ

►► ASRMJ—*jobname*—————►►

Specifies address space SRM data reporting for a particular job. Specify the job you want to measure in **jobname**.

CHANNEL

▶▶ CHANNEL —————▶▶

Specifies channel path activity data reporting.

DEV

▶▶ DEV ———▶▶
┌──DASD──┐
└──type──┘

Specifies data reporting for a table report on I/O device activity, where **type** is all devices in one class, or one or more specific device numbers, volume serial numbers, or storage groups. When you specify the DEV option, a Monitor I session must be measuring any device you request.

Parameter **type** can be one of the following:

- A device class:

DASD

Direct access storage devices

TAPE Magnetic tape devices

COMM

Communication equipment

CHRD

Character reader devices

UNIT

Unit record devices

GRAPH

Graphic devices

- One or more volume serial numbers:

```
      ({aaaaaa      })  
{VOLSER} ({aaaaaa,bbbbbb:zzzzz})  
{V      } ({aaaaaa,bbbbbb,..  })
```

VOLSER requests specific DASD or tape devices, where aaaaaa, bbbbbb, and zzzzzz each represent a volume serial number. You can specify a single volume, in the format aaaaaa, a list of volumes, in the format aaaaa,bbbbbb,...., or a range of volumes, in the format bbbbbb:zzzzz where aaaaa is the first and bbbbbb is the last volume. Your entry cannot exceed 32 characters, including commas and colons. When you specify a range, use a colon as a separator to indicate that the report is to include all volumes from aaaaaa up to and including zzzzzz.

- One or more device numbers:

```
      ({aaaa      })  
{NUMBER} ({aaaa,bbbb:zzzz})  
{N      } ({aaaa,bbbb,...  })
```

NUMBER requests specific device numbers, where aaaa, bbbb, and zzzz each represent hexadecimal device numbers. You can omit leading zeros. You can specify any combination of a single device number, in the format aaaa, a list of device 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. Your specification must not exceed 32 characters, including commas and colons.

- One or more storage group names:


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

SG requests specific storage group names, where aaaaaaaa, bbbbbbbb, and zzzzzzzz each represent 1 to 8 character names, found in SMF type 74 and type 79 records for each DASD device managed by the system-managed storage. 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. RMF reports the devices in sequence by device number within the storage groups.

Example for a display session or a background session:

- To request a Device Activity report for all magnetic tape devices, specify:
DEV TAPE or DEV(TAPE)
- To request a Device Activity report for volumes P50002, P50003, P50004, and P50007, specify:
DEV V(P50002:P50004,P50007) or DEV(V(P50002:P50004,P50007))
- To request a Device Activity report for the storage groups MANF13, MANF14, MANF15, MANF16, MANF17, MANF18, MANF19, and MANF20, specify:
DEV SG(MANF13:MANF20) or DEV(SG(MANF13:MANF20))
- To request all storage groups, specify:
DEV SG or DEV(SG)

DEVV



Specifies data reporting for a row report on a specific direct access device (in contrast to the DEV option which allows you to report on more than one device) where:

VOLSER or V

to request I/O device activity for the specific volume identified by the volume serial number xxxxxx.

NUMBER or N

To request I/O device activity for the specific device identified by the number yyyy.

When you specify DEVV, a Monitor I session must be measuring the device you request. A storage group name is reported for any device that is assigned to one.

Example for a display session or background session:

- To request a Device Activity report for tape device number 1580, specify:
DEVV N(1580) or DEVV(N(1580))
- To request a Device Activity report for direct access device DASD01, specify:
DEVV V(DASD01) or DEVV(V(DASD01))


HFS



Specifies reporting of statistics of hierarchical file system of UNIX System services.

hfsname is the name of the file system for which the statistics data are to be obtained.

ILOCK



Services of the **IMS/VS Resource Lock Manager (IRLM)** are used by IMS to serialize application program requests for data base records to ensure that two programs do not access the same record for update at the same time.

The ILOCK report enables you to identify locking situations that are caused by serialization effects when sharing data among several IMS instances in a sysplex.

To display all blocker and waiters, you have to call the ILOCK command with the parameter ALL, otherwise TOP BLOCKERS will be shown, only.

There is no data gathering component for this report. Instead, the retrieval of the IRLM data from the RMF SMF data buffer is done by the reporter. To have the data available in the SMF data buffer (SMF record type 79 subtype 15), it is necessary to specify this option explicitly, for example:

```
S RMF,,, (SMFBUF (RECTYPE(70:78,79(15))))
```

For details, refer to “Controlling the SMF buffer” on page 47.

Data collection is initiated by the operator who enters at the console the **runtimeo-exit** for one system in the sysplex:

```
F irlmid,RUNTIMEO
```


The command will be propagated automatically to all other systems.

When the SMF records are eventually written by the IRLMs in the data sharing group, the reporter can fetch these SMF records out of the RMF SMF data buffer.

As a consequence, you have to ask the operator to issue this command if you get informed that there is no data available for the report.

Note: Access to the SMF data buffer requires appropriate security authorization. Refer to “Controlling access to RMF data for the sysplex data services” on page 20 for details.

IOQUEUE



Requests reporting on I/O queuing. If you specify IOQUEUE, a Monitor I session must be measuring any I/O queuing activity. Parameter **type** can be any one of the following:

- A device class:

DASD

Direct access storage devices

TAPE Magnetic tape devices

COMM

Communication equipment

CHRDR

Character reader devices

UNITR

Unit record devices

GRAPH

Graphic devices

- One or more logical control unit (LCU) numbers:

```
{aaaa      }
{NUMBER} {aaaa,bbbb:zzzz}
{N      } {aaaa,bbbb,... }
```


NUMBER requests specific logical control unit numbers, where aaaa, bbbb, and zzzz each represent hexadecimal device numbers. Leading zeroes can be omitted. You can specify any combination of a single number, a list of numbers, or a range of numbers, as long as your entry does not exceed 32 characters, including commas and colons. When you specify a range of numbers, use a colon as a separator to indicate that the report is to consist of all numbers from bbbb up to and including zzzz.

Note: When your system is running as a guest under VM, RMF cannot gather data for it. In this case, the IOQUEUE report is not available.

Example for a display session or background session:

- To request an I/O Queuing Activity report for LCUs representing all magnetic tape devices, specify:
IOQUEUE TAPE or IOQUEUE(TAPE)
- To request an I/O Queuing Activity report for LCU numbers D, E, F, 4E, and 55, specify:
IOQUEUE N(D:F,4E,55) or IOQUEUE(N(D:F,4E,55))

LLI



This report provides different program library listing. All operands are optional, and have the following meanings:

LNK

List the link library information. This is the default operand, and takes effect if neither LNK, LPA nor APF is specified. The libraries that will be reported on are those whose names are specified in the LNKLSSTxx Parmlib members.

LPA

List information for libraries defined to the link pack area.

APF

List information about authorized programs defined in the link library.

,A All the information on the selected libraries is to be listed, including the device number, device type and volume serial numbers of the devices where they reside.

It is advisable to specify this operand only when you really want to have the device information listed, because the services used to retrieve this information are expensive in terms of performance.

OPT

▶▶—OPT—▶▶

Invokes a report about the active OPT PARMLIB member and the settings of all OPT parameters.

PGSP

▶▶—PGSP—PAGE—▶▶

Requests that data on page data set activity should be reported, where:

PAGE Indicates page data set activity


A Monitor I session monitoring page activity must be active when you specify the PGSP option.

SDS

▶▶—SDS—▶▶

Requests an RMF Sysplex Data Server activity report to be generated. This report can be generated only if the RMF address space has already been started.

SENQ



Specifies reporting of enqueue contention activity. The operands describe the type of data you require. You can specify only one operand. The meaning of each operand field is:

- S** Specifies a summary report. For each resource that had contention activity, the summary report includes the scope of the resource, the number of tasks that own the resource, the number of tasks waiting for exclusive use of the resource, and the number of tasks waiting for shared use of the resource. If you specify SENQ without operands, S is the default for the reporting routine. For the data gathering routine, the default is D.
- D** Specifies a detail report. For each resource that had enqueue contention, the detail report includes the scope of the resource, the name and address space identifier of each job owning or waiting for the resource, and the type and status of each job's request for the resource.

A,sysname

Specifies a report that includes all resources that a specific system holds in a global resource serialization complex, where *sysname* indicates the system for which the report is requested. Use this report when attempting to recover an inactive processor in a global resource serialization complex. You can request this report from an active processor in the complex and determine from the report the resources that the inactive processor held.

E,sysname

Specifies a report that includes all exclusively-owned resources that a specific system held in a global resource serialization complex, where *sysname* indicates the system for which the report is requested. This report is useful when attempting to recover an inactive processor in a global resource serialization complex. You can request this report from an active processor in the complex and determine from the report the resources that the inactive processor held.

majorname[,minorname]

Specifies a detail report for a specific resource that had contention. The **majorname** field, which corresponds to the **qname** field in the ENQ and DEQ macro instructions, contains the one to eight character major name of a serially-reusable resource. Optionally, the major name is followed by a comma and a minor name. The **minorname** field, which corresponds to the **rname** field in the ENQ and DEQ macro instructions, contains the minor name of the resource.

The maximum length of the field is 32 characters, including the comma. Because the major name is 1 to 8 characters, the minor name can be from 1 to 30 characters, depending on the length of the major name. If you want a report on a minor name, but the **majorname,minorname** operand exceeds 32 characters, you must specify only the major name. RMF then collects data for all resources grouped under the major name.

RMF treats the single character A, D, E, or S as a request for a specific type of report, such as a summary report or a detail report. Therefore, do not use A or E as a major name; do not use S or D as a major name unless you also specify a minor name.

Note: If you intend to run a Postprocessor detail or summary report, keep in mind that if you specify a resource by name or by system on the SENQ option, RMF collects data only for the specified resources. The Postprocessor formats a report containing only the specified resources. For example, if the session option is SENQ(SYSDW), and the Postprocessor option is SENQ(D), the Postprocessor formats a detail report for SYSDW only. Also, if the session option identified a specific resource by name and the Postprocessor option identifies a different specific resource, RMF issues a message to tell you that no data is currently available to meet your selection criteria.

Examples for a display session or background session:

- To obtain a summary report for all resources that have contention, enter:
SENQ
- To obtain a detail report for all resources that have contention, enter:
SENQ D or SENQ(D)
- To obtain a report for all resources that system C303 hold in a global resource serialization complex, enter:
SENQ A,C303 or SENQ(A,C303)
- To obtain a detail report for all resources grouped under the major name of SYSCTLG, enter:
SENQ SYSCTLG or SENQ(SYSCTLG)
- To obtain a detail report for all resources with the major name of SYSI and the minor name of OPENUADS, enter:
SENQ SYSI,OPENUADS or SENQ(SYSI,OPENUADS)

SENQR



Specifies reporting of reserve activity. The operands describe the type of data you require. Only one operand can be specified. The meaning of each operand field is:

ALLVSR

requests data describing all reserve requests. If you specify SENQR without operands, ALLVSR is the default.

volser

requests data describing the reserve requests for a particular device, where **volser** is the one to six character volume serial number of the volume.

Note: If you identify a specific device by specifying a volume serial number on the SENQR session option, RMF collects data only for the device identified. In this case, no data on other devices is available to the Postprocessor, and the Postprocessor cannot produce a reserve activity report for all devices. If, for example, the session option is SENQR(TSO200) and the Postprocessor option is SENQR(ALLVSR), the Postprocessor report includes data only for the device on which the volume TSO200 is mounted. Also, if the session option identifies a

specific device and the Postprocessor option identifies a different specific device, RMF issues a message to inform you that no data is currently available to meet your selection criteria.

Examples for a display session:

- To obtain a report on reserve activity for all devices, enter:
SENQR ALLVSR or SENQR
- To obtain a reserve activity report for the device on which the volume TSO200 is mounted, enter:
SENQR TSO200

Examples for a background session:

- To obtain reserve activity data for all devices, enter:
SENQR(ALLVSR) or SENQR
- To obtain reserve activity data for the device on which the volume TSO200 is mounted, enter:
SENQR(TSO200)

SPAG

►► SPAG ◄◄

Specifies reporting of paging activity.

SRCS

►► SRCS ◄◄

Specifies reporting of central storage/processor/SRM activity.

USER

►► USER ◄◄

Specifies that user-specified activity is to be reported. Your installation must supply a corresponding data-gathering module and data-reporting module before USER can take effect. See the *z/OS RMF Programmer's Guide* for more information.

Conflicting session options

After you enter the START session command from the operator console to start a background session, RMF processes the session options in a certain order (see Chapter 5, "How RMF processes session options," on page 59). 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. Messages notify the operator of the selections made. The possible conflicts are:

Conflict	Problem	RMF Resolution
NOREPORT and NORECORD specified	No way for installation to obtain measurement data	Change NOREPORT to REPORT (DEFER)

Conflict	Problem	RMF Resolution
STOP value specified is less than SINTV	Indicates session termination before obtaining any data	Set STOP value equal to SINTV value
REPORT(DEFER) and NOSTOP specified	SYSOUT becomes cluttered with unprinted reports	Change NOSTOP to STOP set equal to SINTV value

Chapter 15. Long-term reporting with the Postprocessor

This information unit deals with the Postprocessor, which generates reports from data gathered by the RMF monitors.

It describes:

- the Postprocessor capabilities
- how to prepare SMF records for processing
- how to start the Postprocessor
- the Postprocessor data sets
- all available options

You can use the Postprocessor to combine data from one, several, or all of the systems in the sysplex in one report. There are two prerequisites for this:

- Data must be gathered on all systems
- The gatherers on all systems must be synchronized

Recommendation: Specify option **SYNC(SMF)** for all gatherer sessions.

Postprocessor reporting

You can use the Postprocessor to create various reports based on SMF data. In addition, you can create records for further processing with spreadsheet applications on the workstation.

- Interval and duration reports:

Use the options **REPORTS** and **SYSRPTS** to get single-system and sysplex reports with a specified or a default measurement interval (interval reports). With the **DINTV** control statement (see “**DINTV**” on page 215) you can create duration reports combining data from several measurement intervals into one report interval. In “**REPORTS**” on page 223 and “**SYSRPTS**” on page 232, you find a description of all options that you can select to get the desired reports.

- Summary report:

With the option **SUMMARY**, you get a Summary report providing key performance data for a single system.

- Overview report:

The **OVW** option offers the capability of tailoring summary-like reports according to your requirements. You can create your own single-system and sysplex reports that show exactly the information you need for your performance management tasks. The tables in “Overview and exception conditions” on page 238 provide information about all data that is available for overview processing, related to the SMF data which is the source for these values. Because of compatibility reasons, you also can use option **EXCEPT** for single-system reports.

Furthermore, in *z/OS RMF Report Analysis* you find the relationship for each report field in a Postprocessor report to a name you can specify as data field in an Overview report.

- Overview records:

In the same way, as you can create Overview reports, you can create Overview records, just by specifying an additional option **OVERVIEW(RECORD)**. You can

download the records to the workstation for further processing in spreadsheets, or you can use the Spreadsheet Reporter to create and submit Postprocessor jobs without logging on to the host system. In addition, several spreadsheets are available to create a wide range of reports. See Chapter 18, “RMF Spreadsheet Reporter,” on page 293 for a detailed description.

- Exception report:

For exception reporting, you can use the same names as described for overview reporting. The difference is that you specify thresholds for these values as criterion whether the data will be reported or not.

Preparing SMF records for postprocessing

The Postprocessor generates reports based on data gathered in SMF records by the following data gatherer functions:

- RMF Monitor I, Monitor II, Monitor III — gathering SMF records type 70 - 79
- HTTP Server — gathering SMF records type 103
- Lotus Domino[®] Server — gathering SMF records type 108

Because SMF produces VSAM data sets or SMF log streams, but the Postprocessor cannot read these formats, you must copy the SMF records into sequential data sets. You should do this by using the IFASMFDL program for SMF data sets or the IFASMFDL program for SMF log streams. For more information about IFASMFDL or IFASMFDL see *z/OS MVS System Management Facilities (SMF)*. Using other utilities to copy SMF records often results in truncated or unusable records. The input data sets to the Postprocessor should not be compressed.

Sorting is required: The SMF records **must** be sorted by RMF interval start date and interval start time in the data set. If you want to combine SMF records from several data sets, you must sort the records together to ensure correct reports.

RMF provides two SORT exits (ERBPPE15 and ERBPPE35) that should be used when running the SORT program. Use the sample job supplied with RMF in SYS1.SAMPLIB(ERBSAMPP) for sorting Postprocessor input:

Example:

```
//ERBSAMPP JOB job information,REGION=0M
//RMFSORT EXEC PGM=JCL
//SORTIN DD DISP=SHR,DSN=<input_smfdata_system1>
// DD DISP=SHR,DSN=<input_smfdata_system2>
// :
// :
// DD DISP=SHR,DSN=<input_smfdata_systemN>
//SORTOUT DD DISP=disp,DSN=<sorted_smfdata>,...
//SORTWK01 DD DISP=(NEW,DELETE),UNIT=SYSDA,SPACE=(CYL,(10))
//SORTWK02 DD DISP=(NEW,DELETE),UNIT=SYSDA,SPACE=(CYL,(10))
//SORTWK03 DD DISP=(NEW,DELETE),UNIT=SYSDA,SPACE=(CYL,(10))
//SYSPRINT DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//SYSIN DD *
SORT FIELDS=(11,4,CH,A,7,4,CH,A),EQUALS
MODS E15=(ERBPPE15,36000,,N),E35=(ERBPPE35,3000,,N)
```

You can then start the Postprocessor as described in “Starting the Postprocessor” on page 204. If you want to start the Postprocessor with JCL, you can use the statements shown in the example under “Starting with JCL” on page 205 as a second step in the sort job.

Sysplex reporting across time zones

If you have a sysplex with several processors running in different time zones, this results in different time stamps in the SMF records. Therefore, the Postprocessor does not recognize that these records belong to the same interval when creating a sysplex report. You can solve this problem by calling program ERBCHGMT which updates the time stamps during Postprocessor processing.

A sample job supplied with RMF in SYS1.SAMPLIB(ERBGMTTP) for updating the Postprocessor input is available.

ERBGMTTP contains the following job steps:

1. SMFDUMP extracts the SMF records from the input data sets and makes sure that the output file has the correct DCB attributes.
If you want to process SMF records from SMF log streams, use the IFASMFDL program instead of IFASMFDL for SMF data sets.
If you want to process SMF records from the SMF data buffer, use the RMF-supplied program ERBAPPL.
2. CHGGMT changes the GMT offset in the input data sets to a common value. The parameter in the PARM field of program ERBCHGMT determines the GMT offset (in minutes) for the output data set.
Examples:
 - PARM='+0' creates GMT time in all time stamps in the Postprocessor.
 - PARM='-300' creates GMT time minus 5 hours (Eastern time US, Canada)
 - PARM='+60' creates GMT time plus 1 hour (Middle European winter time)
3. RMFSORT sorts the modified input data set(s) by RMF date and time. This step is required.
4. RMFPP invokes the Postprocessor.

Working with SMF data sets

```
//ERBGMTTP JOB job information,REGION=0M
//*
//* DUMP SMF DATASET(S)
//*
//SMFDUMP EXEC PGM=IFASMFDP
//IDD1 DD DISP=SHR,DSN=<input_smfdata_system1>
//IDD2 DD DISP=SHR,DSN=<input_smfdata_system2>
//
//
//IDDN DD DISP=SHR,DSN=<input_smfdata_systemN>
//SMFDATA DD DISP=(NEW,PASS),SPACE=(CYL,(10,10),RLSE),
// UNIT=SYSDA,DCB=(RECFM=VBS,LRECL=32760,BLKSIZE=0)
//SYSPRINT DD SYSOUT=*
//SYSIN DD *
        INDD(IDD1,OPTIONS(DUMP))
        INDD(IDD2,OPTIONS(DUMP))
        :
        INDD(IDDN,OPTIONS(DUMP))
        OUTDD(SMFDATA,TYPE(70:79,103,108))
//*
//* UPDATE GMT OFFSET
//*
//CHGGMT EXEC PGM=ERBCHGMT,PARM='-300'
//SMFDATA DD DISP=(OLD,DELETE),DSN=*.SMFDUMP.SMFDATA
//SMFCHGMT DD DISP=(NEW,PASS),SPACE=(CYL,(10,10),RLSE),
// UNIT=SYSDA,DCB=(RECFM=VBS,LRECL=32760,BLKSIZE=0)
//SYSPRINT DD SYSOUT=*
//*
//* SORT THE SMF RECORDS
//*
//RMFSORT EXEC PGM=SORT
//SORTIN DD DISP=(OLD,DELETE),DSN=*.CHGGMT.SMFCHGMT
.
.
//SYSIN DD *
        SORT FIELDS=(11,4,CH,A,7,4,CH,A),EQUALS
        MODS E15=(ERBPPE15,36000,,N),E35=(ERBPPE35,3000,,N)
//*
//* RMF POSTPROCESSING
//*
//RMFPP EXEC PGM=ERBRMFPP
//MFPINPUT DD DISP=(OLD,DELETE),DSN=*.RMFSORT.SORTOUT
:
```

Working with the SMF buffer (first step)

```
//SMFDUMP EXEC PGM=ERBAPPL,PARM='*/*/70:79,103,108'
//SMFDATA DD DISP=(NEW,PASS),UNIT=SYSDA,SPACE=(CYL,(10,10))
//ERBLIST DD DUMMY
```

Figure 36. JCL example for preparing SMF records

Starting the Postprocessor

Enablement: RMF is an optional feature 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. The following Postprocessor start procedures will not work, and you will receive the message:

ERB111I RMF IS NOT ENABLED TO RUN ON THIS SYSTEM

The Postprocessor executes as a background job. You can start the batch job from ISPF by choosing *Postprocessor* in the *RMF - Performance Management* menu.

The Postprocessor must have access to the system data set SYS1.SCEERUN, which contains run-time modules for the Language Environment. If you have specified this data set in the LINKLST of your system, you need take no further action in this respect (see “Ensure linkage to language environment” on page 16.) Just follow the instructions described in this information unit.

If SYS1.SCEERUN is not in the LINKLST, you must specify it as the STEPLIB of the job you submit to start the Postprocessor.

The SMF records that the Postprocessor uses as input can be in:

- SMF data sets

If you include an MFPINPUT DD statement in the start-up job, the records from the associated data set are used.

- an SMF buffer in each system of the sysplex

If you omit the MFPINPUT DD statement, the Postprocessor uses the Sysplex Data Services to access the SMF buffers.

Note: This access requires the appropriate access authorization, please refer to “Controlling access to RMF data for the sysplex data services” on page 20 for details.

Starting with JCL

To start processing from an SMF data set, use this sample JCL:

```
//EXAMPLE JOB job information,REGION=0M
//POST EXEC PGM=ERBRMFPP
//MFPINPUT DD <sorted SMF records>
//SYSIN DD *
control statements
/*
```

Figure 37. JCL example for starting the Postprocessor

Provide the SMF record data to be postprocessed in the data set specified on the MFPINPUT DD statement. Because RMF can generate spanned SMF records, this DD statement must contain DCB parameters if the input is an unlabeled tape. The DCB parameters are as follows:

```
RECFM=VBS,LRECL=32756,BLKSIZE=xxx
```

If the input is on a labelled tape or DASD, do not specify any DCB parameters.

The SMF records must be sorted. If this has not already been done, you can combine the sort step from the example from Figure 36 on page 204 with the Postprocessor start step shown in Figure 37 in a two-step job. Use the SORTOUT data set from the SORT step as the MFPINPUT data set for the ERBRMFPP step.

To start processing from the SMF buffers, use this sample JCL:

```
//EXAMPLE JOB job information,REGION=0M
//POST EXEC PGM=ERBRMFPP
//SYSIN DD *
control statements
/*
```

Example:

For information about control statements, see “How to use control statements” on page 213 and “Examples of control statements” on page 234.

When you specify control statements::

- You can specify the control statements in any order
- Specify data only in columns 1 to 72
- Do not continue statements over two or more lines. Repeat a control statement until all required options are specified

Starting from ISPF

When you select “Postprocessor” in the RMF Performance Management menu, you get the Postprocessor Setup panel. This ISPF Postprocessor interface consists of a series of panels presented in sequence. Here you can specify the parameters mentioned in “Starting with JCL” on page 205 as input on the panels.

Postprocessor Setup panel

```
RMF - Postprocessor Setup
Command ==>

Input Data      ==> DATASET  DATASET, SDS (Sysplex Data Server Buffers)
                  or SMFLOG (SMF Log Streams)
Output Data     ==> YES      YES or NO (NO to route output to SYSOUT)

Report Profile  ==> _____

Edit generated JCL ==> NO_    YES or NO

Job Statement Information:
==> //uidP      JOB (ACCOUNT),'PGMRNAME',CLASS=A,REGION=32M
==> /**
==> /**
==> /**

Complete this panel and press ENTER to continue, or END to exit.
To return to RMF Primary Menu without saving input, enter CANCEL.
```

Figure 38. Postprocessor - Setup Panel

The panel accepts this information:

- **Input Data.** Use this part of the panel to indicate the type of input data that should be used for the report. Specify:
 - DATASET - to display the Postprocessor Input panel for SMF data sets (see Figure 39 on page 208) where you may enter up to 14 SMF data set names
 - SMFLOG - to display the Postprocessor Input panel for SMF log streams where you may enter up to 14 SMF log stream names.
 - SDS - to have the report generated using Sysplex Data Services to access the SMF buffers
- **Output Data.** If you specify NO, then all output will be routed to SYSOUT. You can define your own output data sets by specifying YES, this leads you to the Postprocessor Output panel (see Figure 40 on page 208).

If you plan to let the Postprocessor create Overview records, you have to specify YES to define the appropriate data set.
- **Report Profile.** This is the name of the data set containing the Postprocessor control statements that define the report details, namely:

- EXCEPT
- EXRPTS
- OVW
- REPORTS
- SYSRPTS

It is used as SYSIN on the ERBRMFPP step of the Postprocessor call. Therefore, the DCB parameters have to be as follows:

```
RECFM=FB,LRECL=80,BLKSIZE=xxx
```

The options that are generated from the Postprocessor Options panel (see Figure 41 on page 209) are added to the control statements that you define in your profile data set. The values from the Option panel are handled with first priority if there is an overlap in both definitions. Note that the values in the profile data set must not contain JCL delimiter characters.

Example: If you have defined the Postprocessor control statements in your data set userid.RMF.SYSIN(DAILYREP), then you set

```
Report Profile ==> rmf.sysin(dailyrep)
```

The specification of all data sets is made following the standard TSO naming conventions.

- **Edit generated JCL.** Use this field to indicate whether or not you would like to edit the JCL for this report before submitting it.
If you enter YES, your JCL is displayed in edit mode after the Postprocessor Options panel. When in edit mode, you may either:
 - Make any necessary changes, then press END to submit the job and return to the Postprocessor Setup panel.
 - Enter CANCEL to cancel the job without saving changes and return to the Postprocessor Setup panel.
- **Job Statement Information.** Here you enter the information that is needed for the JOB statement of the batch job that is to start the Postprocessor. RMF generates the other job-control statements on the basis of your entries.

Postprocessor input: SMF data sets or log streams

If you specified DATASET or SMFLOG as input type in the Postprocessor setup, RMF displays the Postprocessor Input panel. Figure 39 on page 208 shows an example of an input panel for SMF data sets. If you specified SMFLOG as input type in the Postprocessor setup, the Postprocessor Input panel prompts you for the log stream names instead of SMF data sets.

RMF - Postprocessor Input

Command ===>

SMF Data Sets

===>

===>

===>

===>

===>

===>

===>

===>

===>

===>

===>

===>

===>

===>

===>

===>

===>

Sort Input Data

===> YES

YES or NO

Press ENTER to continue. To return to previous panel, press END.

To return to the Postprocessor Setup Menu, enter CANCEL.

Figure 39. Postprocessor Input for SMF Data Sets

For both types of SMF input, you can specify whether the input data is to be sorted or not. It is mandatory to pass sorted SMF records to the Postprocessor, therefore, you can bypass sorting only if you provide sorted records as input.

Postprocessor output data sets

RMF - Postprocessor Output

Command ===>

Overview record data set:

DD-Name

Data Set Name

===> PPOVWREC

===>

Report and message data sets:

===>

===>

===>

===>

===>

===>

===>

Complete this panel and press ENTER to continue.

To return to previous panel, press END.

To return to the Postprocessor Setup Menu, enter CANCEL.

Figure 40. Postprocessor - Output Data Sets

This panel allows you to specify output data sets.

The definition of the ddname PPOVWREC is mandatory if you request Overview records for further processing.

In addition, you can allocate existing data sets to the various report and message files that are created by the Postprocessor. For details about output data set allocation, please refer to “Defining output data sets” on page 210.

Example: Please assume that the Postprocessor is requested to create one Overview and one Summary report, and also Overview records, then you can route this output to your own existing data sets userid.OVERVIEW.RECORD, userid.OVERVIEW.LIST, and userid.SUMMARY.LIST by specifying:

```
====> PPOVWREC    ====> overview.record
====> PPORP001    ====> overview.list
====> PPSUM001    ====> summary.list
```

If you want in addition, that all interval reports should be combined into data set userid.REPORTS.LIST, than you can specify:

```
====> PPRPTS      ====> reports.list
```

Postprocessor Options panel

When you have completed the Postprocessor Setup panel with input, profile and job control data, you will see the Postprocessor Options panel that allows you to specify report options:

RMF - Postprocessor Options

Reporting	(DATE)	Start	====> <u> </u> . <u> </u>	End	====> <u> </u> . <u> </u>	yy.ddd
		or Start	====> <u>01</u> / <u>01</u> / 1958	End	====> <u>12</u> / <u>31</u> / 2057	mm/dd/yyyy
Duration	(DINTV)		====> <u> </u> :			
Exception	(ETOD)	Start	====> <u>00</u> : <u>00</u>	End	====> <u>24</u> : <u>00</u>	hh:mm
Interval	(RTOD)	Start	====> <u>00</u> : <u>00</u>	End	====> <u>24</u> : <u>00</u>	hh:mm
Summary	(STOD)	Start	====> <u>00</u> : <u>00</u>	End	====> <u>24</u> : <u>00</u>	hh:mm

Summary	(SUMMARY)	====> INT,TOT	NO, INT, TOT, or TOT,INT
Overview	(OVERVIEW)	====> <u> </u>	RECORD, REPORT (or both)

DELTA	====> NO_	YES or NO
EXITS	====> NO_	YES or NO
SYSOUT	====> A	Sysout Class

SESSION	====> <u> </u>	Session ID Monitor II
SYSID	====> <u> </u>	System identifier

To (edit and) submit Postprocessor job, press ENTER.
 To return to previous panel, press END.
 To return to the Postprocessor Setup Menu, enter CANCEL.

Command ==>

Figure 41. Postprocessor Options panel

Here you can enter:

- The start and end dates of the reporting period
- The time value for the duration report
- The time ranges for the reports which require them
- The scope of the Summary report
- The output type of the Overview report
- Values for miscellaneous options, with a prompt text to remind you of the valid format for each

From your entered data, RMF generates job control statements. If you have requested to edit the generated JCL, you will enter edit mode when pressing ENTER, otherwise the job will be submitted.

You can use the following commands on this panel:

RESET

Reset all parameters to the default values

SAVE Save all values you have entered (if you do not want to submit the job)

Defining output data sets

RMF dynamically allocates all Postprocessor message and report data sets to SYSOUT. You can route output data to permanent data sets rather than to SYSOUT by allocating the data sets in the JCL of the Postprocessor job. The report data sets for Monitor I, OMVS, and XCF interval reports and duration reports use the ddnames MFRnnnnn and MFEnnnnn. The ddnames of these and the other report data sets and the message data set that the Postprocessor uses are:

Table 22. Postprocessor ddnames - text output

ddname	Contents	Allocations	Notes
MFEnnnnn	Report output after a recoverable abnormal end	One ddname for one data set to be allocated for each of the intervals contained in the input data.	RMF uses this data set to re-allocate report data sets after a recoverable ABEND. nnnnn is a decimal number from 00001 to 99999; successively generated. For example, if a session has 15 intervals, ddnames are MFE00001 through MFE00015.
MFPMSGDS	Message output	One MFPMSGDS data set is allocated each time the Postprocessor is executed.	To change the SYSOUT class parameter for this data set, you must preallocate the data set. You cannot change it in the RMF options.
MFRnnnnn	Report output	One ddname for one data set to be allocated for each of the intervals contained in the input data.	nnnnn is a decimal number from 00001 to 99999; successively generated. For example, if a session has 15 intervals, ddnames are MFR00001 through MFR00015. For creating one output data set, use ddname PPRPTS.
MFXnnnnn	Exception interval report data	One ddname for one data set to be allocated for each of the intervals included in the exception report.	nnnnn is a decimal number from 00001 to 99999. The first interval is assigned the ddname MFX00001. The second MFX00002, and each subsequent interval is assigned a ddname in ascending numerical order. If no exception interval reports are produced in a given interval, the data set for that interval is empty. For creating one output data set, use ddname PPXRPTS.
PPEXTnnn	Exception report output	One ddname for one data set to be allocated for each system included in the input data.	nnn is a decimal number from 001 to 999. The first system encountered is assigned to ddname PPEXT001, and each subsequent system is assigned a ddname in ascending numerical order.
PPORPnnn	Overview report output	One ddname for one data set to be allocated for each of the systems included in the input data.	nnn is a decimal number from 001 to 999. The first system encountered is assigned to ddname PPORP001, and each subsequent system is assigned a ddname in ascending numerical order.
PPRPTS	Combined interval report	One ddname for one data set to be allocated to contain all reports of all intervals included in the input data.	There is no dynamic allocation of this ddname. You must define it explicitly if you want to get all interval reports into one data set or output class. You should not use the subparameter DEFER for this ddname. If you define this ddname, no MFRnnnnn files will be created.

Table 22. Postprocessor ddnames - text output (continued)

ddname	Contents	Allocations	Notes
PPSUMnnn	Summary report output	One ddname for one data set to be allocated for each of the systems included in the input data.	nnn is a decimal number from 001 to 999. The first system encountered is assigned the ddname PPSUM001, and each subsequent system is assigned a ddname in ascending numerical order.
PPXRPTS	Combined exception report	One ddname for one data set to contain all exception reports for all intervals contained in the input data.	There is no dynamic allocation of this ddname. You must define it explicitly if you want to get all exception reports into one data set or output class. You should not use the subparameter DEFER for this ddname. If you define this ddname, no MFXXnnnnn files will be created.
PPXSRPTS	Sysplex report output	One ddname for one data set to be allocated to contain all sysplex reports included in the input data.	All sysplex reports are written to this data set.
RMFPnnnn	Monitor II session interval report output	One ddname is generated and one data set is created for each report for each session-identifier included in the reporting.	nnnn is a decimal number from 0001 to 9999. When the Postprocessor is to generate reports for more than one system, a separate data set is allocated for each report for each system. When operands for a Monitor II session are not specified on the REPORTS 1 statement, the Postprocessor uses the operands in the SMF record, and a separate data set is allocated each time the operands change.

Note:

1. If you omit the DCB characteristics for the mentioned message and report data sets, the characteristics used are:

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

If you omit the data set control block (DCB) characteristics for a data set allocated to the PPXRPTS ddname of the Postprocessor job, the following characteristics are used:

RECFM=VBA, LRECL=1028

For BLKSIZE, the optimal block size as determined by the system is used.

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

2. Please consider that the number of dynamically allocated data sets is limited to 1635.

Table 23. Postprocessor ddname - Overview record output

ddname	Contents	Allocations	Notes
PPOVWREC	Overview record output	One ddname for one data set to be allocated.	Not created automatically. The records in this data set can be used for other applications, for example, for conversion to spreadsheet.

Note: Define this data set explicitly in the JCL for the Postprocessor. Use these data set characteristics:

DCB=(RECFM=VB,LRECL=32756,BLKSIZE=32760)

Besides text reports, the Postprocessor can generate certain reports in XML format. The XML format is generated by specifying the following Postprocessor ddnames for the output data.

Table 24. Postprocessor ddnames - XML output

ddname	Contents	Allocations	Notes
XPOVWRPT	Combined Overview report in XML format	One ddname for one data set to contain all overview reports for each system included in the input data.	There is no dynamic allocation of this ddname. You must define it explicitly if you want to get all overview reports in XML format into one data set or output class. If you define this ddname, no PPORPnnn files are created.
XPRPTS	Combined single system report in XML format	One ddname for one data set to contain all single system reports for each interval included in the input data.	There is no dynamic allocation of this ddname. You must define it explicitly if you want to get all reports in XML format into one data set or output class. If you define this ddname, no MFRnnnnn files are created. If you define this ddname and PPRPTS, no XML output in file XPRPTS is created.
XPXSRPTS	Combined sysplex-wide report in XML format	One ddname for one data set to contain all sysplex reports for each interval included in the input data.	There is no dynamic allocation of this ddname. You must define it explicitly if you want to get all reports in XML format into one data set or output class. If you define this ddname, no MFRnnnnn files are created. If you define this ddname and PPXSRPTS, no XML output in file XPXSRPTS is created.

Note:

1. If the XML output is to be routed to permanent data sets rather than to SYSOUT, define the data set with RECFM=VB and LRECL between 256 and 8192.
2. For a complete list of Postprocessor reports which you can obtain as XML output in addition to the *Overview* report, refer to information unit *Long-term overview reporting with the Postprocessor* in the *z/OS RMF Report Analysis*.
Any others than the listed reports requested on the REPORTS and/or SYSRPTS control statement are ignored.

How the Postprocessor processes control statements

The Postprocessor verifies the control statements and builds a list of options that control the session. If you omit a statement, RMF substitutes the default value, if there is one, or ignores the option. A statement containing a syntax error causes the Postprocessor to terminate, in most cases.

In some cases, an error in a control statement does not cause the Postprocessor to end. RMF notes the condition, issues a warning message to the output message data set, and continues building an option list for the session. When processing is complete, the Postprocessor issues a message to the output message data set indicating the options in effect.

The option list consists of the options you have entered on control statements and any options for which the default values were used. Each option listed is followed by the input source from which the Postprocessor obtained the option. The possible sources are:

- SYSIN -- the option was specified on a control statement for the Postprocessor.
- DEFAULT -- the option was taken from the control statement defaults.

When RMF detects an invalid value and substitutes a default value, a warning message is issued, and DEFAULT appears in the option list.

Defining the reporting period

You can control the length of the reporting period with:

1. Control statements (DATE, RTOD, ETOD, and STOD) that indicate a specific range of dates and specific ranges of times.
2. The SMF record data set. The control statement defaults for the reporting period include all dates and all times in the SMF record data set. If you omit control statements, the Postprocessor generates reports for all dates and times included in the SMF record data set.

Postprocessor completion

When the Postprocessor has generated all requested reports, it issues a return code and ends the session. Any messages are available in the preallocated MFPMSGDS data set.

Among the messages issued, there may be some with the prefixes CEE and EDC. These are Language Environment messages issued by routines that the Postprocessor uses. Please see the *z/OS RMF Programmer's Guide* for details.

The return codes from the Postprocessor are:

Code Meaning

- | | |
|----|---|
| 0 | Normal completion -- reports generated as requested |
| 4 | Normal completion -- no RMF input records found or no RMF input records found that meet the user requirements specified in the control statements |
| 8 | Error -- see accompanying RMF message |
| 12 | Terminating error -- see accompanying RMF message |

How to use control statements

This topic describes purpose and syntax of the control statements that you can use to create Postprocessor reports. Table 25 gives a summary of the available control statements. It also indicates which control statements are required and which you can omit to accept a default value.

Supply the control statements after the SYSIN DD statement in the job you submit to start the Postprocessor.

Table 25. Postprocessor Control Statement Summary

Processor Control Statement	Interval Report	Duration Report	Summary Report	Exception Report	Overview Report	Notes	See
DATE	★	★	★	★	★	1	"DATE" on page 214
DELTA	★					1	"DELTA" on page 215
DINTV		★				2,3	"DINTV" on page 215
ETOD				★	★	1	"ETOD" on page 218
EXCEPT				★	★	2	"EXCEPT" on page 218
EXITS	★	★	★	★	★	1	"EXITS" on page 220
EXRPTS				★		2	"EXRPTS" on page 220

Table 25. Postprocessor Control Statement Summary (continued)

Processor Control Statement	Interval Report	Duration Report	Summary Report	Exception Report	Overview Report	Notes	See
OVERVIEW					★	2	"OVERVIEW" on page 221
OVW					★	2	"OVW" on page 222
REPORTS	★	★				2	"REPORTS" on page 223
RTOD	★	★				1	"RTOD" on page 230
SESSION	★					1	"SESSION" on page 230
STOD			★			1	"STOD" on page 231
SUMMARY			★			1	"SUMMARY" on page 231
SYSID	★	★	★	★	★	1	"SYSID" on page 232
SYSOUT	★	★	★	★	★	1	"SYSOUT" on page 232
SYSRPTS	★	★				2	"SYSRPTS" on page 232
Note: 1. If the default value is acceptable, you need not specify the control statement explicitly. 2. You must specify the control statement explicitly. 3. You cannot request duration reports concurrently with interval reports; each type of report requires a separate Postprocessor session. However, you can request duration reports concurrently with exception generated interval reports and summary and exception reports concurrently with either duration or interval reports.							

The remainder of this topic describes the control statements in alphabetical order.

DATE

The DATE control statement specifies the start and end date of the reporting period for all reports.

The syntax of the statement is:

```
DATE(yyddd,yyddd)
or
DATE(mmddyyyy,mmddyyyy)
```

where:

- yy is the last two digits of the year.
- ddd is the day of the year

or

- mm is the month (01 to 12)
- dd is the day of the month
- yyyy is the year in full-century form, for example, 2013

RMF supports a sliding window which covers the time frame:

Current Year - 50 ↔ Current Year + 49

This sliding window will be used to define the correct value of the century, if not defined explicitly. Write the dates in the full format with leading zeroes. Do not mix the two formats in one control statement. The first date is the starting date, and the second date is the ending date. Specify the dates in chronological order.

The default value is in the four-digit year format and ensures that all dates found in the input data set are reported.

If the entire reporting period falls within one calendar day, you specify the same date twice.

Example: To request reports for data collected on June 7, 2013, specify
DATE(06072013,06072013)

DELTA

The DELTA control statement specifies whether certain fields in Monitor II background session reports are to reflect total values or, after the first report, changed delta values. The fields that are affected by delta mode are described for each report in *z/OS RMF Report Analysis*. The syntax of the statement is:

DELTA/NODELTA

DELTA indicates that the affected fields are to reflect changed delta values, that is, the reports are generated in delta mode. NODELTA indicates that the affected fields are to reflect total values. NODELTA is the default.

DINTV

The DINTV control statement specifies that the Postprocessor is to generate duration reports and indicates the length of the duration interval. The Postprocessor uses the value specified with this control statement also for processing **Overview** and **Exception** reports. This allows you to produce trend reports over long time periods.

The duration interval is the length of time each report can cover and should be a multiple of the measurement interval. The syntax of the statement is:

DINTV(hhmm)

where hh denotes the hours and mm denotes the minutes. The maximum is 9960 which is equivalent to 100 hours.

Note:

1. Before z/OS V1.8 RMF, exception/overview duration processing was not implemented. If a user specified Postprocessor reports, RMF used the options specified with the REPORTS control statement as a filter to select the SMF records for producing these reports.

Starting with z/OS V1.8, exception/overview duration processing is supported. If exception/overview duration reporting is used together with normal Postprocessor reports in one job step, the RMF Postprocessor selects all SMF records for producing both the exception/overview reports as well as the normal reports. This leads to uncommon behavior: For example, if a user specifies REPORTS(DEVICE(TAPE)), RMF may additionally produce unwanted DEVICE(DASD) reports.

Therefore it is not recommended to use REPORTS control statements if you request exception/overview duration reports in the same job.


2. For duration reports, it is recommended to specify not only the reporting interval but also the date, even if the SMF input data contains only records for those days you want to report on.

The duration interval can be the same length as, or shorter than, the reporting period. If it is shorter, there will be several duration intervals in a reporting period. Figure 42 on page 217 illustrates how the duration interval relates to the reporting period. Assume a reporting period that covers the twelve hours between 8:00 AM and 8:00 PM over a range of three days. As the figure shows, specifying DINTV(1200) causes the Postprocessor to generate three duration reports, each covering twelve hours of system activity. Specifying DINTV(0600) causes the Postprocessor to generate six duration reports, each covering six hours of system activity. You could also choose a duration interval that does not break the reporting period into equal blocks of time. As the figure shows, specifying DINTV(1000), indicating a duration interval of ten hours, causes the Postprocessor to generate the following reports:


1. 8:00 AM to 6:00 PM on the first day of the reporting period.
2. 6:00 PM to 8:00 PM on the first day of the reporting period, and 8:00 AM to 4:00 PM on the second day.
3. 4:00 PM to 8:00 PM on the second day of the reporting period, and 8:00 AM to 2:00 PM on the third day.
4. 2:00 PM to 8:00 PM on the third day of the reporting period.

In this case, if you wanted to use the reports to compare system performance over the same hours on each day of the reporting period, it would be difficult because each report covers a different time range and some span two days and include a twelve-hour gap when no reporting was done.


Reporting Period



Duration Interval of Twelve Hours - DINTV (1200)



Duration Interval of Six Hours - DINTV (0600)



Duration Interval of Ten Hours - DINTV (1000)




Figure 42. Relationship between Duration Interval and Reporting Period

The syntax of the control statement allows a maximum duration interval of 99 hours and 60 minutes.

For most effective reporting, the duration interval should relate to the length of the reporting period.

The actual length of time included in the resulting duration report depends on the time within the duration interval when data was actually collected.

Example: Request a duration report for channel path activity that has a duration interval of 12 hours (from 8:00 AM to 8:00 PM) on June 7, 2013:

```
DATE(06072013,06072013)
RTOD(0800,2000)
DINTV(1200)
REPORTS(CHAN)
```

However, if channel path activity measurement did not begin until 10:00 AM, and ended at 6:00 PM, the duration report covers the time from 10:00 AM to 6:00 PM. The start time of the first measurement interval and the time when the last interval ended appear in the heading of the report.

Midpoint Processing: Postprocessor duration reporting can put data into the incorrect duration interval when the time stamp in the input record differs from the expected time. For example, if your installation uses 15-minutes intervals and specifies RTOD(0900,1000) and DINTV(0100), the time stamp for the 10:00 to 10:15 interval might indicate that the interval began at 09:59:59.997. In this case, the 10:00 record would be reported in the 9:00 to 10:00 duration interval and is one hour and 15 minutes long.

To avoid this problem, the Postprocessor checks whether the midpoint of each interval falls between the duration interval to ensure that a record will not be used unless at least half of the data belongs to the interval.

Therefore, you should specify RTOD(0900,1000) and **not** RTOD(0859,1000) to get the correct duration report.

This note also applies to Postprocessor interval, exception, and summary reporting.

ETOD

The ETOD control statement specifies the starting time and ending time of the reporting period for an Exception or Overview report for each day in the reporting period.

The syntax of the statement is:

```
ETOD(hhmm, hhmm)
```

where hh denotes the hours and mm denotes the minutes.

When the ETOD statement is omitted, the default value is ETOD(0000,2400); that is, all times are reported. Thus, you would use this control statement when you want a reporting period for an Exception report that is different from the default value.

Because the range of values allowed is from 0000 to 2400, it is not possible to define a reporting period that consists of a single block of time spanning more than one calendar day. Thus, you cannot define a reporting period that, for example, runs from 12 noon on one day to 12 noon on the next day. However, you can define a reporting period that consists of the same block of time over several days. For example, to produce an Exception report using data collected from 8:00 AM to 1:00 PM for the week beginning on January 3, 2013 and ending on January 9, 2013, the required DATE and ETOD statements are:

Example:

```
DATE(01032013,01092013)
ETOD(0800,1300)
```

EXCEPT

The EXCEPT control statement defines a condition that RMF is to test for an exceptional value. The definition consists of a condition (which is a system indicator that RMF recognizes by name), a threshold value for the condition, and an operator that establishes the relation between the condition and the threshold value. RMF compares the threshold value for the condition with the contents of the appropriate SMF record field. If the condition exceeds the threshold, RMF recognizes that an exception has occurred. The conditions that RMF can recognize and test for exceptional values are listed in "Overview and exception conditions" on page 238.

An exception can consist of one or multiple conditions. If you group conditions into a single exception, all conditions must exceed their threshold values to let RMF recognize that the exception has occurred.

If you specify an OVERVIEW control statement (see “OVERVIEW” on page 221) together with EXCEPT control statements, RMF produces an Overview report based on the EXCEPT control statements, but no Exception report.

By supplying an EXRPTS control statement for the exception, you can also cause RMF to generate one or more interval reports when the exception occurs.

Only one condition and threshold can be specified in a single control statement, however, you can specify multiple EXCEPT control statements. The syntax of the EXCEPT statement is:

```
EXCEPT([control-statement-name](condition-name[(qualifier)] ,{LE} ,threshold-value))
```

control-statement-name

Specifies a one to eight-character name, starting with an alphabetic character, that has three uses:

- First, it provides a means of grouping multiple conditions to form a single exception. You group conditions by coding the same control-statement-name on each separate EXCEPT statement that defines each one of the conditions that form the exception. When conditions are grouped, all conditions must be satisfied in order for RMF to recognize the exception.
- Second, the control-statement-name associates the exception with the interval reports, if any, that RMF is to generate when the exception occurs. The EXRPTS statement defines these reports. For more information, see the EXRPTS control statement.

When Exception reports are generated, exceptions are listed in chronological order by control-statement-name. The control-statement-names are listed in alphabetical order. Choosing a meaningful control-statement-name makes it easier to recognize an exception.

- Third, the control-statement-name is used in an Overview report as the header for the column with the corresponding exception data. In this case, the control-statement-name is mandatory, and must be unique for each exception.

condition-name

Specifies the name of the condition that RMF is to test for an exceptional value. All valid condition names are listed in the tables in “Overview and exception conditions” on page 238, which also show the record types that contain the data that RMF compares with the threshold for each condition. RMF performs the test by comparing the contents of the appropriate field in an SMF record with the threshold value. Thus, RMF can recognize an exception only when the required SMF record was produced during the reporting period.

qualifier

Specifies an optional qualifier that can limit the scope of the condition identified by condition-name. The tables in “Overview and exception conditions” on page 238 also list the available qualifiers.

LE or GE

Specifies the operator RMF is to use to determine if the exception has occurred.

LE indicates that any value in the SMF record that is less than or equal to the threshold value causes an exception.

GE indicates that any value in the SMF record that is greater than or equal to the threshold value causes an exception.

threshold-value

Specifies the value that RMF compares to a computed value from the appropriate SMF record fields. You can specify the value as:

- a whole number – where the value can be a one to six digit integer.
- a fraction – up to six digits can appear before and up to three digits can appear after the decimal point. However, the total number of digits specified cannot exceed nine, including the decimal point.
- percentage – where the maximum percentage that can be specified is 100 percent. The value can be expressed as a whole integer or as a fraction. You can specify a percentage only with those conditions that indicate a percent value.

When the SMF record exceeds the threshold value, as indicated by the LE or GE operand, RMF flags the condition for reporting. For each interval, only one line is printed for each exception regardless of the number of times the threshold is exceeded.

Note: For information on exception/overview duration reporting, refer to “DINTV” on page 215.

Examples:

- The Postprocessor is to generate a line in the Exception report for each interval when the total busy percent value for channel path 1 is greater than or equal to ten percent. Use the following control statement:

```
EXCEPT((CHTBSY(01),GE,10))
```

The absence of a control-statement-name indicates that this is a single condition exception and that no interval reports are generated.

- The Postprocessor is to generate a Channel Path Activity report for each interval when the I/O service rate for service class TSO is higher than or equal to 100 service units per second, and the device percent utilized value for device 06D8 is greater than or equal to 3. Use the following control statements:

```
EXCEPT(EXCP1(IOSRV(S.TSO),GE,100))  
EXCEPT(EXCP1(DVUTL(06D8),GE,3))
```

Note that the control-statement-name EXCP1 is used to group the two conditions into one exception. When RMF recognizes both conditions, RMF writes a line for each condition in the Exception report.

EXITS

The EXITS control statement specifies whether or not a user exit routine is to be executed during the Postprocessor session. The syntax of the statement is:

EXITS/NOEXITS

When EXITS is specified, a user exit routine is executed. See the *z/OS RMF Programmer's Guide* for information on how to code a user exit routine for the Postprocessor.

The default NOEXITS indicates that no user exit routine is to be executed.

EXRPTS

The EXRPTS control statement is required when you want the Postprocessor to generate interval reports when a particular exception occurs. Each report specified

must be separated from any other reports by a comma. No continuation statements are permitted. However, you can use multiple EXRPTS statements. The syntax of the statement is:

```
EXRPTS (control-statement-name(report,report,...))
```

control-statement-name

Specifies a one to eight-character name starting with an alphabetic character that associates one or more EXCEPT statements with the EXRPTS statement. The EXCEPT statement defines the exception to RMF and the EXRPTS statement defines the action to be performed by RMF.

report

Specifies any of the reports based on measurements from Monitor I or Monitor III that is acceptable on the REPORTS control statement. The following key-words can be specified for report. For their meaning, see "REPORTS" on page 223:

```
ALL
CACHE(option,option,...)|NOCACHE
CHAN|NOCHAN
CPU|NOCPU
CRYPTO|NOCRYPTO
DEVICE(option,option,...)|NODEVICE
ENQ|NOENQ
ESS|NOESS
FCD|NOFCD
HFS|NOHFS
IOQ|NOIOQ
OMVS|NOOMVS
PAGESP|NOPAGESP
PAGING|NOPAGING
TRACE|NOTRACE
VSTOR(option,list)|NOVSTOR
XCF|NOXCF
```

Note: The EXRPTS control statement can not be used for reports that are only available in XML output format. Hence, for example, no Postprocessor *Serialization Delay* report can be generated, because the keyword SDELAY is not supported in the EXRPTS control statement.

Example: The Postprocessor is to generate a Channel Path Activity report for each interval when the I/O service rate for service class TSO is higher than or equal to 100 service units per second. Use the following control statements:

```
EXCEPT(CHNRPT(IOSRV(S.TSO),GE,100))
EXRPTS(CHNRPT(CHAN))
```

OVERVIEW

The OVERVIEW control statement can be used together with the OVW and the EXCEPT control statement to specify Overview processing.

The syntax of the statement is:

```
OVERVIEW(type[,type])
```

type specifies the output destination and can be:

REPORT

Requesting a report to be written.

RECORD

Requesting data to be written to a data set.

You can specify both types on one OVERVIEW control statement.

OVERVIEW in context with OVW

If you specify OVW control statements (see “OVW”), then OVERVIEW has the default value of REPORT and is required only for creating records, either as OVERVIEW(RECORD) or OVERVIEW(REPORT,RECORD).

OVERVIEW in context with EXCEPT

If you specify EXCEPT control statements (see “EXCEPT” on page 218), then OVERVIEW defines whether Overview reports or Overview records should be created. There is no default OVERVIEW option, and if OVERVIEW is missing, an Exception report will be created.

OVERVIEW in context with OVW and EXCEPT

It is recommended **not** to specify OVW and EXCEPT statements together. If you do so nevertheless, overview data is created using both types of statements, and you must ensure not to use duplicate control statement names.

There is no automatic allocation of the output data set for the Overview records. You have to define the data set explicitly in your JCL (see “Defining output data sets” on page 210) as

```
//PPOVWREC DD DSNAME=data.set.name, ...
```

or in the Postprocessor Output Data Set panel (see “Postprocessor output data sets” on page 208).

The data set should have these characteristics:

```
DCB=(RECFM=VB,LRECL=32756,BLKSIZE=32760)
```

You find a description of the records in *z/OS RMF Report Analysis*.

OVW

You can use the OVW control statement to create Overview reports and Overview records. Depending on the condition names, the SMF data, and the optional parameters SYSTEMS|NOSYSTEMS, you get single-system and sysplex reports or records.

Note: There is a limitation of 253 OVW control statements for one step of the Postprocessor.

The OVW syntax is:

```
OVW(control-statement-name(condition-name(qualifier))[,SYSTEMS])  
or  
OVW(control-statement-name(condition-name(qualifier))[,NOSYSTEMS])
```

If in the past you have performed overview processing with EXCEPT control statements, you should exchange them to OVW control statements because OVW statements enable you to create sysplex reports. The OVW syntax just requires a control statement name, but no threshold, as required for the EXCEPT control statement.

Refer to “EXCEPT” on page 218 for an explanation of the parameters control-statement-name and condition-name(qualifier).

Note: For information on exception/overview duration reporting, refer to “DINTV” on page 215.

SYSTEMS

You get a single-system report for each system of the sysplex, and you get an additional sysplex report for all sysplex-related overview conditions (based on the records for the CF, SDEVICE, and WLMGL reports). This optional parameter (which is the default) is valid only in combination with sysplex-related overview conditions, but is tolerated for single-system overview conditions.

NOSYSTEMS

Only sysplex-wide overview data is generated.

You find a list of all condition names in “Overview and exception conditions” on page 238.

Example: You want to get the sysplex-wide response times for all TSO users running in service class TSOSERV, and you have defined three service class periods. You want to get the data for all intervals between 10:00 AM and 2:00 PM, and you want to create a report as well as records.

You specify the following control statements:

```
ETOD(1000,1400)
OVERVIEW(REPORT,RECORD)
OVW(RTIMEP1(RTIMETOT(S.TSOSERV.1)),NOSYSTEMS)
OVW(RTIMEP2(RTIMETOT(S.TSOSERV.2)),NOSYSTEMS)
OVW(RTIMEP3(RTIMETOT(S.TSOSERV.3)),NOSYSTEMS)
```

Note that you specify your own appropriate control-statement-names (in our example RTIMEP1, ... , RTIMEP3) that appear as column headings for the desired measurements in the resulting report or record.

REPORTS

The REPORTS control statement specifies the reports to be generated by the Postprocessor for a single system. In combination with the control statement DINTV, duration reports are generated. Otherwise, interval reports are generated.

Note: No duration reports are available for enqueue activity (ENQ), tracing activity (TRACE), and serialization delay (SDELAY), as well as for all reports based on Monitor II data. This applies also when specifying ALL together with DINTV.

Each report specified as an option of the REPORTS control statement must be separated from any other reports by a comma. No continuation statements are permitted. However, you can specify multiple REPORTS control statements. The syntax of the REPORTS control statement is:

```
REPORTS(report[,report...])
```

where the **report** option can be any of the reports listed here and described in detail in the following:

ALL	FCD() NOFCD
ARD() NOARD	HFS NOHFS
ARDJ() NOARDJ	HTTP NOHTTP
ASD() NOASD	IOQ NOIOQ

ASDJ() NOASDJ	IOQUEUE NOIOQUEUE
ASRM() NOASRM	OMVS NOOMVS
ASRMJ() NOASRMJ	PAGESP NOPAGESP
CACHE() NOCACHE	PAGING NOPAGING
CHAN NOCHAN	PCIE NOPCIE
CHANNEL NOCHANNEL	PGSP() NOPGSP
CPU NOCPU	SDELAY NOSDELAY
CRYPTO NOCRYPTO	SENQ() NOSENQ
DEV() NODEV	SENQR() NOSENQR
DEVICE() NODEVICE	SPAG NOSPAG
DEVV() NODEVV	SRCS NOSRCS
DOMINO NODOMINO	TRACE NOTRACE
ENQ NOENQ	VSTOR() NOVSTOR
ESS() NOESS	XCF NOXCF

Examples:

```
REPORTS(CPU)
REPORTS(CPU,CRYPTO)
REPORTS(CPU,CRYPTO,DEVICE(COMM,DASD))
```

ALL

Indicates that all of the listed reports are to be generated, if gathered data is available. Any user-supplied Monitor II background session reports are also included when ALL is specified. ALL can be combined with explicit specifications of other options.

Examples:

```
REPORTS(ALL,NOENQ,DEVICE(NOUNITR,NOCOMM))
```

All of the reports are generated, with the exception of enqueue activity and device activity for unit record devices and communication equipment.

```
REPORTS(NOCPU)
or
REPORTS(ALL,NOCPU)
```

The Postprocessor generates all reports except for the CPU Activity report.

```
REPORTS(ALL,CPU)
or
REPORTS(ALL)
```

The Postprocessor generates all reports.

ARD[(class,status)] | NOARD

Specifies the Monitor II Address Space Resource Data report, where suboptions **class** and **status** specify the selection conditions for the address spaces to be included. For detailed information refer to Table 21 on page 187.

ARDJ(jobname) | NOARDJ

Specifies the Monitor II Address Space Resource Data by jobname report, where suboption **jobname** identifies a specific job for the report. If you specify ARDJ for the Postprocessor, make sure you have specified the Monitor II ARDJ report command. You cannot run a Postprocessor ARDJ report with data collected by the Monitor II ARD report command. For detailed information refer to Table 21 on page 187.

ASD[(class,status)] | NOASD

Specifies the Monitor II Address Space State Data report. Suboptions **class**, and **status** are the same as for ARD. For detailed information refer to Table 21 on page 187.

ASDJ(jobname) | NOASDJ

Specifies the Monitor II Address Space State Data by jobname report, where **jobname** identifies a specific job for the report. If you specify ASDJ for the Postprocessor, make sure you have specified the Monitor II ASDJ option. You cannot run a Postprocessor ASDJ report with data collected by the Monitor II ASD option. For detailed information refer to Table 21 on page 187.

ASRM[(class,status)] | NOASRM

Specifies the Monitor II Address Space SRM Data report. Suboptions **class** and **status** are the same as for ARD. For detailed information refer to Table 21 on page 187.

ASRMJ(jobname) | NOASRMJ

Specifies the Monitor II Address Space SRM Data by jobname report, where **jobname** identifies a specific job for the report. If you specify ASRMJ for the Postprocessor, make sure you have specified the Monitor II ASRMJ option. You cannot run a Postprocessor ASRMJ report with data collected by the Monitor II ASRM option. For detailed information refer to Table 21 on page 187.

CACHE([SSID(list)][,EXSSID(list)][,DEVICE|SUBSYS][,SUMMARY]) | NOCACHE

Specifies the Monitor I Cache Subsystem Activity report.

SSID(list)

Specifies an unlimited list of storage subsystem identifiers (SSIDs), identifying the control units to be included in the report. Each element in the list can be:

- A single SSID
- A range of SSIDs, defined by the lowest and the highest SSID, separated by a colon.

EXSSID(list)

Causes the Postprocessor to suppress reports for the control unit or control units with the SSIDs specified. You can specify an unlimited number of elements in the list using the same syntax as for the SSID option.

If EXSSID is not specified, RMF reports on all control units in the SSID option list, or, if the SSID option has not been specified, all control units are reported on.

DEVICE|SUBSYS

Specify **DEVICE** to create a report on device level and additionally a report on subsystem level for each reported control unit.

Specify **SUBSYS** to create reports on subsystem level only.

SUMMARY

Specify **SUMMARY** to create a Summary report. You can specify this option in addition to the other options. If **SUMMARY** is the only parameter, you just get the Summary report.

CHAN | NOCHAN

Specifies the Monitor I Channel Path Activity report.

CHANNEL | NOCHANNEL

Specifies the Monitor II Channel Path Activity report. For detailed information refer to Table 21 on page 187.

CPU | NOCPU

Specifies the Monitor I CPU Activity report.

CRYPTO | NOCRYPTO

Specifies the Monitor I Crypto Hardware Activity report.

DEV [(type)] | NODEV

Specifies the Monitor II table report for I/O device activity. You can request device activity by specifying all devices in one class, or one or more specific device numbers, volume serial numbers, or storage groups. For detailed information refer to Table 21 on page 187.

DEVICE(suboption1[,suboption2,...]) | NODEVICE

Specifies the Monitor I Device Activity reports. You can request device activity by specifying all devices within one or more classes, and, optionally, one or more specific devices.

You can specify any of the suboptions listed below.

Note: The default values for the Postprocessor are listed below. They are different from the Monitor I session default values.

- A device number in the form NMBR(nmbr1,nmbr2) where nmbr1 and nmbr2 are 4-digit hexadecimal numbers.
- Any of the following classes:

CHRRDR|NOCHRRDR

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

- Storage groups in the form SG (aaa,bbb) where aaa and bbb are 1 to 8 character names. The report will be sorted by device number within storage group.

When you omit DEVICE and specify ALL, the device classes defaults underscored above are included in the report. When you specify DEVICE, you must include a list of either device classes, numbers, or both.

When you specify a device class in the **option** field, the reports generated depend on whether you have used the negative value or the positive value of the option. If you use a negative option, you get the device reports with the exception of the option or options you specify. For example, DEVICE(NOTAPE) causes the Postprocessor to generate all Device Activity reports except the report on magnetic tape devices. If you use a positive option, you will get only the device report corresponding to that option. For example, DEVICE(TAPE) causes the Postprocessor to generate the Device Activity report for magnetic tape devices; no other Device Activity reports are printed.

The NMBR field indicates that RMF is to report on the specific devices identified. The numbers can be expressed as a single device or as a range of devices. A range is indicated by specifying the first and last device numbers separated by a colon. Each single number or range is separated by a comma.

For example, to request device reporting for magnetic tape devices 2180, 2183, 2184, 2185, and 2188 as well as all direct access devices and communication equipment, you would specify:

```
REPORTS (DEVICE (COMM, DASD, NMBR (2180, 2183:2185, 2188)))
```

RMF reports on the storage groups you specify in the SG field. You can select one storage group name or a range of storage groups. To select a range of storage groups, specify NODASD, and the first name and the last name with a colon between them. For example to select the range of storage groups from PROC01 to PROC05, specify:

```
REPORTS (DEVICE (NODASD, SG (PROC01:PROC05)))
```

To select one storage group, for example, PROC02, specify:

```
REPORTS (DEVICE (NODASD, SG (PROC02)))
```

DEVV(id) | NODEVV

Specifies the Monitor II row report for device activity, where **id** is either a specific VOLSER or device number. For detailed information refer to Table 21 on page 187.

DOMINO | NODOMINO

Specifies the Monitor I Lotus Domino Server report.

ENQ | NOENQ

Specifies the Monitor I Enqueue Activity report. The level of enqueue activity reporting depends on the level selected at the time the data was gathered. There is no ENQ duration report.

ESS | NOESS

Specifies the Monitor I ESS Enterprise Disk Systems report (link statistics, extent pool and rank statistics).

FCD([option,option,...]) | NOFCD

Specifies the Monitor I FICON Director Activity report, where *option* can be one of the following:

NMBR(list)

Specifies a list of FICON directors to be included into the report identified by their hexadecimal switch device numbers. If NMBR is not specified, all FICON directors are selected except those specified by the EXNMBR option.

EXNMBR(list)

Specifies a list of FICON directors to be excluded from the report identified by their hexadecimal switch device numbers. If EXNMBR is not specified, all FICON directors in the NMBR option list are selected, or, if NMBR has not been specified, all FICON directors are selected.

HFS | NOHFS

Specifies the Monitor I HFS Statistics report.

HTTP | NOHTTP

Specifies the Monitor I HTTP Server report.

IOQ | NOIOQ

Specifies the Monitor I I/O Queuing Activity report.

IOQUEUE(type) | NOIOQUEUE

Specifies the Monitor II I/O Queuing Activity report. For detailed information refer to Table 21 on page 187.

OMVS | NOOMVS

Specifies the Monitor I OMVS Kernel Activity report.

PAGESP | NOPAGESP

Specifies the Monitor I Page Data Set Activity report.

PAGING | NOPAGING

Specifies the Monitor I Paging Activity report.

PCIE | NOPCIE

Specifies the Monitor I PCIE Activity Report.

PGSP | PGSP(option) | NOPGSP

Specifies the Monitor II Page Data Set Activity report. The following option is possible:

PAGE Get Page data set activity

Specifying PGSP is equivalent with PGSP(PAGE). For detailed information refer to Table 21 on page 187.

SDELAY | NOSDELAY

Specifies the Monitor I Serialization Delay Report.

SENQ | SENQ(suboption) | NOSENQ

Specifies the Monitor II System Enqueue Contention report. The operands describe the type of data you require:

S Summary report

D Detail report

A,sysname
Report with all resources

E,sysname
Report with exclusively-owned resources

majorname[,minorname]
Report for a specific resource

Specifying SENQ is equivalent with SENQ(S). You can specify only one operand. For detailed information refer to Table 21 on page 187.

Note:

1. If the session option specified a resource or group of resources by name, the report includes data for only those resources.
2. RMF treats the single character A, D, E or S as a request for the report. Therefore, A or E cannot be used as a major name; S or D cannot be used as a major name unless a minor name is also specified.
3. If the session option identified a different resource or group of resources, RMF issues a message to tell you that no data was available to meet your selection conditions.

SENQR | SENQR(option) | NOSENQR

Specifies the Monitor II System Enqueue Reserve report, where *option* describes the type of data you require:

ALLVSER
Report on all volumes

volser Report on a specific volume

Specifying SENQR is equivalent with SENQR(ALLVSR). For detailed information refer to Table 21 on page 187.

Note: If the session option identified a different specific device, RMF issues a message to tell you that no data was available to meet your selection conditions.

SPAG | NOSPAG

Specifies the Monitor II Paging Activity report. For detailed information refer to Table 21 on page 187.

SRCS | NOSRCS

Specifies the Monitor II Central Storage/Processor/SRM report. For detailed information refer to Table 21 on page 187.

TRACE | NOTRACE

Specifies the Monitor I Trace Activity report. There is no TRACE duration report.

user-report[(operands)]

Specifies a user-supplied Monitor II session report, where **user-report** is the name of the option used to collect data for the report and **operands** are any operands your installation established when the report was designed. When your report has operands but you do not specify any operands on the REPORTS control statement, the Postprocessor uses the menu default, if present, for any omitted operand. When there is no menu default, the Postprocessor takes the operand in effect when the data was collected.

See the *z/OS RMF Programmer's Guide* for a description of how to add a user-supplied report to the Postprocessor. After you have performed the steps that make your report available to the Postprocessor, your report will be printed when you specify the name of the report or ALL on the REPORTS control statement.

VSTOR | VSTOR(operands) | NOVSTOR

Specifies the Monitor I Virtual Storage Activity report. The operands describe the type of data you require:

S Summary report

D [,jobname1,jobname2,...]
Summary and detail report (for specified jobs)

jobname1 [,jobname2,jobname3,...]
Summary report for specified job(s)

RMF can produce common storage summary and detail reports and private area summary and detail reports. When you specify S, either explicitly or by default, RMF produces summary reports; when you specify D, RMF produces both summary reports and detail reports.

The Monitor I session gathers private area data only when you specify a jobname on the VSTOR option during the session. The Postprocessor, however, reports any private area data that it finds in the input records. Thus, it is not necessary to identify specific jobnames for the Postprocessor. (If you identify a specific jobname, the Postprocessor produces a private area report for that job only, and only if private area data for it exists in the input records.) It is a good practice to omit specific jobnames on the Postprocessor control statements. This practice enables you to use the same Postprocessor control

statement to obtain common storage report(s) or to obtain both common storage report(s) and private area report(s) when data exists for private area report(s).

If you specify VSTOR without any operands, RMF produces a summary report for common storage. Examples of other possible combinations are:

- REPORTS(VSTOR(D)) produces a summary and detail report for common storage. The Postprocessor also produces a summary and detail report for any private area data in the input records.
- REPORTS(VSTOR(D,VTAM)) produces a summary and detail report for common storage and a summary and detail report for the private area of the VTAM address space. The Postprocessor does not produce reports for any other private area data in the input records.
- REPORTS(VSTOR(MYJOB)) produces a summary report for common storage and a summary report for the private area of the MYJOB address space. The Postprocessor does not produce reports for any other private area data in the input records.

XCF | NOXCF

Specifies whether the XCF Activity report is to be generated. RMF produces an XCF usage by system, XCF usage by member, and XCF path statistics sections.

RTOD

The RTOD control statement specifies the starting time and ending time of the reporting period for interval or duration reporting for each day included in the reporting period. The syntax of the statement is:

RTOD(hhmm,hhmm)

where **hh** is the hour and **mm** is the minute on a 24-hour clock. Times must be specified in full, including leading zeroes.

The first time specifies the beginning of the reporting period and the second time specifies the end of the reporting period. The second time must be later than the first, or a syntax error occurs. When the RTOD statement is omitted, the default value is RTOD(0000,2400); that is, all times are reported.

Note: Because the range of values allowed is from 0000 to 2400, it is not possible to define a reporting period that consists of a single block of time that spans more than one calendar day. For example, you cannot define a reporting period that runs from 12 noon on one day to 12 noon on the next day. However, you can define a reporting period that consists of the same block of time over several days.

Example: For example, to produce interval reports including data collected every morning from 8:00 AM to 1:00 PM for the week beginning on January 3, 2013, and ending on January 9, 2013, the required DATE and RTOD statements would be:

```
DATE(01032013,01092013)
RTOD(0800,1300)
```

SESSION

The SESSION control statement specifies the particular Monitor II background session that created the SMF records to be included in the reports. The syntax of the statement is:

SESSION(session-id)

where **session-id** is the two-character alphanumeric session identifier of the particular session. If you explicitly specify **SESSION**, you must supply a session identifier. Only one session-id may be reported on during a Postprocessor session.

When you omit the **SESSION** statement, all SMF records that fall within the reporting period and are pertinent to the types of reports specified on the **REPORTS** statement are included in the reports, regardless of the session that created them.

STOD

The **STOD** control statement specifies the starting time and ending time of the reporting period for a Summary report for each day in the reporting period.

The syntax of the statement is:

```
STOD(hhmm,hhmm)
```

where **hh** is the hour and **mm** is the minute on a 24-hour clock. The first time specifies the beginning of the reporting period and the second time specifies the end of the reporting period. The second time must be later than the first, or a syntax error occurs.

When the **STOD** statement is omitted, the default value is **STOD(0000,2400)**; that is, all times are reported. Thus, you would use this control statement when you want a reporting period for a Summary report that is different from the default value.

Because the range of values allowed is from 0000 to 2400, it is not possible to define a reporting period that consists of a single block of time that spans more than one calendar day. Thus, you cannot define a reporting period that, for example, runs from 12 noon on one day to 12 noon on the next day. However, you can define a reporting period that consists of the same block of time over several days. For example, to produce a Summary report using data collected from 8:00 AM to 1:00 PM for the week beginning on January 3, 2013, and ending on January 9, 2013, the required **DATE** and **STOD** statements would be:

Example:

```
DATE(01032013,01092013)  
STOD(0800,1300)
```

SUMMARY

The **SUMMARY** control statement specifies whether a Summary report is to be produced and indicates the type of summary data that you require. The syntax of the statement is:

```
SUMMARY(type)|NOSUMMARY
```

where **type** can be either or both of the following:

INT Indicating that one interval summary line is to be produced for each measurement interval that falls within the reporting period.

TOT Indicating that one total summary data line is to be produced for all the measurement intervals that fall within the reporting period.

When both are specified, **INT** and **TOT** can appear in any order. When you explicitly specify **SUMMARY**, you must specify the type of summary data that you require. Specifying **SUMMARY** without **type** causes a syntax error. When you omit

the SUMMARY statement, the default is SUMMARY(INT,TOT). That is, a Summary report is produced, and the report includes both interval summary data lines and a total summary data line.

When a Summary report consists of more than one page, the headings are repeated for each page. When total summary data is requested, a total summary line is generated for the intervals on each page, and the last page of the report contains a total summary data line that reflects the contents of all pages in the report.

SYSID

The SYSID control statement specifies the one- to four-character system identifier of the single system about which reports are to be generated. It is ignored for SYSRPTS options. The syntax of the statement is:

`SYSID(cccc)`

where **cccc** can be any four alphanumeric and/or special characters that specify the SMF system identifier. When you explicitly specify SYSID, you must supply the system identifier. You can only specify one SYSID control statement per Postprocessor session. Specifying the SYSID control statement causes the Postprocessor to include in the reporting all pertinent SMF records that have a matching system identifier. Omitting SYSID causes the Postprocessor to include in the reporting all SMF records, for all system identifiers. When more than one system identifier is encountered, the Postprocessor produces separate reports for each system encountered. IBM recommends that you do not mix records from different processors with the same system identifier. If you do mix records, the current duration interval for the I/O Queuing duration report will be shortened which will cause the remaining I/O Queuing records from the original duration interval to be skipped.

SYSOUT

The SYSOUT control statement specifies the SYSOUT class for all formatted report output. The syntax of the statement is:

`SYSOUT(class)`

where **class** is the desired SYSOUT class. When you explicitly specify SYSOUT, you must indicate a SYSOUT class. When you omit the SYSOUT statement, the default is SYSOUT class A. The SYSOUT class for Postprocessor messages is not affected by the SYSOUT control statement. The message SYSOUT class can be changed by preallocating MFPMSGDS.

SYSRPTS

The SYSRPTS control statement specifies the sysplex report options for a Postprocessor report.

Note: To get sysplex reports, you have to ensure that data gathering for all systems in the sysplex is synchronized.

Recommendation: Specify the Monitor I option **SYNC(SMF)** for all systems.

You also need to ensure that data from multiple systems is sorted **together** according the information in “Preparing SMF records for postprocessing” on page 202.

The syntax of the statement is:

SYSRPTS(option[,option]...[,option])

where **option** can be the following:

ALL

Specifies the following options: CF - SDEVICE(DASD) - WLMGL(WGPER)

CF | NOCF

Specifies the Coupling Facility Activity report.

To obtain an entire Coupling Facility Activity report, ensure that you supply the Postprocessor with a complete set of SMF 74 subtype 4 records from all systems in the sysplex.

SDEVICE(suboption[,suboption]...[,suboption]) | NOSDEVICE

Specifies whether the Postprocessor should generate Shared Device Activity reports or not.

suboption can be:

DASD | NODASD

Specifies the DASD Shared Device Activity report

TAPE | NOTAPE

Specifies the Magnetic Tape Shared Device Activity report

NMBR(list)

Specifies a list of devices to be included into the report. You can specify as many device numbers as you like. Each element in the list can be:

- A 4-digit device number
- A range of device numbers defined by the lowest and the highest number, separated by a colon. For example, 1234:1243

If the devices belong to the class you specified with the DASD or TAPE option, the NMBR option has no effect. If they belong to the other class, they are reported on *in addition to* the devices of the class you specified.

EXNMBR(list)

Causes the Postprocessor to suppress reports for the device or devices with the device numbers specified. You can specify as many device numbers as you like. Each element in the list can be:

- A 4-digit device number
- A range of device numbers defined by the lowest and the highest number, separated by a colon. For example, 1234:1243

The specified devices are excluded from the set of devices you specified with the DASD, TAPE and NMBR options.

The EXNMBR option has no effect for devices that have not been included in the DASD, TAPE or NMBR options.

WLMGL(suboption[,suboption]...[,suboption]) | NOWLMGL

Specifies whether the Workload Activity report is to be generated.

suboption can have the values listed below. These specify conditions by which the Postprocessor selects the information to be reported.

In the suboptions, **namelist** can be a list of names, or a single name. If you omit **namelist**, the Postprocessor reports on all names that exist for the appropriate condition.

You can specify “wild cards” for names of workload groups, service classes and report classes. A wild card consists of a character string followed immediately by an asterisk (*). Reports are generated for all groups or classes

whose names start with the specified character string. For example, specifying WG1* would produce reports on WG1MINE, WG1YOURS, WG1HIS, WG1HERS, and so on.

POLICY[(namelist)] | NOPOLICY

Specify policy names in namelist. For each policy specified, the Postprocessor issues a summary report.

RCLASS[(namelist)] | NORCLASS

Specify report class names in namelist. The Postprocessor issues reports of the specified classes.

RCPER[(namelist)] | NORCPER

Specify report class names in namelist. The Postprocessor issues a report for each report class period defined for the specified report classes.

RTD | NORTD

Specifies whether the Response Time Distribution section should be displayed in WLMGL Service/Report Class Period reports. This suboption can only be specified together with suboptions SCPER or RCPER; otherwise, it is ignored.

SCLASS[(namelist)] | NOSCLASS

Specify service class names in namelist. For each service class, the Postprocessor issues a summary report.

SCPER[(namelist)] | NOSCPER

Specify service class names in namelist. The Postprocessor issues a report for each service class period defined for the specified service classes. The report includes subsystem states, general execution delays, and a response-time-distribution chart.

SYSNAM[(namelist)]

Specify system names in namelist. The Postprocessor combines data from all the specified systems in one report. If this option is omitted, which is the default, the data from all systems is combined in one report.

WGPER[(namelist)] | NOWGPER

Specify workload group names in namelist. For each workload group you specify, the Postprocessor reports on the associated service classes and their service-class periods.

WGROUP[(namelist)] | NOWGROUP

Specify workload group names. For each workload group specified, the Postprocessor issues a summary report.

Examples of control statements

The examples in this section show various uses of the Postprocessor. All the examples include the DATE statement to illustrate how the value specified for DATE relates to the value specified for RTOD, ETOD, or STOD to define the reporting period. During actual execution of the Postprocessor, your installation might find it more useful to control the dates included in the reports by controlling the contents of the input data set and omitting the DATE statement. Because the default for the DATE statement is a reporting period that encompasses all dates included in the SMF records in the input data set, omitting the DATE statement enables you to establish a set of control statements that can be used on a regular schedule without modification.

Note: Because the EXITS, SYSID, and SYSOUT statements are omitted and their defaults taken in the following examples, no user exits are entered, all systems are included in the reports, and any report and message output is sent to SYSOUT class A.

Single-system report

The Postprocessor is to generate all single-system interval reports except tracing. The reporting period runs from 8:00 AM to 12 noon for the five days from June 6, 2013 to June 10, 2013. Use the following control statements:

Example:

```
DATE(06062013,06102013)
REPORTS(ALL,NOTRACE)
RTOD(0800,1200)
NOSUMMARY
```

Duration report

The Postprocessor is to generate duration reports for CPU activity, channel path activity, and I/O device activity for magnetic tape devices, direct access devices, and communications equipment. The reporting period is the twelve-hour period from 6:00 AM to 6:00 PM on June 24, 2013. The duration interval is six hours, causing two duration reports to be produced for each specified activity. Use the following control statements:

Example:

```
DATE(06242013,06242013)
DINTV(0600)
REPORTS(CPU,CHAN)
REPORTS(DEVICE(TAPE,DASD,COMM))
RTOD(0600,1800)
NOSUMMARY
```

Note: For information on exception/overview duration reporting, refer to “DINTV” on page 215.

Sysplex report

The Postprocessor is to generate sysplex reports. The reporting period runs from 8:00 AM to 6:00 PM for the five days from June 6, 2013 to June 10, 2013. Use the following control statements:

Example:

```
DATE(06062013,06102013)
RTOD(0800,1800)
NOSUMMARY
```

Create a Coupling Facility Activity report:

Example:

```
SYSRPTS(CF)
```

Create a Workload Activity reports and assume that all CICS applications run in the three workload groups CICSPROD, CICSTEST, and CICSADMN. Get the Workload Group report for all groups:

Example:

```
SYSRPTS(WLMGL(WGROUP(CICS*)))
```

Get detailed data for service class TSOPROD by requesting the Service Class Period report:

Example:

```
SYSRPTS(WLMGL(WGPER(TSOPROD)))
```

Get duration reports for two-hour intervals for the Policy Summary report. Assume that only one policy was active during the range to be reported, therefore no policy-name parameter is required.

Example:

```
DINTV(0200)  
SYSRPTS(WLMGL(POLICY))
```

Get a Shared DASD Activity report for all DASDs in the address range 0700 — 071F and 1220 — 123F (the example assumes that only DASD devices are configured in these ranges):

Example:

```
SYSRPTS(SDEVICE(NMBR(0700:071F,1220:123F)))
```

Exception report

The reporting period is the eight-hour interval from 8:00 AM to 4:00 PM for the week beginning June 13, 2013 and ending June 17, 2013:

Example:

```
DATE(06132013,06172013)  
ETOD(0800,1600)
```

A line in the **Exception report** when the percent device utilization for device 06D8 is greater than or equal to 3.

A **Channel Path Activity report** and a Device Activity report for the DASD device class if this condition is met:

Example:

```
EXCEPT(IORATE(DVUTL(06D8),GE,3))  
EXRPTS(IORATE(CHAN,DEVICE(DASD)))
```

A line in the **Exception report** if the total busy percent value for channel path 01 is greater than or equal to ten percent:

Example:

```
EXCEPT(CHANEX01(CHTBSY(01),GE,10))
```

A line in the **Exception report** if the percent busy for CPU 0 is less than or equal to 80 percent busy and the average number of TSO users is greater than ten:

Example:

```
EXCEPT(USERWORK(CPUBSY(0),LE,80))  
EXCEPT(USERWORK(AVGTSO,GE,10))
```

Overview report

The **Overview report** requires that you specify the output format, either as record or report. If you want to get records for further spreadsheet processing and to get a report to be printed, you can specify:

Example:

```
OVERVIEW(RECORD,REPORT)
```

You want to get an overview of the TSO activity in your sysplex for all intervals between 10:00 AM and 2:00 PM. The following control statements assume that all TSO users run in service class TSOSERV and that you have defined three service class periods.

You specify the following control statements:

Example: You use the ETOD statement to specify the time range:

```
ETOD(1000,1400)
```

The exception-condition name T0TSRV specifies the total service units, the qualifier S.TSOSERV refers to service class TSOSERV, and suboption NOSYSTEMS defines sysplex reporting:

```
OVW(SERVUNIT(T0TSRV(S.TSOSERV)),NOSYSTEMS)
```

The exception-condition name RTIMETOT specifies the average response time:

```
OVW(RTIMEP1(RTIMETOT(S.TSOSERV.1)),NOSYSTEMS)
OVW(RTIMEP2(RTIMETOT(S.TSOSERV.2)),NOSYSTEMS)
OVW(RTIMEP3(RTIMETOT(S.TSOSERV.3)),NOSYSTEMS)
```

With the exception-condition name PI, you specify the performance index:

```
OVW(PIP1(PI(S.TSOSERV.1)),NOSYSTEMS)
OVW(PIP2(PI(S.TSOSERV.2)),NOSYSTEMS)
```

The exception-condition name TRANS specifies the transaction rate:

```
OVW(TRXP1(TRANS(S.TSOSERV.1)),NOSYSTEMS)
OVW(TRXP2(TRANS(S.TSOSERV.2)),NOSYSTEMS)
OVW(TRXP3(TRANS(S.TSOSERV.3)),NOSYSTEMS)
```

Cache report

You may want to start with the Cache Summary report which provides a summary of all subsystems and a list of volumes that need special attention. Probably, you would select an interval with a very high I/O activity on your system. If this is during night shift between 8:00 PM and 10:00 PM on May 13, 2013, then you define

Example:

```
DATE(05132013,05132013)
RTOD(2000,2200)
REPORTS(CACHE(SUMMARY))
```

With this information, you can continue by either getting reports with some more details, or by creating an Overview report for the most interesting subsystems and devices.

You get a Subsystem Activity report for the SSIDs 0044 and 0058 with:

Example:

```
REPORTS(CACHE(SSID(0044,0058)))
```

If you see from the list of the top-20 devices that the volumes DATA01 with device number 06F3, DB2PRD with device number 0722, and CICS14 with device number

0734 have the highest cache miss rates in the reported interval, you can create an Overview report with several details for a longer period, for example, for two complete days:

Example:

```
DATE(05132013,05142013)
OVERVIEW(RECORD,REPORT)
OVW(DATA01A(CADRT(06F3)))
OVW(DATA01B(CADRH(06F3)))
OVW(DATA01C(CADSTG(06F3)))
OVW(DB2PRDA(CADRT(0722)))
OVW(DB2PRDB(CADRH(0722)))
OVW(DB2PRDC(CADSTG(0722)))
OVW(CICS14A(CADRT(0734)))
OVW(CICS14B(CADRH(0734)))
OVW(CICS14C(CADSTG(0734)))
```

In this example, the Overview report will provide values about the total I/O rate, the read cache hit rate, and the DASD staging I/O rate. Refer to “OVW” on page 222 for more information about the syntax of the OVW statement.

Overview and exception conditions

The following tables show the condition names that RMF recognizes on OVW and EXCEPT control statements, grouped by SMF record type.

CPU Activity - SMF record type 70-1

One of the following qualifiers is possible:

cluster

Name of the sysplex or cluster

coreid A processor identifier (one or two hexadecimal digits) that either identifies a logical core (when LOADxx PROCVIEW CORE is in effect) or a logical processor (when LOADxx PROCVIEW CPU is in effect).

If the qualifier is omitted, the values represent the average of all logical processors or cores.

cpuid A processor identifier which must be in the format *cpuid*[.*threadid*]

cpuid is a processor identifier (one or two hexadecimal digits) that either identifies a logical core (when LOADxx PROCVIEW CORE is in effect) or a logical processor (when LOADxx PROCVIEW CPU is in effect).

threadid is an optional thread identifier (0 or 1) that identifies a thread that is executing on the logical core designated by *cpuid*. It is ignored when LOADxx PROCVIEW CPU is in effect. If LOADxx PROCVIEW CORE is in effect and *threadid* is omitted, the values represent the average of all threads executing on the logical core.

Examples: 0A, 3F.0, A.1

If the qualifier is omitted, the values represent the average of all logical processors or cores.

lpar Logical partition name

group Group of logical partitions managed towards a common group capacity limit

Table 26. CPU Activity - Conditions Based on SMF Record Type 70–1

Condition	Condition Name	Qualifier	Source	Algorithm
Percent general purpose processor busy	CPUBSY	coreid	SMF70WAT SMF70INT SMF70ONT SMF70PDT	Refer to the field description of “TIME % LPAR BUSY” in the RMF Postprocessor CPU Activity report
Percent zAAP busy	CPUIBSY, AAPBSY	coreid	SMF70WAT SMF70INT SMF70ONT SMF70PDT SMF70_LPAR_BUSY	Same as for CPUBSY but applied to zAAPs
Percent zIIP busy	IIPBSY	coreid	SMF70WAT SMF70INT SMF70ONT SMF70PDT SMF70_LPAR_BUSY	Same as for CPUBSY but applied to zIIPs
Percent MVS busy for general purpose processors	MVSBSY	cpuid	Same as for CPUBSY	Same as for CPUBSY
Percent MVS busy for zAAPs	MVSIBSY, AAPMBSY	cpuid	Same as for MVSBSY	Same as for MVSBSY, but applied to zAAPs
Percent MVS busy for zIIPs	IIPMBSY	cpuid	Same as for MVSBSY	Same as for MVSBSY, but applied to zIIPs
Maximum number of batch users	MXBATCH	none	SMF70BMM	Value or comparison
Maximum number of started tasks	MXSTC	none	SMF70SMM	Value or comparison
Maximum number of TSO users	MXTSO	none	SMF70TMM	Value or comparison
Maximum number of APPC/MVS transaction scheduler (ASCH) users	MXASCH	none	SMF70PMM	Value or comparison
Average number of batch jobs	AVGBATCH	none	SMF70BTT SMF70SAM	BTT/SAM
Average number of started tasks	AVGSTC	none	SMF70STT SMF70SAM	STT/SAM
Average number of TSO users	AVGTSO	none	SMF70TTT SMF70SAM	TTT/SAM
Average number of APPC/MVS transaction scheduler (ASCH) users	AVGASCH	none	SMF70PTT SMF70SAM	PTT/SAM
Average number of in and ready users	AVGIARDY	none	SMF70RTT SMF70SAM	RTT/SAM
Average number of out and ready users	AVGOARDY	none	SMF700TT SMF70SAM	0TT/SAM
Maximum number of OMVS address spaces	MXOMVS	none	SMF70XMM	Value or comparison
Average number of OMVS address spaces	AVGOMVS	none	SMF70XTT SMF70SAM	XTT/SAM
Average number of general purpose processors online during the reporting interval	NUMPROC	none	SMF70ONT SMF70INT	Sum(ONT) / INT
Number of logical zAAP processors or cores online at the end of the reporting interval	NUMIFA, NUMAAP	none	SMF70IFA	Value or comparison
Number of logical zIIP processors or cores online at the end of the reporting interval	NUMIIP	none	SMF70SUP	Value or comparison
Percentage of the reporting interval during which at least one job could not be dispatched	OCPU1	none	SMF70Q01 SMF70Q02 ... SMF70Q12 SMF70SAM	(Q01+Q02+...+Q12)/SAM x 100
Percentage of the reporting interval during which at least two jobs could not be dispatched	OCPU2	none	SMF70Q02 ... SMF70Q12 SMF70SAM	(Q02+Q03+...+Q12)/SAM x 100
Percentage of the reporting interval during which at least three jobs could not be dispatched	OCPU3	none	SMF70Q03 ... SMF70Q12 SMF70SAM	(Q03+Q04+...+Q12)/SAM x 100

Table 26. CPU Activity - Conditions Based on SMF Record Type 70–1 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Percentage of the reporting interval during which at least four jobs could not be dispatched	OCPU4	none	SMF70Q04 ... SMF70Q12 SMF70SAM	$(Q04+Q05+...+Q12)/SAM \times 100$
Percentage of the reporting interval during which more than five jobs could not be dispatched	OCPU5	none	SMF70Q05 ... SMF70Q12 SMF70SAM	$(Q05+Q06+...+Q12)/SAM \times 100$
Percentage of the reporting interval during which more than ten jobs could not be dispatched	OCPU10	none	SMF70Q06 ... SMF70Q12 SMF70SAM	$(Q06+Q07+...+Q12)/SAM \times 100$
Percentage of the reporting interval during which more than 15 jobs could not be dispatched	OCPU15	none	SMF70Q07 ... SMF70Q12 SMF70SAM	$(Q07+Q08+...+Q12)/SAM \times 100$
Percentage of the reporting interval during which more than 20 jobs could not be dispatched	OCPU20	none	SMF70Q08 ... SMF70Q12 SMF70SAM	$(Q08+Q09+...+Q12)/SAM \times 100$
Percentage of the reporting interval during which more than 30 jobs could not be dispatched	OCPU30	none	SMF70Q09 ... SMF70Q12 SMF70SAM	$(Q09+Q10+Q11+Q12)/SAM \times 100$
Percentage of the reporting interval during which more than 40 jobs could not be dispatched	OCPU40	none	SMF70Q10 SMF70Q11 SMF70Q12 SMF70SAM	$(Q10+Q11+Q12)/SAM \times 100$
Percentage of the reporting interval during which more than 60 jobs could not be dispatched	OCPU60	none	SMF70Q11 SMF70Q12 SMF70SAM	$(Q11+Q12)/SAM \times 100$
Percentage of the reporting interval during which more than 80 jobs could not be dispatched	OCPU80	none	SMF70Q12 SMF70SAM	$Q12 / SAM \times 100$
Average number of logical ready users	AVGULRDY	none	SMF70LTT	LTT/SAM
Average number of logical wait users	AVGULWT	none	SMF70ATT	ATT/SAM
Average number of in users	AVGUIN	none	SMF70ITT	ITT/SAM
Online time percentage of general purpose processors	CONTPER	coreid	SMF70ONT SMF70INT	$(ONT/INT) \times 100$
Defined weighting for the cluster	WDEFC	cluster	SMF70BPS	Value or comparison
Defined weighting of the partition for general purpose processors	WDEFL	lpar	SMF70BPS	Value or comparison
Defined weighting of the partition for zAAPs	WDEFLAAP	lpar	SMF70BPS	Value or comparison
Defined weighting of the partition for zIIPs	WDEFLIIP	lpar	SMF70BPS	Value or comparison
Actual weighting of the partition for general purpose processors	WACTL	lpar	SMF70ACS SMF70DSA	ACS / DSA
Minimum weighting of the partition	WMINL	lpar	SMF70MIS	Value or comparison
Minimum weighting percentage of the partition	WMIPL	lpar	SMF70NSI SMF70DSA	$(NSI/DSA) \times 100$
Maximum weighting of the partition	WMAXL	lpar	SMF70MAS	Value or comparison
Maximum weighting percentage of the partition	WMAPL	lpar	SMF70NSA SMF70DSA	$(NSA/DSA) \times 100$
Number of defined logical processors or cores for the cluster	NLDEFC	cluster	SMF70BDN	Value or comparison
Number of defined logical processors or cores for the partition	NLDEFL	lpar	SMF70BDN	Value or comparison
Number of defined general purpose processors for the partition	NLDEFLCP	lpar	SMF70BDN	Value or comparison
Number of defined logical zAAP processors or cores for the partition	NLDEFLAP	lpar	SMF70BDN	Value or comparison

Table 26. CPU Activity - Conditions Based on SMF Record Type 70–1 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Number of defined logical zIIP processors or cores for the partition	NLDEFLIP	lpar	SMF70BDN	Value or comparison
Number of actual logical processors or cores for the partition	NLACTL	lpar	SMF70ONT SMF70INT	Sum(ONT) / INT
Number of actual general purpose processors for the partition	NLACTLCP	lpar	SMF70ONT SMF70INT	Sum(ONT) / INT
Number of actual logical zAAP processors or cores for the partition	NLACTLAP	lpar	SMF70ONT SMF70INT	Sum(ONT) / INT
Number of actual logical zIIP processors or cores for the partition	NLACTLIP	lpar	SMF70ONT SMF70INT	Sum(ONT) / INT
Logical processor busy percentage for general purpose processors for the partition	LBUSYL	lpar	SMF70PDT SMF70ONT	Sum(PDT) / ONT
Logical processor/core busy percentage for zAAPs for the partition	LBUSYLAP	lpar	SMF70PDT SMF70ONT	Sum(PDT) / ONT
Logical processor/core busy percentage for zIIPs for the partition	LBUSYLIP	lpar	SMF70PDT SMF70ONT	Sum(PDT) / ONT
Physical processor busy percentage for general purpose processors for the partition	PBUSYL	lpar	SMF70PDT SMF70INT	Sum(PDT) / INT
Physical processor busy percentage for zAAPs for the partition	PBUSYLAP	lpar	SMF70PDT SMF70INT	Sum(PDT) / INT
Physical processor busy percentage for zIIPs for the partition	PBUSYLIP	lpar	SMF70PDT SMF70INT	Sum(PDT) / INT
Logical processor/core average busy percentage for the cluster	LBUSYC	cluster	SMF70PDT SMF70ONT	Sum(PDT) / ONT
Physical processor average busy percentage for the cluster	PBUSYC	cluster	SMF70PDT SMF70INT	Sum(PDT) / INT
Defined capacity limit in units of MSU	LDEFMSU	lpar	SMF70MSU	Value or comparison
Actual number of consumed MSUs	LACTMSU	lpar	SMF70PDT SMF70CPA_ actual SMF70CPA_ scaling_factor SMF70INT	$(\text{Sum PDT}) * 3600 * 16 * \text{CPA_scaling_factor} / (\text{CPA_actual} * \text{INT} * 1000000)$
Percentage of WLM capping of the partition	WCAPPER	lpar	SMF70NSW SMF70DSA	$(\text{NSW} / \text{DSA}) * 100$
Actual MSU consumption for the capacity group	GCMSUACT	group	SMF70PDT SMF70CPA_ actual SMF70CPA_ scaling_factor SMF70INT	$\text{Sum}(\text{Sum(PDT)} * 3600 * 16 * \text{CPA_scaling_factor} / (\text{CPA_actual} * \text{INT} * 1000000))$
Group weight	GCWEIGHT	group	SMF70PFL SMF70BPS SMF70ACS	Sum(BPS) if PFL(Bit 3) is set for all partitions in capacity group, otherwise Sum(ACS)
Minimum entitlement	MINENT	lpar	SMF70PFL SMF70BPS SMF70ACS SMF70GMU SMF70MSU	$\text{MIN}(\text{MSU}, \text{GMU} * \text{BPS} / \text{Sum(BPS)})$ if PFL(Bit 3) is set for all partitions in capacity group, otherwise $\text{MIN}(\text{MSU}, \text{GMU} * \text{ACS} / \text{Sum(ACS)})$
Maximum entitlement	MAXENT	lpar	SMF70GMU SMF70MSU	$\text{MIN}(\text{GMU}, \text{MSU})$
Number of TCB dispatches on general purpose processors per second	CPTCB	cpuid	SMF70TCB SMF70INT	Sum(TCB) / INT
Number of TCB dispatches on zAAPs per second	AAPTCB	cpuid	SMF70TCB SMF70INT	Same as CPTCB but applied to zAAPs
Number of TCB dispatches on zIIPs per second	IIPTCB	cpuid	SMF70TCB SMF70INT	Same as CPTCB but applied to zIIPs
Number of SRB dispatches on general purpose processors per second	CPSRB	cpuid	SMF70SRB SMF70INT	Sum(SRB) / INT
Number of SRB dispatches on zAAPs per second	AAPSRB	cpuid	SMF70SRB SMF70INT	Same as CPSRB but applied to zAAPs
Number of SRB dispatches on zIIPs per second	IIPSRB	cpuid	SMF70SRB SMF70INT	Same as CPSRB but applied to zIIPs

Table 26. CPU Activity - Conditions Based on SMF Record Type 70–1 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Number of I/Os requested by general purpose processors per second	CPNIO	cpuid	SMF70NIO SMF70INT	Sum(NIO) / INT
Number of I/Os requested by zAAPs per second	AAPNIO	cpuid	SMF70NIO SMF70INT	Same as CPNIO but applied to zAAPs
Number of I/Os requested by zIIPs per second	IIPNIO	cpuid	SMF70NIO SMF70INT	Same as CPNIO but applied to zIIPs
Number of CPs with high HiperDispatch share for partition	HDCPHIGH	lpar	SMF70POI SMF70CIX	Number of logical processors with POI=3 and CIX=1
Number of logical zAAP processors or cores with high HiperDispatch share for partition	HDAPHIGH	lpar	SMF70POI SMF70CIX	Number of logical processors with POI=3 and CIX=3
Number of logical zIIP processors or cores with high HiperDispatch share for partition	HDIPHIGH	lpar	SMF70POI SMF70CIX	Number of logical processors with POI=3 and CIX=6
Number of CPs with medium HiperDispatch share for partition	HDCPMED	lpar	SMF70POI SMF70CIX	Number of logical processors with POI=2 and CIX=1
Number of logical zAAP processors or cores with medium HiperDispatch share for partition	HDAPMED	lpar	SMF70POI SMF70CIX	Number of logical processors with POI=2 and CIX=3
Number of logical zIIP processors or cores with medium HiperDispatch share for partition	HDIPMED	lpar	SMF70POI SMF70CIX	Number of logical processors with POI=2 and CIX=6
Number of CPs with low HiperDispatch share for partition	HDCPLow	lpar	SMF70POI SMF70CIX	Number of logical processors with POI=1 and CIX=1
Number of logical zAAP processors or cores with low HiperDispatch share for partition	HDAPLOW	lpar	SMF70POI SMF70CIX	Number of logical processors with POI=1 and CIX=3
Number of logical zIIP processors or cores with low HiperDispatch share for partition	HDIPLOW	lpar	SMF70POI SMF70CIX	Number of logical processors with POI=1 and CIX=6
The percentage of time that the general purpose processor was parked.	CPARKPER	cpuid	SMF70PAT SMF70INT	(PAT/INT) * 100
Maximum number of in-ready work units for general purpose processors	MXWUCP	none	SMF70CMM	Value or comparison
Maximum number of in-ready work units for zAAPs	MXWUAAP	none	SMF70DMM	Value or comparison
Maximum number of in-ready work units for zIIPs	MXWUIIP	none	SMF70EMM	Value or comparison
Average number of in-ready work units for general purpose processors	AVGWUCP	none	SMF70CTT SMF70SRM	CTT / SRM
Average number of in-ready work units for zAAPs	AVGWUAAP	none	SMF70DTT SMF70SRM	DTT / SRM
Average number of in-ready work units for zIIPs	AVGWUIIP	none	SMF70ETT SMF70SRM	ETT / SRM
Mean time to wait for general purpose processors in microseconds	MTTWCP	cpuid	SMF70EDT SMF70WTD	EDT / WTD or EDT (if WTD=0)
Mean time to wait for zAAPs in microseconds	MTTWAAP	cpuid	SMF70EDT SMF70WTD	same as MTTWCP, but applied to zAAPs
Mean time to wait for zIIPs in microseconds	MTTWIIP	cpuid	SMF70EDT SMF70WTD	same as MTTWCP, but applied to zIIPs
Percentage of the reporting interval during which at least one work unit could not be dispatched	WCPU1	none	SMF70U01 ... SMF70U15 SMF70SRM	(U01+...+U15) / SRM x 100
Percentage of the reporting interval during which at least two work units could not be dispatched	WCPU2	none	SMF70U02 ... SMF70U15 SMF70SRM	(U02+...+U15) / SRM x 100
Percentage of the reporting interval during which at least three work units could not be dispatched	WCPU3	none	SMF70U03 ... SMF70U15 SMF70SRM	(U03+...+U15) / SRM x 100

Table 26. CPU Activity - Conditions Based on SMF Record Type 70–1 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Percentage of the reporting interval during which at least four work units could not be dispatched	WCPU4	none	SMF70U04 ... SMF70U15 SMF70SRM	$(U04+...+U15) / SRM \times 100$
Percentage of the reporting interval during which at least five work units could not be dispatched	WCPU5	none	SMF70U05 ... SMF70U15 SMF70SRM	$(U05+...+U15) / SRM \times 100$
Percentage of the reporting interval during which at least ten work units could not be dispatched	WCPU10	none	SMF70U06 ... SMF70U15 SMF70SRM	$(U06+...+U15) / SRM \times 100$
Percentage of the reporting interval during which at least 15 work units could not be dispatched	WCPU15	none	SMF70U07 ... SMF70U15 SMF70SRM	$(U07+...+U15) / SRM \times 100$
Percentage of the reporting interval during which at least 20 work units could not be dispatched	WCPU20	none	SMF70U08 ... SMF70U15 SMF70SRM	$(U08+...+U15) / SRM \times 100$
Percentage of the reporting interval during which at least 30 work units could not be dispatched	WCPU30	none	SMF70U09 ... SMF70U15 SMF70SRM	$(U09+...+U15) / SRM \times 100$
Percentage of the reporting interval during which at least 40 work units could not be dispatched	WCPU40	none	SMF70U10 ... SMF70U15 SMF70SRM	$(U10+...+U15) / SRM \times 100$
Percentage of the reporting interval during which at least 60 work units could not be dispatched	WCPU60	none	SMF70U11 ... SMF70U15 SMF70SRM	$(U11+...+U15) / SRM \times 100$
Percentage of the reporting interval during which at least 80 work units could not be dispatched	WCPU80	none	SMF70U13 SMF70U12 ... SMF70U15 SMF70SRM	$(U12+...+U15) / SRM \times 100$
Percentage of the reporting interval during which at least 100 work units could not be dispatched	WCPU100	none	SMF70U13 ... SMF70U15 SMF70SRM	$(U13+...+U15) / SRM \times 100$
Percentage of the reporting interval during which at least 120 work units could not be dispatched	WCPU120	none	SMF70U14 SMF70U15 SMF70SRM	$(U14+U15) / SRM \times 100$
Percentage of the reporting interval during which at least 150 work units could not be dispatched	WCPU150	none	SMF70U15 SMF70SRM	$U15 / SRM \times 100$
Nominal processor capacity available to the CPC	NOMCAPAC	none	SMF70NCR	Value or comparison
Effective capacity percentage	EFFCAPAC	none	SMF70MCR SMF70NCR	$(MCR / NCR) * 100$
Long-term average of CPU service (millions of service units) at the end of the reporting interval	LACS	none	SMF70LAC	Value or comparison
The percentage of times PR/SM™ issued a warning-track interruption to a general purpose processor and z/OS was able to return it to PR/SM within the grace period	WTRKCP	cpuid	SMF70WTS SMF70WTU	$WTS / (WTS + WTU)$
The percentage of times PR/SM issued a warning-track interruption to a zAAP and z/OS was able to return it to PR/SM within the grace period	WTRKAAP	cpuid	SMF70WTS SMF70WTU	Same as WTRKCP but applied to zAAP

Table 26. CPU Activity - Conditions Based on SMF Record Type 70–1 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
The percentage of times PR/SM issued a warning-track interruption to a zIIP and z/OS was able to return it to PR/SM within the grace period	WTRKIIP	cpuid	SMF70WTS SMF70WTU	Same as WTRKCP but applied to zIIP
Time in milliseconds that a general purpose processor was yielded to PR/SM due to warning-track processing	WTRKTCP	cpuid	SMF70WTI	Value or comparison
Time in milliseconds that a zAAP was yielded to PR/SM due to warning-track processing	WTRKTAAP	cpuid	SMF70WTI	Value or comparison
Time in milliseconds that a zIIP was yielded to PR/SM due to warning-track processing	WTRKTIIP	cpuid	SMF70WTI	Value or comparison
Initial Capping for general purpose processors	INICAP	lpar	SMF70VPF	1 if VPF(Bit 3) is set, otherwise 0
Absolute physical hardware capacity limit in numbers of CPUs for general purpose processors	LIMCPU	lpar	SMF70HW_Cap_Limit	HW_Cap_Limit/100
Percent multithreading core productivity for zIIPs	IIPPROD	coreid	SMF70_PROD	PROD / 1024
Percent multithreading core productivity for zAAPs	AAPPROD	coreid	SMF70_PROD	PROD / 1024
Percent multithreading core utilization for zIIPs	IIPUTIL	coreid	SMF70_PROD SMF70PDT SMF70ONT SMF70_LPAR_BUSY	PROD / 1024 multiplied by value of Overview Condition IIPBSY
Percent multithreading core utilization for zAAPs	AAPUTIL	coreid	SMF70_PROD SMF70PDT SMF70ONT SMF70_LPAR_BUSY	PROD / 1024 multiplied by value of Overview Condition AAPBSY

Crypto Hardware Activity - SMF record type 70-2

One of the following qualifiers is possible:

- ccid** cryptographic coprocessor index (one or two decimal digits). If the qualifier is omitted, performance data is reported for coprocessor index 0.
- caid** cryptographic accelerator index (one or two decimal digits). If the qualifier is omitted, performance data is reported for accelerator index 0.

Table 27. Crypto Hardware Activity - Conditions Based on SMF Record Type 70–2

Condition	Condition Name	Qualifier	Source	Algorithm
Cryptographic coprocessor total rate	CRYCTR	ccid	R7023C0 SMF70INT	3C0 / INT
Cryptographic coprocessor total utilization	CRYCTU	ccid	R7023T0 R7023SF SMF70INT	3T0 * 3SF * 100 / INT
Cryptographic coprocessor total avg execution time	CRYCTE	ccid	R7023C0 R7023T0 R7023SF	3T0 * 3SF / 3C0
Cryptographic coprocessor key-gen rate	CRYCKR	ccid	R7023C1 SMF70INT	3C1 / INT
Cryptographic accelerator 1024bit-ME rate	CRYAM1R	caid	R7021MEC SMF70INT	Sum(1MEC) / INT
Cryptographic accelerator 1024bit-ME utilization	CRYAM1U	caid	R7024EN R7021MET R7024SF SMF70INT	Sum(1MET) * 4SF * 100 / (INT * 4EN)
Cryptographic accelerator 1024bit-ME avg execution time	CRYAM1E	caid	R7021MEC R7021MET R7024SF	Sum(1MET) * 4SF / Sum(1MEC)

Table 27. Crypto Hardware Activity - Conditions Based on SMF Record Type 70–2 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Cryptographic accelerator 2048bit-ME rate	CRYAM2R	caid	R7022MEC SMF70INT	Sum(2MEC) / INT
Cryptographic accelerator 2048bit-ME utilization	CRYAM2U	caid	R7024EN R7022MET R7024SF SMF70INT	Sum(2MET) * 4SF * 100 / (INT * 4EN)
Cryptographic accelerator 2048bit-ME avg execution time	CRYAM2E	caid	R7022MEC R7022MET R7024SF	Sum(2MET) * 4SF / Sum(2MEC)
Cryptographic accelerator 4096bit-ME rate	CRYAM3R	caid	R7023MEC SMF70INT	Sum(3MEC) / INT
Cryptographic accelerator 4096bit-ME utilization	CRYAM3U	caid	R7024EN R7023MET R7024SF SMF70INT	Sum(3MET) * 4SF * 100 / (INT * 4EN)
Cryptographic accelerator 4096bit-ME avg execution time	CRYAM3E	caid	R7023MEC R7023MET R7024SF	Sum(3MET) * 4SF / Sum(3MEC)
Cryptographic accelerator 1024bit-CRT rate	CRYAC1R	caid	R7021CRC SMF70INT	Sum(1CRC) / INT
Cryptographic accelerator 1024bit-CRT utilization	CRYAC1U	caid	R7021CRT R7024EN R7024SF SMF70INT	Sum(1CRT) * 4SF * 100 / (INT * 4EN)
Cryptographic accelerator 1024bit-CRT avg execution time	CRYAC1E	caid	R7021CRC R7021CRT R7024SF	Sum(1CRT) * 4SF / Sum(1CRC)
Cryptographic accelerator 2048bit-CRT rate	CRYAC2R	caid	R7022CRC SMF70INT	Sum(2CRC) / INT
Cryptographic accelerator 2048bit-CRT utilization	CRYAC2U	caid	R7022CRT R7024EN R7024SF SMF70INT	Sum(2CRT) * 4SF * 100 / (INT * 4EN)
Cryptographic accelerator 2048bit-CRT avg execution time	CRYAC2E	caid	R7022CRC R7022CRT R7024SF	Sum(2CRT) * 4SF / Sum(2CRC)
Cryptographic accelerator 4096bit-CRT rate	CRYAC3R	caid	R7023CRC SMF70INT	Sum(3CRC) / INT
Cryptographic accelerator 4096bit-CRT utilization	CRYAC3U	caid	R7023CRT R7024EN R7024SF SMF70INT	Sum(3CRT) * 4SF * 100 / (INT * 4EN)
Cryptographic accelerator 4096bit-CRT avg execution time	CRYAC3E	caid	R7023CRC R7023CRT R7024SF	Sum(3CRT) * 4SF / Sum(3CRC)
Single DES encryption rate	CRYISDER	none	R702SNEC SMF70INT	SNEC / INT
Single DES encryption size	CRYISDES	none	R702SNEB R702SNEC	SNEB / SNEC
Single DES number of instructions used to encipher the data	CRYISDEI	none	R702SNEI	
Triple DES encryption rate	CRYITDER	none	R702TNEC SMF70INT	TNEC / INT
Triple DES encryption size	CRYITDES	none	R702TNEB R702TNEC	TNEB / TNEC
Triple DES number of instructions used to encipher the data	CRYITDEI	none	TNEI	
Single DES decryption rate	CRYISDDR	none	R702SNDC SMF70INT	SNDC / INT
Single DES decryption size	CRYISDDS	none	R702SNDB R702SNDC	SNDB / SNDC

Table 27. Crypto Hardware Activity - Conditions Based on SMF Record Type 70–2 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Single DES number of instructions used to decipher the data	CRYISDDI	none	R702SNDI	
Triple DES decryption rate	CRYITDDR	none	R702TNDC SMF70INT	TNDC / INT
Triple DES decryption size	CRYITDDS	none	R702TNDB R702TNDC	TNDB / TNDC
Triple DES number of instructions used to decipher the data	CRYITDDI	none	R702TNDI	
Rate of AES encryption service calls sent to a coprocessor	CRYIAER	none	R702AESC SMF70INT	AESC / INT
Rate of AES decryption service calls sent to a coprocessor	CRYIADR	none	R702ASDC SMF70INT	ASDC / INT
Average number of bytes processed per AES encryption service call handled by a coprocessor	CRYIAES	none	R702AESB R702AESC	AESB / AESC
Average number of bytes processed per AES decryption service call handled by a coprocessor	CRYIADS	none	R702ASDB R702ASDC	ASDB / ASDC
Average number of times the coprocessor was called to complete the AES encryption service calls	CRYIAEO	none	R702AESI R702AESC	AESI / AESC
Average number of times the coprocessor was called to complete the AES decryption service calls	CRYIADO	none	R702ASDI R702ASDC	ASDI / ASDC
MAC generation rate	CRYIMGR	none	R702NMGC SMF70INT	NMGC / INT
MAC generation size	CRYIMGS	none	R702NMGB R702NMGC	NMGB / NMGC
Number of instructions used to MAC generate	CRYIMGI	none	R702NMGI	
MAC verify rate	CRYIMVR	none	R702NMVC SMF70INT	NMVC / INT
MAC verify size	CRYIMVS	none	R702NMVB R702NMVC	NMVB / NMVC
Number of instructions used to MAC verify	CRYIMVI	none	R702NMVI	
Hashing rate using the SHA-1 algorithm	CRYIHAR	none	R702NHAC SMF70INT	NHAC / INT
Hashing size using the SHA-1 algorithm	CRYIHAS	none	R702NHAB R702NHAC	NHAB / NHAC
Number of instructions used to hash data with the SHA-1 algorithm	CRYIHAI	none	R702NHAI	Value or comparison
Hashing rate using the SHA-224 or the SHA-256 algorithm	CRYIH2R	none	R702NH2C SMF70INT	NH2C / INT
Hashing size using the SHA-224 or the SHA-256 algorithm	CRYIH2S	none	R702NH2B R702NH2C	NH2B / NH2C
Number of instructions used to hash data with the SHA-224 or the SHA-256 algorithm	CRYIH2I	none	R702NH2I	Value or comparison
Hashing rate using the SHA-384 or the SHA-512 algorithm	CRYIH5R	none	R702NH5C SMF70INT	NH5C / INT
Hashing size using the SHA-384 or the SHA-512 algorithm	CRYIH5S	none	R702NH5B R702NH5C	NH5B / NH5C
PIN translation rate	CRYIPTR	none	R702NPTC SMF70INT	NPTC / INT
PIN verify rate	CRYIPVR	none	R702NPVC SMF70INT	NPVC / INT
AES MAC generation rate	CRYIAMGR	none	R702AMGC SMF70INT	AMGC/INT

Table 27. Crypto Hardware Activity - Conditions Based on SMF Record Type 70–2 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
AES MAC generation size	CRYIAMGS	none	R702AMGB R702AMGC	AMGB / AMGC
Number of instructions used to generate AES MACs	CRYIAMGI	none	R702AMGI	Value or comparison
AES MAC verify rate	CRYIAMVR	none	R702AMVC SMF70INT	AMVC/INT
AES MAC verify size	CRYIAMVS	none	R702AMVB R702AMVC	AMVB/AMVC
Number of instructions used to verify AES MACs	CRYIAMVI	none	R702AMVI	Value or comparison
RSA digital signature generation rate	CRYIDRGR	none	R702DRGC SMF70INT	DRGC/INT
RSA digital signature verify rate	CRYIDRVR	none	R702DRVC SMF70INT	DRVC/INT
ECC digital signature generation rate	CRYIDEGR	none	R702DEGC SMF70INT	DEGC/INT
ECC digital signature verify rate	CRYIDEVVR	none	R702DEVVC SMF70INT	DEVVC/INT
FPE encipher rate	CRYIFPER	none	R702FPEC SMF70INT	FPEC / INT
FPE encipher size	CRYIFPES	none	R702FPFB R702FPEC	FPFB / FPEC
Number of instructions used to encipher data using FPE	CRYIFPEI	none	R702FPEI	Value or comparison
FPE decipher rate	CRYIFPDR	none	R702FPDC SMF70INT	FPDC / INT
FPE decipher size	CRYIFPDS	none	R702FPDB R702FPDC	FPDB / FPDC
Number of instructions used to decipher data using FPE	CRYIFPDI	none	R702FPDI	Value or comparison
FPE translate rate	CRYIFPTR	none	R702FPTC SMF70INT	FPTC / INT
FPE translate size	CRYIFPTS	none	R702FPTB R702FPTC	FPTB / FPTC
Number of instructions used to translate data using FPE	CRYIFPTI	none	R702FPTI	Value or comparison
Cryptographic PKCS11 coprocessor total rate	CRYPTR	ccid	R7025SAC R7025FAC R7025SPC R7025SCC R7025AGC SMF70INT	(5SAC+5FAC+ 5SPC+ 5SCC+5AGC)/INT
Cryptographic PKCS11 coprocessor total utilization	CRYPTU	ccid	R7025FAT R7025AGT R7025SAT R7025SPT R7025SCT R7025SF SMF70INT	(5SAT+5FAT+5SPT+ 5SCT+5AGT)*5SF*100/INT
Cryptographic PKCS11 coprocessor total average execution time (in milliseconds)	CRYPTE	ccid	R7025SAC R7025FAC R7025SPC R7025SCC R7025AGC R7025SAT R7025FAT R7025SPT R7025SCT R7025AGT SMF70INT	(5SAT+5FAT+5SPT+ 5SCT+5AGT)*5SF*1000/ (5SAC+5FAC+ 5SPC+ 5SCC+5AGC)
Rate of operations executed by slow asymmetric-key functions	CRYP SAR	ccid	R7025SAC SMF70INT	5SAC/INT

Table 27. Crypto Hardware Activity - Conditions Based on SMF Record Type 70–2 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Utilization of operations executed by slow asymmetric-key functions	CRYPSAU	ccid	R7025SAT R7025SF SMF70INT	5SAT*5SF*100/INT
Average execution time of operations executed by slow asymmetric-key functions (in milliseconds)	CRYPSAE	ccid	R7025SAC R7025SAT R7025SF	5SAT*5SF*1000/5SAC
Rate of operations executed by fast asymmetric-key functions	CRYPFAR	ccid	R7025FAC SMF70INT	5FAC/INT
Utilization of operations executed by fast asymmetric-key functions	CRYPFAU	ccid	R7025FAT R7025SF SMF70INT	5FAT*5SF*100/INT
Average execution time of operations executed by fast asymmetric-key functions (in milliseconds)	CRYPFAE	ccid	R7025FAC R7025FAT R7025SF	5FAT*5SF*1000/5FAC
Rate of operations executed by symmetric-key functions that return partial or incremental results	CRYPSPR	ccid	R7025SPC SMF70INT	5SPC/INT
Utilization of operations executed by symmetric-key functions that return partial or incremental results	CRYPSPU	ccid	R7025SPT R7025SF SMF70INT	5SPT*5SF*100/INT
Average execution time of operations executed by symmetric-key functions that return partial or incremental results (in milliseconds)	CRYPSPE	ccid	R7025SPC R7025SPT R7025SF	5SPT*5SF*1000/5SPC
Rate of operations executed by symmetric-key functions that return a complete or final result	CRYPSCR	ccid	R7025SCC SMF70INT	5SCC/INT
Utilization of operations executed by symmetric-key functions that return a complete or final result	CRYPSCU	ccid	R7025SCT R7025SF SMF70INT	5SCT*5SF*100/INT
Average execution time of operations executed by symmetric-key functions that return a complete or final result (in milliseconds)	CRYPSCCE	ccid	R7025SCC R7025SCT R7025SF	5SCT*5SF*1000/5SCC
Rate of operations executed by asymmetric-key generation function	CRYPAGR	ccid	R7025AGC SMF70INT	5AGC/INT
Utilization of operations executed by asymmetric-key generation function	CRYPAGU	ccid	R7025AGT R7025SF SMF70INT	5AGT*5SF*100/INT
Average execution time of operations executed by asymmetric-key generation function (in milliseconds)	CRYPAGE	ccid	R7025AGC R7025AGT R7025SF	5AGT*5SF*1000/5AGC

Paging Activity - SMF record type 71

Table 28. Paging Activity - Conditions Based on SMF Record Type 71

Condition	Condition Name	Qualifier	Source	Algorithm
Size of central storage (K)	STORAGE	none	SMF71TFC SMF71FIN	TFC+FIN
Total number of pages per second	TPAGRT	none	SMF71PIN SMF71POT SMF71SIN SMF71SOT SMF71VIN SMF71VOT SMF71BLP SMF71INT SMF71HOT SMF71HIN	(PIN+POT+SIN+SOT +VIN+VOT+BLP +HOT+HIN)/INT
Number of page faults per second	PAGERT	none	SMF71PIN SMF71INT	PIN/INT

Table 28. Paging Activity - Conditions Based on SMF Record Type 71 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Demand paging per second	DPAGRT	none	SMF71PIN SMF71POT SMF71INT	(PIN+POT)/INT
Percent successful swap-out	PLSWAPOU	none	SMF71TOT(k) SMF71AXD(k) SMF71ESD(k) SMF71LES(k) SMF71LAX(k)	$(TOT(k)-AXD(k)-ESD(k)-LES(k)-LAX(k))*100 / (TOT(k)-AXD(k)-ESD(k))$
Maximum number of SQA fixed frames	MXSQA	none	SMF71MXQ	Value or comparison
Average number of SQA fixed frames	AVGSQA	none	SMF71AVQ	Value or comparison
Maximum number of CSA fixed-frames	MXCSAF	none	SMF71MXC	Value or comparison
Average number of CSA fixed frames	AVGCSAF	none	SMF71AVC	Value or comparison
Maximum number of central storage CSA frames	MXCSAT	none	SMF71MXP	Value or comparison
Average number of central storage CSA frames	AVGCSAT	none	SMF71AVP	Value or comparison
Maximum number of VIO allocated local page data set slots	MAXVIOF	none	SMF71MXV	Value or comparison
Average number of VIO allocated local page data set slots	AVGVIOF	none	SMF71LVV	Value or comparison
Page move rate	PGMVRT	none	SMF71PMV SMF71INT	PMV/INT
Average high unreferenced interval count for central storage frames	AVGHUIC	none	SMF71UAC	Value or comparison
Maximum high unreferenced interval count for central storage frames	MXHUIC	none	SMF71UHC	Value or comparison
Page movement rate to expanded storage	PTES	none	SMF71PES	PES/INT
Migration rate from expanded to auxiliary storage	ESMRLW	none	SMF71PEA	PEA/INT
Minimum number of available CS frames	CSTORAVM	none	SMF71CAM	Value or comparison
Maximum number of available CS frames	CSTORAVX	none	SMF71CAX	Value or comparison
Average number of available CS frames	CSTORAVA	none	SMF71CAA	Value or comparison
Minimum number of low-impact CS frames	CSTORLIM	none	SMF71CLM	Value or comparison
Maximum number of low-impact CS frames	CSTORLIX	none	SMF71CLX	Value or comparison
Average number of low-impact CS frames	CSTORLIA	none	SMF71CLA	Value or comparison
Minimum number of medium-impact CS frames	CSTORMIM	none	SMF71CMM	Value or comparison
Maximum number of medium-impact CS frames	CSTORMIX	none	SMF71CMX	Value or comparison
Average number of medium-impact CS frames	CSTORMIA	none	SMF71CMA	Value or comparison
Minimum number of high-impact CS frames	CSTORHIM	none	SMF71CHM	Value or comparison
Maximum number of high-impact CS frames	CSTORHIX	none	SMF71CHX	Value or comparison
Average number of high-impact CS frames	CSTORHIA	none	SMF71CHA	Value or comparison
Minimum number of available Expanded Storage frames	ESTORAVM	none	SMF71EAM	Value or comparison
Maximum number of available Expanded Storage frames	ESTORAVX	none	SMF71EAX	Value or comparison
Average number of available Expanded Storage frames	ESTORAVA	none	SMF71EAA	Value or comparison
Minimum number of low-impact Expanded Storage frames	ESTORLIM	none	SMF71ELM	Value or comparison
Maximum number of low-impact Expanded Storage frames	ESTORLIX	none	SMF71ELX	Value or comparison
Average number of low-impact Expanded Storage frames	ESTORLIA	none	SMF71ELA	Value or comparison
Minimum number of medium-impact Expanded Storage frames	ESTORMIM	none	SMF71EMM	Value or comparison
Maximum number of medium-impact Expanded Storage frames	ESTORMIX	none	SMF71EMX	Value or comparison
Average number of medium-impact Expanded Storage frames	ESTORMIA	none	SMF71EMA	Value or comparison
Minimum number of high-impact Expanded Storage frames	ESTORHIM	none	SMF71EHM	Value or comparison
Maximum number of high-impact Expanded Storage frames	ESTORHIX	none	SMF71EHX	Value or comparison
Average number of high-impact Expanded Storage frames	ESTORHIA	none	SMF71EHA	Value or comparison
Minimum number of VIO pages in central storage	RSVIOM	none	SMF71MVI	Value or comparison
Maximum number of VIO pages in central storage	RSVIOX	none	SMF71XVI	Value or comparison
Average number of VIO pages in central storage	RSVIOA	none	SMF71AVI	Value or comparison

Table 28. Paging Activity - Conditions Based on SMF Record Type 71 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Minimum number of hyperspace pages in central storage	RSHSPM	none	SMF71MHI	Value or comparison
Maximum number of hyperspace pages in central storage	RSHSPX	none	SMF71XHI	Value or comparison
Average number of hyperspace pages in central storage	RSHSPA	none	SMF71AHI	Value or comparison
Number of VIO pages written to central storage	RSVIOW	none	SMF71VWS	Value or comparison
Number of VIO pages read from central storage	RSVIOR	none	SMF71VRS	Value or comparison
Number of hyperspace pages written to central storage	RSHSPW	none	SMF71HWS	Value or comparison
Number of hyperspace pages read from central storage	RSHSPR	none	SMF71HRS	Value or comparison
Minimum number of pages fixed between 16MB and 2GB	FXBETWM	none	SMF71MFB	Value or comparison
Maximum number of pages fixed between 16MB and 2GB	FXBETWX	none	SMF71XFB	Value or comparison
Average number of pages fixed between 16MB and 2GB	FXBETWA	none	SMF71AFB	Value or comparison
Average number of shared page groups in the system	SHRPT	none	SMF71AGT	Value or comparison
Average number of shared page groups in central storage	SHRPC	none	SMF71AGC	Value or comparison
Average number of shared page groups in auxiliary storage	SHRPA	none	SMF71AGA	Value or comparison
Average number of shared pages fixed	SHRPF	none	SMF71AGF	Value or comparison
Average number of shared pages fixed below 16 MB	SHRPB	none	SMF71AGB	Value or comparison
Number of page-ins from Auxiliary Storage for shared pages	SHRPI	none	SMF71ASI	Value or comparison
Number of page-outs to Auxiliary Storage for shared pages	SHRPO	none	SMF71ASO	Value or comparison
Average number of shared pages in the system with a virtual storage address above the bar	SHRPTH	none	SMF71PTH	Value or comparison
Average number of shared pages in Central Storage with a virtual storage address above the bar	SHRPCH	none	SMF71PCH	Value or comparison
Average number of shared pages in Auxiliary Storage with a virtual storage address above the 2GB bar	SHRPAH	none	SMF71PAH	Value or comparison
Peak number of shared pages from virtual storage above the 2GB bar	SHRPBLG	none	SMF71BLG	Value or comparison
Number of page-ins from Auxiliary Storage for shared pages with a virtual storage address above the bar	SHRPIH	none	SMF71PIH	Value or comparison
Number of page-outs to Auxiliary Storage for shared pages with a virtual storage address above the bar	SHRPOH	none	SMF71POH	Value or comparison
Average number of memory objects allocated in the high virtual common storage of the system	CMOA	none	SMF71COA	Value or comparison
Average number of memory objects allocated in the high virtual shared storage of the system	SMOA	none	SMF71SOA	Value or comparison
Average number of fixed memory objects that are allocated in the system and can be backed in 1 MB frames (only available with Enhanced DAT Architecture)	LMOA	none	SMF71LOA	Value or comparison
Average number of 1 MB pages backed in central storage	LFRA	none	SMF71LRA	Value or comparison
Average number of pages from high virtual common storage that are backed in central storage (in units of 4 KB)	CFRA	none	SMF71CRA	Value or comparison
Average number of fixed pages from high virtual common storage that are backed in central storage (in units of 4 KB)	CFRA	none	SMF71CFA	Value or comparison
Average number of pages from high virtual shared storage that are backed in central storage (in units of 4 KB)	SFRA	none	SMF71SRA	Value or comparison
Average number of auxiliary storage slots used for high virtual common pages backed on DASD	CAUXSA	none	SMF71CSA	Value or comparison
Total number of logical swaps	LSWAPTOT	none	SMF71TLS	Value or comparison
Total number of 1 MB frames that can be used by fixed memory objects	LFFRTA	none	SMF71L1A	Value or comparison
Average number of 1 MB frames in the Large Frame Area that are in-use by fixed memory objects	LFFRUA	none	SMF71L3A	Value or comparison
Total number of 1 MB frames that can be used by pageable and DREF memory objects	LPFRTA	none	SMF71L4A	Value or comparison
Average number of 1 MB frames that are not in-use by pageable and DREF memory objects	LPFRAA	none	SMF71L5A	Value or comparison
Average number of 1 MB frames that are in-use by pageable and DREF memory objects	LPFRUA	none	SMF71L6A	Value or comparison

Table 28. Paging Activity - Conditions Based on SMF Record Type 71 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Average number of 1 MB frames in the Large Frame Area that are not in-use	LFFRNA	none	SMF71L7A	Value or comparison
Average total number of high virtual shared storage pages	SFRTA	none	SMF71S1A	Value or comparison
Average number of auxiliary storage slots used for high virtual shared pages that are backed on DASD	SAUXSA	none	SMF71S5A	Value or comparison
Average number of auxiliary storage slots used for high virtual shared pages that are backed on SCM storage	SAUXSSA	none	SMF71S6A	Value or comparison
Average total number of high virtual common storage pages	CFRTA	none	SMF71C1A	Value or comparison
Average number of auxiliary storage slots used for high virtual common memory pages that are backed on SCM storage	CAUXSSA	none	SMF71C4A	Value or comparison
Average number of shared memory pages backed on SCM storage	SHRPASCM	none	SMF71S7A	Value or comparison

Workload Activity - SMF record type 72-3

The following table is valid only for overview processing, not for exception reporting. Depending on the OVW suboption SYSTEMS|NOSYSTEMS, reports or records will be created for each single system in addition to sysplex reporting.

One qualifier is possible:

type This qualifier can have one of the following values:

- S.scname.period Service class period
- S.scname Service class
- R.rcname.period Report class period
- R.rcname Report class
- W.wname Workload
- POLICY Policy

period This qualifier can have one of the following values:

- S.scname.period Service class period
- R.rcname.period Report class period

Table 29. Workload Activity (Goal Mode) - Conditions Based on SMF Record Type 72-3

Condition	Condition Name	Qualifier	Source	Algorithm
Total service per second	TOTSRV	type	R723CSRV Interval	Sum(R723CSRV) / Interval
I/O service per second	IOSRV	type	R723CIOC Interval	Sum(R723CIOC) / Interval
CPU service per second	CPUSRV	type	R723CCPU Interval	Sum(R723CCPU) / Interval
SRB service per second	SRBSRV	type	R723CSRB Interval	Sum(R723CSRB) / Interval
Storage service per second	MSOSRV	type	R723CMSO Interval	Sum(R723CMSO) / Interval
Ended transactions per second	TRANS	type	R723CRCP Interval	Sum(R723CRCP) / Interval
Transaction execution time	RTIME	type	R723CXET R723CRCP	Sum(R723CXET) / Sum(R723CRCP)
Number swaps per transaction	SPERTRA	type	R723CSWC R723CRCP	Sum(R723CSWC) / Sum(R723CRCP)
Absorption rate	ABSRPTN	type	R723CSRV R723CTRR	Sum(R723CSRV) / Sum(R723CTRR)
Transaction service rate	TRXSERV	type	R723CSRV R723CTAT	Sum(R723CSRV) / Sum(R723CTAT)
Execution velocity	EXVEL	period	R723CTOU R723CTOT	Sum(R723CTOU) / (Sum(R723CTOU) + Sum(R723CTOT)) * 100
TCB seconds	TCBSEC	type	R723CCPU R723MCPJ R723MADJ	Sum((R723CCPU * R723MADJ) / (1600 * R723MCPJ))

Table 29. Workload Activity (Goal Mode) - Conditions Based on SMF Record Type 72-3 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
SRB seconds	SRBSEC	type	R723CSRB R723MSRB R723MADJ	$\text{Sum}((R723CSRB * R723MADJ) / (1600 * R723MSRB))$
Region Control Task (RCT) seconds	RCTSEC	type	R723CRCT	$\text{Sum}(R723CRCT)$
I/O interrupt (IIT) seconds	IITSEC	type	R723CIIT	$\text{Sum}(R723CIIT)$
Hiperspace™ service (HST) seconds	HSTSEC	type	R723CHST	$\text{Sum}(R723CHST)$
Application execution time on general purpose processors in seconds	APPLSEC	type	R723CCPU R723CSRB R723CRCT R723CIIT R723CHST R723MCPU R723MSRB R723MADJ R723CIFA R723CSUP R723NFFI R723NFFS	$\text{Sum}((R723CCPU * R723MADJ) / (1600 * R723MCPU) + (R723CSRB * R723MADJ) / (1600 * R723MSRB) + (R723CRCT + R723CIIT + R723CHST) - ((R723CIFA * R723MADJ) / (1600 * R723MCPU) * R723NFFI / 256) - ((R723CSUP * R723MADJ) / (1600 * R723MCPU) * R723NFFS / 256))$
Percentage of processor time used by task and preemptable-class SRB work	TCBPER	type	R723CCPU R723CIFA R723CSUP R723NFFI R723NFFS R723MCF R723MCFI R723MCFS R723MCPU R723MADJ Interval	$\text{Sum}(((R723CCPU - R723CIFA * R723NFFI / 256 - R723CSUP * R723NFFS / 256) / (R723MCF / 1024) + R723CIFA / (R723MCFI / 1024) + R723CSUP / (R723MCFS / 1024)) * R723MADJ) / (1600 * R723MCPU)) / \text{Interval} * 100$
Percentage of processor time used by non-preemptable SRB work	SRBPER	type	R723CSRB R723MSRB R723MADJ R723MCF Interval	$\text{Sum}((R723CSRB * R723MADJ) / (1600 * R723MCPU)) * 100 / (\text{Interval} * (R723MCF / 1024))$
Percentage of general purpose processor time used	APPLPER	type	R723CCPU R723CSRB R723CRCT R723CIIT R723CHST R723MCPU R723MSRB R723MADJ R723CIFA R723CSUP R723NFFI R723NFFS R723MCF Interval	$\text{Sum}((R723CCPU * R723MADJ) / (1600 * R723MCPU) + (R723CSRB * R723MADJ) / (1600 * R723MSRB) + (R723CRCT + R723CIIT + R723CHST) - ((R723CIFA * R723MADJ) / (1600 * R723MCPU) * R723NFFI / 256) - ((R723CSUP * R723MADJ) / (1600 * R723MCPU) * R723NFFS / 256))) / (\text{Interval} * (R723MCF / 1024)) * 100$
Page-in rate from auxiliary storage	SINGLE	type	R723CPIR R723CTRR R723CIEA	$\text{Sum}(R723CPIR) / \text{Sum}(R723CTRR - R723CIEA)$
Block page-in rate from auxiliary storage	BLOCK	type	R723CBPI R723CTRR R723CIEA	$\text{Sum}(R723CBPI) / \text{Sum}(R723CTRR - R723CIEA)$
Page-in rate from expanded storage	EXPSNGL	type	R723CPIE R723CTRR R723CIEA	$\text{Sum}(R723CPIE) / \text{Sum}(R723CTRR - R723CIEA)$
Block page-in rate from expanded storage	EXPBLK	type	R723CBPE R723CTRR R723CIEA	$\text{Sum}(R723CBPE) / \text{Sum}(R723CTRR - R723CIEA)$
Hiperspace page-in rate	HSP	type	R723CHPI R723CTRR R723CIEA	$\text{Sum}(R723CHPI) / \text{Sum}(R723CTRR - R723CIEA)$
ESO-hiperspace read miss rate	HSPMISS	type	R723CCRM R723CTRR R723CIEA	$\text{Sum}(R723CCRM) / \text{Sum}(R723CTRR - R723CIEA)$

Table 29. Workload Activity (Goal Mode) - Conditions Based on SMF Record Type 72-3 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Shared storage page-in rate from auxiliary storage	SHARED	type	R723CSPA R723CTRR R723CIEA	Sum(R723CSPA) / Sum(R723CTRR-R723CIEA)
Shared storage page-in rate from expanded storage	EXPSHR	type	R723CSPE R723CTRR R723CIEA	Sum(R723CSPE) / Sum(R723CTRR-R723CIEA)
Number of EXCPs	EXCP	type	R723CIOC R723MIOC	Sum(R723CIOC) / Sum(R723MIOC)
EXCP rate	EXCPRT	type	R723CIOC R723MIOC Interval	Sum(R723CIOC) / Sum(R723MIOC) / Interval
CS frames of all swapped-in transactions	STOCEN	type	R723CPRS R723CERS Interval	(Sum(R723CPRS)-Sum(R723CERS)) / Interval
ES frames of all swapped-in transactions	STOEXP	type	R723CERS Interval	Sum(R723CERS) / Interval
Shared frames of all swapped-in transactions	STOSHR	type	R723CSRS Interval	Sum(R723CSRS) / Interval
Total frames of all swapped-in transactions	STOTOT	type	R723CPRS Interval	Sum(R723CPRS) / Interval
Ended transactions	TRANSTOT	type	R723CRCP Interval	Sum(R723CRCP)
Average number of swapped-in transactions	TRANSMPL	type	R723CTRR Interval	Sum(R723CTRR) / Interval
Average number of active transactions	TRANSAVG	type	R723CTAT Interval	Sum(R723CTAT) / Interval
Transaction response time	RTIMETOT	type	R723CTET R723CRCP	Sum(R723CTET) / Sum(R723CRCP)
Transaction queue time	RTIMEQUE	type	R723CQDT R723CRCP	Sum(R723CQDT) / Sum(R723CRCP)
Transaction ineligible queue time	TRANSIQT	type	R723CIQT R723CRCP	Sum(R723CIQT) / Sum(R723CRCP)
Transaction r/s affinity delay time	TRANSADT	type	R723CADT R723CRCP	Sum(R723CADT) / Sum(R723CRCP)
Transaction JCL conversion time	TRANSCVT	type	R723CCVT R723CRCP	Sum(R723CCVT) / Sum(R723CRCP)
Start subchannel rate	SSCHRT	type	R723CIRC Interval	Sum(R723CIRC) / Interval
Average DASD response time	RESP	type	R723CICT R723CIWT R723CIDT R723CIOT R723CIRC	Sum(R723CIRC + R723CIWT + R723CIDT + R723CIOT) / Sum(R723CIRC)
Average DASD connect time	CONN	type	R723CICT R723CIRC	Sum(R723CICT) / Sum(R723CIRC)
Average DASD disconnect time	DISC	type	R723CIDT R723CIRC	Sum(R723CIDT) / Sum(R723CIRC)
Average DASD pending time	QPEND	type	R723CIWT R723CIRC	Sum(R723CIWT) / Sum(R723CIRC)
Average DASD IOS queue time	IOSQ	type	R723CIOT R723CIRC	Sum(R723CIOT) / Sum(R723CIRC)
Performance index	PI	period	R723CTOU R723CTOT R723CTET R723CRCP R723CVAL R723CPCT	Depending on goal definition See z/OS RMF Report Analysis for the calculation rules.
CPU Using %	CPUUSGP	period	R723CCUS R723CTSA	Sum(R723CCUS) / Sum(R723CTSA) * 100
CPU Delay %	CPUDLYP	period	R723CCDE R723CTSA	Sum(R723CCDE) / Sum(R723CTSA) * 100
Crypto Using %	CRYUSGP	period	R723CAMU R723APU R723CTSA	Sum(R723CAMU + R723APU) / Sum(R723CTSA + R723CAMU + R723CAMD + R723APU + R723APD)

Table 29. Workload Activity (Goal Mode) - Conditions Based on SMF Record Type 72-3 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Crypto CAP Using %	CAPUSGP	period	R723CAMU R723CTSA	$\text{Sum}(\text{R723CAMU}) / \text{Sum}(\text{R723CTSA} + \text{R723CAMU} + \text{R723CAMD} + \text{R723APU} + \text{R723APD})$
Crypto AP Using %	APUSGP	period	R723APU R723CTSA	$\text{Sum}(\text{R723APU}) / \text{Sum}(\text{R723CTSA} + \text{R723CAMU} + \text{R723CAMD} + \text{R723APU} + \text{R723APD})$
Crypto Delay %	CRYDLYP	period	R723CAMD R723APD R723FQD R723CTSA	$\text{Sum}(\text{R723CAMD} + \text{R723APD} + \text{R723FQD}) / \text{Sum}(\text{R723CTSA} + \text{R723CAMU} + \text{R723CAMD} + \text{R723APU} + \text{R723APD})$
Crypto CAP Delay %	CAPDLYP	period	R723CAMD R723CTSA	$\text{Sum}(\text{R723CAMD}) / \text{Sum}(\text{R723CTSA} + \text{R723CAMU} + \text{R723CAMD} + \text{R723APU} + \text{R723APD})$
Crypto AP Delay %	APDLYP	period	R723APD R723CTSA	$\text{Sum}(\text{R723APD}) / \text{Sum}(\text{R723CTSA} + \text{R723CAMU} + \text{R723CAMD} + \text{R723APU} + \text{R723APD})$
Crypto FQ Delay %	FQDLYP	period	R723FQD R723CTSA	$\text{Sum}(\text{R723FQD}) / \text{Sum}(\text{R723CTSA} + \text{R723CAMU} + \text{R723CAMD} + \text{R723APU} + \text{R723APD})$
Resource contention Using %	RCUSGP	period	R723RCOU	$\text{Sum}(\text{R723RCOU}) / \text{Sum}(\text{R723CTSA} + \text{R723RCOD} + \text{R723RCOU}) * 100$
Resource contention Delay %	RCDLYP	period	R723RCOD	$\text{Sum}(\text{R723RCOD}) / \text{Sum}(\text{R723CTSA} + \text{R723RCOD} + \text{R723RCOU}) * 100$
I/O Using %	IOUSGP	period	R723CIOU R723CTSA	$\text{Sum}(\text{R723CIOU}) / \text{Sum}(\text{R723CTSA}) * 100$
I/O Delay %	IODLYP	period	R723CIOD R723CTSA	$\text{Sum}(\text{R723CIOD}) / \text{Sum}(\text{R723CTSA}) * 100$
Swap-in delay %	SWINP	period	R723CSWI R723CTSA	$\text{Sum}(\text{R723CSWI}) / \text{Sum}(\text{R723CTSA}) * 100$
MPL delay %	MPLP	period	R723CMPL R723CTSA	$\text{Sum}(\text{R723CMPL}) / \text{Sum}(\text{R723CTSA}) * 100$
Queue %	QUEUEP	period	R723CQ R723CTSA	$\text{Sum}(\text{R723CQ}) / \text{Sum}(\text{R723CTSA}) * 100$
Capping %	CAPP	period	R723CCCA R723CTSA	$\text{Sum}(\text{R723CCCA}) / \text{Sum}(\text{R723CTSA}) * 100$
Storage %	STOP	period	R723CAPR R723CACO R723CAXM R723CVIO R723CHSP R723CCHS R723CTSA	$\text{Sum}(\text{R723CAPR} + \text{R723CACO} + \text{R723CAXM} + \text{R723CVIO} + \text{R723CHSP} + \text{R723CCHS}) / \text{Sum}(\text{R723CTSA}) * 100$
Server delay %	SERVP	period	R723CSPV R723CSVI R723CSHS R723CSMP R723CSSW R723CTSA	$\text{Sum}(\text{R723CSPV} + \text{R723CSVI} + \text{R723CSHS} + \text{R723CSMP} + \text{R723CSSW}) / \text{Sum}(\text{R723CTSA}) * 100$
Idle %	IDLEP	period	R723CIDL R723CTSA	$\text{Sum}(\text{R723CIDL}) / \text{Sum}(\text{R723CTSA}) * 100$
Unknown %	UNKP	period	R723CUNK R723CTSA	$\text{Sum}(\text{R723CUNK}) / \text{Sum}(\text{R723CTSA}) * 100$
Average number of independent enclaves during the interval (contained in TRANSAVG)	ENCAVG	type	R723CIEA	$\text{Sum}(\text{R723CIEA}) / \text{Interval}$
Average number of foreign enclaves during the interval	ENCREM	type	R723CFEA	$\text{Sum}(\text{R723CFEA}) / \text{Interval}$
Average number of multi-system enclaves during the interval	ENCMS	type	R723CXEA	$\text{Sum}(\text{R723CXEA}) / \text{Interval}$
zAAP service per second	AAPSRV	type	R723CIFA Interval	$\text{Sum}(\text{R723CIFA}) / \text{Interval}$
zAAP on CP service per second	AAPCPSRV	type	R723CIFC Interval	$\text{Sum}(\text{R723CIFC}) / \text{Interval}$

Table 29. Workload Activity (Goal Mode) - Conditions Based on SMF Record Type 72-3 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
zAAP service time in seconds	IFASEC, AAPSEC	type	R723CIFA R723MADJ R723MCPU	$\text{Sum}((\text{R723CIFA} * \text{R723MADJ}) / (1600 * \text{R723MCPU}))$
zAAP service time in seconds (normalized)	IFANSEC, AAPNSEC	type	R723CIFA R723MADJ R723MCPU R723NFFI	$\text{Sum}((\text{R723CIFA} * \text{R723MADJ}) / (1600 * \text{R723MCPU})) * \text{R723NFFI} / 256$
zAAP on CP service time in seconds	IFACPSEC, AAPCPSEC	type	R723CIFC R723MADJ R723MCPU	$\text{Sum}((\text{R723CIFC} * \text{R723MADJ}) / (1600 * \text{R723MCPU}))$
Percentage of zAAP time used	APPLIFA, APPLAAP	type	R723CIFA R723MADJ R723MCPU R723MCFI Interval	$\text{Sum}((\text{R723CIFA} * \text{R723MADJ}) / (1600 * \text{R723MCPU})) / (\text{Interval} * (\text{R723MCFI} / 1024)) * 100$
Percentage of general purpose processor time used by zAAP eligible transactions	APPLIFCP, APPLAPCP	type	R723CIFC R723MADJ R723MCPU R723MCF Interval	$\text{Sum}((\text{R723CIFC} * \text{R723MADJ}) / (1600 * \text{R723MCPU})) / (\text{Interval} * (\text{R723MCF} / 1024)) * 100$
zAAP Using %	IFAUSGP, AAPUSGP	period	R723IFAU R723CTSA	$\text{R723IFAU} / \text{R723CTSA} * 100$
zAAP on CP Using %	IFCUSGP, APCUSGP	period	R723IFCU R723CTSA	$\text{R723IFCU} / \text{R723CTSA} * 100$
zAAP Delay %	IFADLYP, AAPDLYP	period	R723IFAD R723CTSA	$\text{R723IFAD} / \text{R723CTSA} * 100$
zIIP service per second	IIPSRV	type	R723CSUP Interval	$\text{Sum}(\text{R723CSUP}) / \text{Interval}$
zIIP on CP service per second	IIPCPSRV	type	R723CSUC Interval	$\text{Sum}(\text{R723CSUC}) / \text{Interval}$
zIIP service time in seconds	IIPSEC	type	R723CSUP R723MADJ R723MCPU	$\text{Sum}((\text{R723CSUP} * \text{R723MADJ}) / (1600 * \text{R723MCPU}))$
zIIP service time in seconds (normalized)	IIPNSEC	type	R723CSUP R723MADJ R723MCPU R723NFFS	$\text{Sum}((\text{R723CSUP} * \text{R723MADJ}) / (1600 * \text{R723MCPU})) * \text{R723NFFS} / 256$
zIIP on CP service time in seconds	IIPCPSEC	type	R723CSUC R723MADJ R723MCPU	$\text{Sum}((\text{R723CSUC} * \text{R723MADJ}) / (1600 * \text{R723MCPU}))$
Percentage of zIIP time used	APPLIIP	type	R723CSUP R723MADJ R723MCPU R723MCFI Interval	$\text{Sum}((\text{R723CSUP} * \text{R723MADJ}) / (1600 * \text{R723MCPU})) / (\text{Interval} * (\text{R723MCFI} / 1024)) * 100$
Percentage of general purpose processor time used by zIIP eligible transactions	APPLIPCP	type	R723CSUC R723MADJ R723MCPU R723MCF Interval	$\text{Sum}((\text{R723CSUC} * \text{R723MADJ}) / (1600 * \text{R723MCPU})) / (\text{Interval} * (\text{R723MCF} / 1024)) * 100$
zIIP Using %	IIPUSGP	period	R723SUPU R723CTSA	$\text{R723SUPU} / \text{R723CTSA} * 100$
zIIP on CP Using %	IPCUSGP	period	R723SUCU R723CTSA	$\text{R723SUCU} / \text{R723CTSA} * 100$
zIIP Delay %	IIPDLYP	period	R723SUPD R723CTSA	$\text{R723SUPD} / \text{R723CTSA} * 100$
Percentage of general purpose processor time (including normalized zIIP and zAAP time) consumed while dispatching priority of work with low importance was temporarily raised to help blocked workloads	PROMPER	type	R723TPDP R723MCF Interval	$\text{Sum}(\text{R723TPDP}) / (\text{Interval} * (\text{R723MCF} / 1024)) * 100$
CPU time (in seconds) consumed while dispatching priority of work with low importance was temporarily raised to help blocked workloads	PROMSEC	type	R723TPDP	$\text{Sum}(\text{R723TPDP})$

Table 29. Workload Activity (Goal Mode) - Conditions Based on SMF Record Type 72-3 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Percentage of general purpose processor time (including normalized zIIP and zAAP time) consumed while dispatching priority was temporarily raised by enqueue management because the work unit held a resource that other work needed	EPROMPER	type	R723ECTC R723MCF Interval	$\text{Sum}(\text{R723ECTC}) / (\text{Interval} * (\text{R723MCF} / 1024)) * 100$
CPU time (in seconds) consumed while dispatching priority was temporarily raised by enqueue management because the work unit held a resource that other work needed	EPROMSEC	type	R723ECTC	$\text{Sum}(\text{R723ECTC})$
Percentage of general purpose processor time (including normalized zIIP and zAAP time) consumed while dispatching priority was temporarily raised by chronic resource contention management because the work unit held a resource that other work needed	CPROMPER	type	R723CPDP R723MCF Interval	$\text{Sum}(\text{R723CPDP}) / (\text{Interval} * (\text{R723MCF} / 1024)) * 100$
CPU time (in seconds) consumed while dispatching priority was temporarily raised by chronic resource contention management because the work unit held a resource that other work needed	CPROMSEC	type	R723CPDP	$\text{Sum}(\text{R723CPDP})$
Percentage of general purpose processor time (including normalized zIIP and zAAP time) consumed while dispatching priority was temporarily raised to shorten the lock hold time for a local suspend lock held by the work unit. Only valid in HiperDispatch mode.	LPROMPER	type	R723LPDP R723MCF Interval	$\text{Sum}(\text{R723LPDP}) / (\text{Interval} * (\text{R723MCF} / 1024)) * 100$
CPU time (in seconds) consumed while dispatching priority was temporarily raised to shorten the lock hold time for a local suspend lock held by the work unit. Only valid in HiperDispatch mode.	LPROMSEC	type	R723LPDP	$\text{Sum}(\text{R723LPDP})$
Percentage of general purpose processor time (including normalized zIIP and zAAP time) consumed while dispatching priority for a work unit was temporarily raised by the z/OS supervisor to a higher dispatching priority than assigned by WLM.	SPROMPER	type	R723SPDP R723MCF Interval	$\text{Sum}(\text{R723SPDP}) / (\text{Interval} * (\text{R723MCF} / 1024)) * 100$
CPU time (in seconds) consumed while dispatching priority for a work unit was temporarily raised by the z/OS supervisor to a higher dispatching priority than assigned by WLM.	SPROMSEC	type	R723SPDP	$\text{Sum}(\text{R723SPDP})$

Channel Path Activity - SMF record type 73

One qualifier is possible:

cpid A two-digit hexadecimal number that identifies a channel path.

cptype

Channel path type (as contained in SMF73ACR) enclosed in quotes.

For overview processing, one qualifier is required. If it is omitted for exception reporting, the threshold applies to all channel paths in the SMF record.

For most conditions, there exist two condition names. The second in the list has always the prefix CHG (instead of CH for the first one — to be used with the qualifier **cpid**) and has to be used with the qualifier **cptype**. This condition can be used for channels that are under control of Dynamic Channel Path Management

(DCM). All channels of the specified type will be accumulated and then processed. Therefore, the formulas for the algorithm contains summary values instead of single-channel values, for example:

$(SMF73TUT*100)/SMF73PTI*8$ see CHTBSY(cpid)
 $(Sum(SMF73TUT)*100)/Sum(SMF73PTI*8)$ see CHGTBSY(cptype)

Table 30. Channel Path Activity - Conditions Based on SMF Record Type 73

Condition	Condition Name	Qualifier	Source	Algorithm (depending on channel type)
For processors earlier than z990, use the following condition if CPMF is not available or for CPMF compatibility mode:				
Percent channel busy	CHPBSY CHGPBSY	cpid cptype	SMF73BSY SMF73SMP	$(BSY*100)/SMP$ Note: In case of monitoring z990 hardware, both fields are zero.
Use the following conditions for CPMF extended mode:				
Channel total busy %	CHTBSY CHGTBSY	cpid cptype	SMF73TUT SMF73PTI or SMF73TUC SMF73MCU SMF73PTI	$(TUT*100) / PTI*8$ or $(TUC*100) / (MCU*Int/10^6)$ Int = $PTI*1024$
Channel partition busy %	CHLBSY CHGLBSY	cpid cptype	SMF73PUT SMF73PTI or SMF73PUC SMF73MCU SMF73PTI	$(PUT*100) / PTI*8$ or $(PUC*100) / (MCU*Int/10^6)$
Channel bus total %	CHBTOT CHGBTOT	cpid cptype	SMF73TBC SMF73MBC SMF73PTI	$(TBC*100) / (MBC*Int/10^6)$
Channel total read rate (MB/SEC)	CHTREAD CHGTREAD	cpid cptype	SMF73TRU SMF73US SMF73PTI	$(TRU*US) / Int$
Channel partition read rate (MB/SEC)	CHLREAD CHGLREAD	cpid cptype	SMF73PRU SMF73US SMF73PTI	$(PRU*US) / Int$
Channel total write rate (MB/SEC)	CHTWRITE CHGTWRIT	cpid cptype	SMF73TWU SMF73US SMF73PTI	$(TWU*US) / Int$
Channel partition write rate (MB/SEC)	CHLWRITE CHGLWRIT	cpid cptype	SMF73PWU SMF73US SMF73PTI	$(PWU*US) / Int$
Channel partition write rate for HiperSockets (B/SEC)	CHLWRITE	cpid	SMF73PDS, SMF73PDU, SMF73PTI	$PDS*PDU / (Int/10^6)$
Channel total write rate for HiperSockets (B/SEC)	CHTWRITE	cpid	SMF73TDS, SMF73TDU, SMF73PTI	$TDS*TDU / (Int/10^6)$
Number of native FICON operations per second.	CHFRATE	cpid	SMF73EOC	EOC / Int
Average number of native FICON operations that are concurrently active.	CHFACTV	cpid	SMF73EOS SMF73EOC	EOS / EOC
Number of deferred native FICON operations per second that could not be initiated by the channel due to the lack of available resources.	CHFDPER	cpid	SMF73EOD	EOD / Int
Number of zHPF (High Performance FICON) operations per second.	CHFXRATE	cpid	SMF73ETC	ETC / Int
Average number of zHPF operations that are concurrently active.	CHFXACTV	cpid	SMF73ETS SMF73ETC	ETS / ETC
Number of deferred zHPF operations per second that could not be initiated by the channel due to the lack of available resources.	CHFXDFER	cpid	SMF73ETD	ETD / Int

Table 30. Channel Path Activity - Conditions Based on SMF Record Type 73 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm (depending on channel type)
Channel partition message sent rate	CHLMSGST	cpid	SMF73PMS, SMF73PUM, SMF73PTI	PMS*PUM / (Int/10 ⁶)
Channel total message sent rate	CHTMSGST	cpid	SMF73TMS, SMF73TUM, SMF73PTI	TMS*TUM / (Int/10 ⁶)
Average Channel partition message size (in bytes)	CHLMSGSZ	cpid	SMF73PDS, SMF73PDU, SMF73PMS, SMF73PUM	PDS*PDU/ PMS*PUM
Average Channel Total message size (in bytes)	CHTMSGSZ	cpid	SMF73TDS, SMF73TDU, SMF73TMS, SMF73TUM	TDS*TDU/ TMS*TUM
Channel partition message failed rate	CHLMSGF	cpid	SMF73PUS, SMF73PTI	PUS/(Int/10 ⁶)
Channel partition receive failed rate	CHLRECF	cpid	SMF73PUB, SMF73PTI	PUB/(Int/10 ⁶)
Channel total receive failed rate	CHTRECF	cpid	SMF73TUB, SMF73PTI	TUB/(Int/10 ⁶)

Device Activity - SMF record type 74-1

One qualifier is required, otherwise a syntax error occurs and RMF will not process the condition.

devnmbr

A one- to four-digit hexadecimal device number in the range 0000 through FFFF. Example: (012F)

volser A one- to six-character volume serial number enclosed in quotes. Example: ('012345')

stg grp

A one- to eight-character storage group name in parentheses, preceded by the keyword SG. Example: (SG(COMMON01))

class Any of the six valid device classes for Monitor I device activity measurements.

For OVW statements, only the qualifiers devnmbr and volser are valid. If you have selected a shared device in the sysplex, you will receive a value which reflects the sysplex view (not possible for DNOTRDY, DHPAVNM, DHPAVLSS, and DVCAP).

All times reported are in milliseconds, unless otherwise noted.

Table 31. Device Activity - Conditions Based on SMF Record Type 74-1

Condition	Condition Name	Qualifier	Source	Algorithm
Percent not ready	DNOTRDY	devnmbr, volser, stg grp, or class	SMF74NRD SMF74SAM	(NRD*100)/SAM (no sysplex view)
Percent reserved	DR	devnmbr, volser, stg grp, or class	SMF74RSV SMF74SAM	(RSV*100)/SAM
Percent mount pending	DMTPEND	devnmbr, volser, stg grp, or class	SMF74MTP SMF74SAM	(MTP*100)/SAM

Table 31. Device Activity - Conditions Based on SMF Record Type 74-1 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Percent device utilization	DVUTL	devnmbr, volser, stg grp, or class	SMF74CNN SMF74DIS SMF74INT SMF74UTL SMF74SAM	$((((CNN+DIS)/INT) + (UTL/SAM)) * 100$
Device activity rate	DART	devnmbr, volser, stg grp, or class	SMF74SSC SMF74INT	SSC/INT in seconds
Average connect time	DCTAVG	devnmbr, volser, stg grp, or class	SMF74CNN SMF74MEC	CNN/MEC
Average disconnect time	DDTAVG	devnmbr, volser, stg grp, or class	SMF74DIS SMF74MEC	DIS/MEC
Average pending time	DPTAVG	devnmbr, volser, stg grp, or class	SMF74PEN SMF74MEC	PEN/MEC
Average IOS queue time	DQTAVG	devnmbr, volser, stg grp, or class	SMF74QUE SMF74SAM SMF74SSC SMF74INT	$(QUE/SAM)/(SSC/INT)$
Average response time	DRTAVG	devnmbr, volser, stg grp, or class	SMF74ATV SMF74MEC SMF74SSC SMF74INT SMF74QUE SMF74SAM	$(ATV/MEC)+(QUE/SAM) / (SSC/INT)$
Average device busy delay time	DBDL	devnmbr, volser, stg grp, or class	SMF74DVB SMF74MEC	DVB/MEC
Average command response time	CMRDL	devnmbr	SMF74CMR SMF74MEC	CMR/MEC
Average number of HyperPAV devices	DHPAVNM	devnmbr, volser, stg grp, or class	SMF74NUX SMF74PSM	NUX/PSM
Number of configured HyperPAV aliases for the LSS of the device	DHPAVLSS	devnmbr, volser, stg grp, or class	SMF74HPC	HPC
DASD volume capacity (in cylinders)	DVCAP	devnmbr, volser	SMF74CAP	CAP
Average interrupt delay time	INTDL	devnmbr, volser, stg grp, or class	SMF74IDT SMF74MEC	IDT/MEC

Coupling Facility Activity - SMF record type 74-4

Due to the structure of the Coupling Facility Activity report, the scope of the results of overview processing is different and is indicated in column **Scope**:

- S** Overview column created only for each single system, not for sysplex
- X** Overview column created only for sysplex, not for each single system
- B** Overview column created for single systems as well as sysplex

There is no exception reporting for Coupling Facility records.

One qualifier is possible:

struct Mandatory - Coupling Facility structure name.

cfname

Mandatory - Coupling Facility name.

Table 32. Coupling Facility Activity - Conditions Based on SMF Record Type 74-4

Condition	Condition Name	Qualifier	Source	Algorithm	Scope
Average service time of SYNC operations	SYNCST	struct	R744SSTM R744SSRC	Sum(R744SSTM) / Sum(R744SSRC)	B
SYNC operation rate	SYNCR	struct	R744SSRC Interval	Sum(R744SSRC) / Interval	B
Average service time of ASYNC operations	ASYNCS	struct	R744SATM R744SARC	Sum(R744SATM) / Sum(R744SARC)	B
Ended ASYNC operation rate	ASYNCR	struct	R744SARC Interval	Sum(R744SARC) / Interval	B
Percentage of changed operations	CHNGDP	struct	R744SSRC R744SARC R744SSTA	Sum(R744SSTA) / (Sum(R744SSRC) + Sum(R744SARC)) * 100	B
Changed operation rate	CHNGDRT	struct	R744SSTA Interval	Sum(R744SSTA) / Interval	B
Path busy rate	PBSY	cfname	R744FPBC Interval	Sum(R744FPBC) / Interval	S
Percent delayed requests	DREQP	cfname	R744SSRC R744SARC R744SSTA R744FSCC R744SQR	(Sum(R744FSCC) + Sum(R744SQR)) / (Sum(R744SSRC + R744SARC + R744SSTA)) * 100	S
CF processor utilization	CFUTIL	cfname	R744PBSY R744PWAI	Sum(R744PBSY) / (Sum(R744PBSY) + Sum(R744PWAI)) * 100 Summation over all processors ← unweighted average	X
Directory reclaims	DIRRCLM	struct	R744CDER	Sum(R744CDER)	X
List/directory entry current to total ratio	LDECTR	struct	R744SLEL R744SLEM R744SDEC R744CDEC	R744SLEM / R744SLEL (for List/Lock structure) R744CDEC / R744SDEC (for Cache structure)	X
Data elements current to total ratio	DECTR	struct	R744SMAE R744SCUE R744SDEL R744CDAC	R744SCUE / R744SMAE (for List/Lock structure) R744CDAC / R744SDEL (for Cache structure)	X
Lock entries current to total ratio	LECTR	struct	R744SLTL R744SLTM	(R744SLTM / R744SLTL)	X
Cache read request rate	CREADRT	struct	R744CRHC Interval	R744CRHC / Interval	X
Cache write request rate	CWRTERT	struct	R744CWH0 R744CWH1 Interval	(R744CWH0 + R744CWH1) / Interval	X
Cache castout rate	CCOUTRT	struct	R744CCOC Interval	R744CCOC / Interval	X
Cache cross invalidation rate	CXIRT	struct	R744CXDR R744CXFW R744CXNI R744CXRL R744CXCI Interval	(R744CXDR + R744CXFW + R744CXNI + R744CXRL + R744CXCI) / Interval	X
Total requests to lock structure or serialized list structure	LCKREQ	struct	R744STRC	Sum(R744STRC)	B
Contention on lock structure	LCKCONT	struct	R744SCN	Sum(R744SCN)	B
False contention on lock structure	LCKFCNT	struct	R744SFCN	Sum(R744SFCN)	B
CF utilization percentage	STUTILP	struct	R744SETM R744PBSY	R744SETM * 100 / Sum(R744PBSY)	X

Table 32. Coupling Facility Activity - Conditions Based on SMF Record Type 74-4 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm	Scope
Subchannel busy percentage	SUBCHBP	cfname	R744SSTM R744SATM R744FSCU	$(\text{SUM}(\text{R744SSTM}) + \text{SUM}(\text{R744SATM})) * 100 / \text{Interval} * \text{R744FSCU}$	S
Percentage of storage class memory in use	SCMIUP	struct	R744MIUS R744MSMA	$\text{R744MIUS} * 100 / \text{R744MSMA}$	X
Percentage of augmented space in use	AUGMIUP	struct	R744MIUA R744MEMA	$\text{R744MIUA} * 100 / \text{R744MEMA}$	X
SCM list entry current to total ratio	SCMLCTR	struct	R744MENL R744MEML	$\text{R744MENL} / \text{R744MEML}$	X
SCM list element current to total ratio	SCMLECTR	struct	R744MENE R744MEME	$\text{R744MENE} / \text{R744MEME}$	X
Average service time per SCM read operation	SCMRST	struct	R744MRST R744MRFC R744MRPC	$\text{R744MRST} / (\text{R744MRFC} + \text{R744MRPC})$	X
Average service time per SCM write operation	SCMWST	struct	R744MWST R744MSWC	$\text{R744MWST} / \text{R744MSWC}$	X
SCM auxiliary enabled commands to total request ratio	SCMAUXR	struct	R744MAEC R744SSRC R744SARC	$\text{R744MAEC} / \text{Sum}(\text{R744SSRC} + \text{R744SARC})$	X
SCM delayed faults to total request ratio	SCMDFR	struct	R744MSRL R744MSRR R744MSRM R744SSRC R744SARC	$(\text{R744MSRL} + \text{R744MSRR} + \text{R744MSRM}) / \text{Sum}(\text{R744SSRC} + \text{R744SARC})$	X

Cache Activity - SMF record type 74-5

The following qualifiers are possible:

ssid SSID number

devn Device number

rrid RAID rank identifier

To define a subsystem-related exception, you specify SSID(ssid)

To define a device-related exception, you specify SSID(ssid),DEVN(devn)

To define an exception for RAID rank data, you specify SSID(ssid),RRID(rrid)

For exception processing only:

The conditions CASSC, CADSC, and CASSNVS can be used only with the operator EQ, not with LE or GE.

Table 33. Cache Activity - Conditions Based on SMF Record Type 74-5

Condition	Condition Name	Qualifier	Source	Algorithm
Subsystem Status: Device Status: CACHING	CASSC CADSC	SSID(ssid) DEVN(devn)	R745SOS R745DSDV	active = R745SOS all 3 bits zero active = R745DSDV, both bits zero
Subsystem Status: NON-VOLATILE STORAGE	CASSNVS	SSID(ssid)	R745SVSS	active = Bit 0 to 4 zero
Subsystem Overview: Device Activity: TOTAL I/O	CASTOT CADTOT	SSID(ssid) DEVN(devn)	R745DRCR R745DRSR R745DRNR R745DWRC R745DWSR R745DWNR R745DICL R745DBCR	Sum of these counts

Table 33. Cache Activity - Conditions Based on SMF Record Type 74-5 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Subsystem Overview: Device Activity: CACHE I/O	CASCTOT CADCTOT	SSID(ssid) DEVN(devn)	R745DRCR R745DRSR R745DRNR R745DWRC R745DWSR R745DWNR	Sum of these counts
Subsystem Overview: CACHE OFFLINE	CASCOFF	SSID(ssid)	R745DRCR R745DRSR R745DRNR R745DWRC R745DWSR R745DWNR R745DICL R745DBCR	Sum of these counts
Subsystem Overview: Device Activity: TOTAL H/R	CASHRT CADHRT	SSID(ssid) DEVN(devn)	HITS = R745DCRH + R745DRSH + R745DNRH + R745DWCH + R745DWSH + R745DWNH TOTAL = R745DRCR + R745DRSR + R745DRNR + R745DWRC + R745DWSR + R745DWNR + R745DICL + R745DBCR	HITS / TOTAL
Subsystem Overview: Device Activity: CACHE H/R	CASHR CADHR	SSID(ssid) DEVN(devn)	HITS = R745DCRH + R745DRSH + R745DNRH + R745DWCH + R745DWSH + R745DWNH CACHE I/O = R745DRCR + R745DRSR + R745DRNR + R745DWRC + R745DWSR + R745DWNR	HITS / CACHE I/O
Subsystem Overview: Device Activity: READ I/O REQUESTS RATE NORMAL	CASRN CADRN	SSID(ssid) DEVN(devn)	R745DRCR R745CINT	R745DRCR / R745CINT
Subsystem Overview: Device Activity: READ I/O REQUESTS RATE SEQUENTIAL	CASRS CADRS	SSID(ssid) DEVN(devn)	R745DRSR R745CINT	R745DRSR / R745CINT
Subsystem Overview: Device Activity: READ I/O REQUESTS RATE CFW DATA	CASRC CADRC	SSID(ssid) DEVN(devn)	R745DRNR R745CINT	R745DRNR / R745CINT
Subsystem Overview: Device Activity: READ I/O REQUESTS RATE TOTAL	CASRT CADRT	SSID(ssid) DEVN(devn)	R745CINT TOTAL = R745DRCR + R745DRSR + R745DRNR	TOTAL / R745CINT
Subsystem Overview: Device Activity: READ I/O REQUESTS HITS RATE NORMAL	CASRHN CADRHN	SSID(ssid) DEVN(devn)	R745DCRH R745CINT	R745DCRH / R745CINT
Subsystem Overview: Device Activity: READ I/O REQUESTS HITS RATE SEQUENTIAL	CASRHS CADRHS	SSID(ssid) DEVN(devn)	R745DRSH R745CINT	R745DRSH / R745CINT
Subsystem Overview: Device Activity: READ I/O REQUESTS HITS RATE CFW DATA	CASRHC CADRHC	SSID(ssid) DEVN(devn)	R745DNRH R745CINT	R745DNRH / R745CINT
Subsystem Overview: Device Activity: READ I/O REQUESTS HITS RATE TOTAL	CASRHT CADRHT	SSID(ssid) DEVN(devn)	R745CINT TOTAL = R745DCRH + R745DRSH + R745DNRH	TOTAL / R745CINT
Subsystem Overview: Device Activity: READ I/O REQUESTS H/R NORMAL	CASRHRN CADRHRN	SSID(ssid) DEVN(devn)	R745DCRH R745DRCR	R745DCRH / R745DRCR

Table 33. Cache Activity - Conditions Based on SMF Record Type 74-5 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Subsystem Overview: Device Activity: READ I/O REQUESTS H/R SEQUENTIAL	CASRHRS CADRHRS	SSID(ssid) DEVN(devn)	R745DRSH R745DRSR	R745DRSH / R745DRSR
Subsystem Overview: Device Activity: READ I/O REQUESTS H/R CFW DATA	CASRHRC CADRHRC	SSID(ssid) DEVN(devn)	R745DNRH R745DRNR	R745DNRH / R745DRNR
Subsystem Overview: Device Activity: READ I/O REQUESTS H/R TOTAL	CASRHRT CADRHRT	SSID(ssid) DEVN(devn)	HITS = R745DCRH + R745DRSH + R745DNRH TOTAL = R745DRCR + R745DRSR + R745DRNR	HITS / TOTAL
Subsystem Overview: Device Activity: WRITE I/O REQUESTS RATE NORMAL	CASWN CADWN	SSID(ssid) DEVN(devn)	R745DWRC R745CINT	R745DWRC / R745CINT
Subsystem Overview: Device Activity: WRITE I/O REQUESTS RATE SEQUENTIAL	CASWS CADWS	SSID(ssid) DEVN(devn)	R745DWSR R745CINT	R745DWSR / R745CINT
Subsystem Overview: Device Activity: WRITE I/O REQUESTS RATE CFW DATA	CASWC CADWC	SSID(ssid) DEVN(devn)	R745DWNR R745CINT	R745DWNR / R745CINT
Subsystem Overview: Device Activity: WRITE I/O REQUESTS RATE TOTAL	CASWT CADWT	SSID(ssid) DEVN(devn)	R745CINT TOTAL = R745DWRC + R745DWSR + R745DWNR	TOTAL / R745CINT
Subsystem Overview: Device Activity: WRITE I/O REQUESTS FAST WRITE RATE NORMAL	CASWFN CADWFN	SSID(ssid) DEVN(devn)	R745DFWC R745CINT	R745DFWC / R745CINT
Subsystem Overview: Device Activity: WRITE I/O REQUESTS FAST WRITE RATE SEQUENTIAL	CASWFS CADWFS	SSID(ssid) DEVN(devn)	R745DFWS R745CINT	R745DFWS / R745CINT
Subsystem Overview: Device Activity: WRITE I/O REQUESTS FAST WRITE RATE CFW DATA	CASWFC CADWFC	SSID(ssid) DEVN(devn)	R745DWNR R745CINT	R745DWNR / R745CINT
Subsystem Overview: Device Activity: WRITE I/O REQUESTS FAST WRITE RATE TOTAL	CASWFT CADWFT	SSID(ssid) DEVN(devn)	R745CINT TOTAL = R745DFWC + R745DFWS + R745DWNR	TOTAL / R745CINT
Subsystem Overview: Device Activity: WRITE I/O REQUESTS HITS RATE NORMAL	CASWHN CADWHN	SSID(ssid) DEVN(devn)	R745DWCH R745CINT	R745DWCH / R745CINT
Subsystem Overview: Device Activity: WRITE I/O REQUESTS HITS RATE SEQUENTIAL	CASWHS CADWHS	SSID(ssid) DEVN(devn)	R745DWSH R745CINT	R745DWSH / R745CINT
Subsystem Overview: Device Activity: WRITE I/O REQUESTS HITS RATE CFW DATA	CASWHC CADWHC	SSID(ssid) DEVN(devn)	R745DWNH R745CINT	R745DWNH / R745CINT
Subsystem Overview: Device Activity: WRITE I/O REQUESTS HITS RATE TOTAL	CASWHT CADWHT	SSID(ssid) DEVN(devn)	R745CINT TOTAL = R745DWCH + R745DWSH + R745DWNH	TOTAL / R745CINT
Subsystem Overview: Device Activity: WRITE I/O REQUESTS H/R NORMAL	CASWHRN CADWHRN	SSID(ssid) DEVN(devn)	R745DWCH R745DWRC	R745DWCH / R745DWRC
Subsystem Overview: Device Activity: WRITE I/O REQUESTS H/R SEQUENTIAL	CASWHRS CADWHRS	SSID(ssid) DEVN(devn)	R745DWSH R745DWSR	R745DWSH / R745DWSR
Subsystem Overview: Device Activity: WRITE I/O REQUESTS H/R CFW DATA	CASWHRC CADWHRC	SSID(ssid) DEVN(devn)	R745DWNH R745DWNR	R745DWNH / R745DWNR
Subsystem Overview: Device Activity: WRITE I/O REQUESTS H/R TOTAL	CASWHRT CADWHRT	SSID(ssid) DEVN(devn)	HITS = R745DWCH + R745DWSH + R745DWNH TOTAL = R745DWRC + R745DWSR + R745DWNR	HITS / TOTAL

Table 33. Cache Activity - Conditions Based on SMF Record Type 74-5 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Subsystem Overview: Device Activity: % READ NORMAL	CASRWN CADRWN	SSID(ssid) DEVN(devn)	R745DRCR R745DWRC	$R745DRCR * 100 / (R745DRCR + R745DWRC)$
Subsystem Overview: Device Activity: % READ SEQUENTIAL	CASRWS CADRWS	SSID(ssid) DEVN(devn)	R745DRSR R745DWSR	$R745DRSR * 100 / (R745DRSR + R745DWSR)$
Subsystem Overview: Device Activity: % READ CFW DATA	CASRWC CADRWC	SSID(ssid) DEVN(devn)	R745DRNR R745DWNR	$R745DRNR * 100 / (R745DRNR + R745DWNR)$
Subsystem Overview: Device Activity: % READ TOTAL	CASRWT CADRWT	SSID(ssid) DEVN(devn)	TOTAL_READ = R745DRCR + R745DRSR + R745DRNR TOTAL_WRITE = R745DWRC + R745DWSR + R745DWNR	$TOTAL_READ * 100 / (TOTAL_READ + TOTAL_WRITE)$
Subsystem Overview: Device Activity: CACHE MISSES READ RATE NORMAL	CASMRN CADMRN	SSID(ssid) DEVN(devn)	R745DRCR R745DCRH R745CINT	$(R745DRCR - R745DCRH) / R745CINT$
Subsystem Overview: Device Activity: CACHE MISSES READ RATE SEQUENTIAL	CASMRS CADMRS	SSID(ssid) DEVN(devn)	R745DRSR R745DRSH R745CINT	$(R745DRSR - R745DRSH) / R745CINT$
Subsystem Overview: Device Activity: CACHE MISSES READ RATE CFW DATA	CASMRC CADMRC	SSID(ssid) DEVN(devn)	R745DRNR R745DNRH R745CINT	$(R745DRNR - R745DNRH) / R745CINT$
Subsystem Overview: Device Activity: CACHE MISSES WRITE RATE NORMAL	CASMWN CADMWN	SSID(ssid) DEVN(devn)	R745DWRC R745DWCH R745CINT	$(R745DWRC - R745DWCH) / R745CINT$
Subsystem Overview: Device Activity: CACHE MISSES WRITE RATE SEQUENTIAL	CASMWS CADMWS	SSID(ssid) DEVN(devn)	R745DWSR R745DWSH R745CINT	$(R745DWSR - R745DWSH) / R745CINT$
Subsystem Overview: Device Activity: CACHE MISSES WRITE RATE CFW DATA	CASMWC CADMWC	SSID(ssid) DEVN(devn)	R745DWNR R745DWNH R745CINT	$(R745DWNR - R745DWNH) / R745CINT$
Subsystem Overview: Device Activity: CACHE MISSES TRACKS RATE NORMAL	CASMTN CADMTN	SSID(ssid) DEVN(devn)	R745DNTD R745CINT	$R745DNTD / R745CINT$
Subsystem Overview: Device Activity: CACHE MISSES TRACKS RATE SEQUENTIAL	CASMTS CADMTS	SSID(ssid) DEVN(devn)	R745DTC R745CINT	$R745DTC / R745CINT$
Subsystem Overview: Device Activity: CACHE MISSES RATE TOTAL	CASMT CADMT	SSID(ssid) DEVN(devn)	R745CINT TOTAL = R745DRCR + R745DRSR + R745DRNR + R745DWRC + R745DWSR + R745DWNR HITS = R745DCRH + R745DRSH + R745DNRH + R745DWCH + R745DWSH + R745DWNH	$(TOTAL - HITS) / R745CINT$
Subsystem Overview: Device Activity: MISC (Miscellaneous) DFW BYPASS RATE	CASDFWB CADDFWB	SSID(ssid) DEVN(devn)	R745DFWB R745CINT	$R745DFWB / R745CINT$
Subsystem Overview: Device Activity: MISC (Miscellaneous) CFW BYPASS RATE	CASCFWB CADCFWB	SSID(ssid) DEVN(devn)	R745DFWR R745CINT	$R745DFWR / R745CINT$
Subsystem Overview: Device Activity: MISC (Miscellaneous) DFW INHIBIT RATE	CASDFWI CADDFWI	SSID(ssid) DEVN(devn)	R745CINT TOTAL_ WRITES = R745DWRC + R745DWSR + R745DWNR FAST_ WRITES = R745DFWC + R745DFWS + R745DWNR	$(TOTAL_WRITES - FAST_WRITES) / R745CINT$

Table 33. Cache Activity - Conditions Based on SMF Record Type 74-5 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Subsystem Overview: Device Activity: MISC (Miscellaneous) ASYNC(TRKS) RATE	CASASYNC CADASYNC	SSID(ssid) DEVN(devn)	R745DCTD R745CINT	R745DCTD / R745CINT
Subsystem Overview: Device Activity: NON CACHE I/O ICL RATE	CASNCICL CADNCICL	SSID(ssid) DEVN(devn)	R745DICL R745CINT	R745DICL / R745CINT
Subsystem Overview: Device Activity: NON CACHE I/O BYPASS RATE	CASNCB CADNCB	SSID(ssid) DEVN(devn)	R745DBCR R745CINT	R745DBCR / R745CINT
Subsystem Overview: Device Activity: NON CACHE I/O TOTAL RATE	CASNCT CADNCT	SSID(ssid) DEVN(devn)	R745DICL R745DBCR R745CINT	(R745DICL + R745DBCR) / R745CINT
Subsystem Overview: Device Activity: Host Adapter Activity - Average number of bytes per read request (normal+sequential+CFW)	CASBRR CADBRR	SSID(ssid) DEVN(devn)	R7451CT1 R745DRCR R745DRSR R745DRNR	Sum(R7451CT1) / Sum(R745DRCR + R745DRSR + R745DRNR)
Subsystem Overview: Device Activity: Host Adapter Activity - Average number of bytes read per second	CASBRS CADBRS	SSID(ssid) DEVN(devn)	R7451CT1 R745CINT	Sum(R7451CT1) / R745CINT
Subsystem Overview: Device Activity: Host Adapter Activity - Average number of bytes per write request (normal+sequential+CFW)	CASBWR CADBWR	SSID(ssid) DEVN(devn)	R7451CT2 R745DRCR R745DRSR R745DRNR	Sum(R7451CT2) / Sum(R745DRCR + R745DRSR + R745DRNR)
Subsystem Overview: Device Activity: Host Adapter Activity - Average number of bytes written per second	CASBWS CADBWS	SSID(ssid) DEVN(devn)	R7451CT2 R745CINT	Sum(R7451CT2) / R745CINT
Subsystem Overview: Device Activity: Disk Activity - Response time in milliseconds per read request	CASDRRT CADDRRT	SSID(ssid) DEVN(devn)	R7452PRT R7452PRO	Sum(R7452PRT) / Sum(R7452PRO)
Subsystem Overview: Device Activity: Disk Activity - Average number of bytes per read request	CASDRBR CADDRBR	SSID(ssid) DEVN(devn)	R7452PBR R7452PRO	Sum(R7452PBR) / Sum(R7452PRO)
Subsystem Overview: Device Activity: Disk Activity - Average number of bytes read per second	CASDRBS CADDRBS	SSID(ssid) DEVN(devn)	R7452PBR R745CINT	Sum(R7452PBR) / R745CINT
Subsystem Overview: Device Activity: Disk Activity - Response time in milliseconds per write request	CASDWRT CADDWRT	SSID(ssid) DEVN(devn)	R7452PWT R7452PWO	Sum(R7452PWT) / Sum(R7452PWO)
Subsystem Overview: Device Activity: Disk Activity - Average number of bytes per write request	CASDWBR CADDWBR	SSID(ssid) DEVN(devn)	R7452PBW R7452PWO	Sum(R7452PBW) / Sum(R7452PWO)
Subsystem Overview: Device Activity: Disk Activity - Average number of bytes written per second	CASDWBS CADDWBS	SSID(ssid) DEVN(devn)	R7452PBW R745CINT	Sum(R7452PBW) / R745CINT
Subsystem Overview: Device Activity: CFW % of all requests	CASCFWP CADCFWP	SSID(ssid) DEVN(devn)	R745DWNR R745DRNR R745DRCR R745DRSR R745DRNR R745DWRC R745DWSR R745DWNR	Sum(R745DWNR + R745DRNR) * 100 / Sum(R745DRCR + R745DRSR + R745DRNR + R745DWRC + R745DWSR + R745DWNR)
Device Overview *CACHE: Sequential Stage Rate	CASCSTGS	SSID(ssid)	R745DWSR R745DWSH R745DRSR R745DRSH R745DWSR R745CINT	Sum((R745DWSR - R745DWSH) + (R745DRSR - R745DRSH) - R745DWSR + R745DFWS) / R745CINT
Device Overview: I/O RATE	CADT	DEVN(devn)	R745CINT TOTAL = R745DRCR + R745DRSR + R745DRNR + R745DWRC + R745DWSR + R745DWNR + R745DICL + R745DBCR	TOTAL / R745CINT

Table 33. Cache Activity - Conditions Based on SMF Record Type 74-5 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Device Overview: DASD I/O RATE STAGE	CADSTG	DEVN(devn)	R745CINT TOTAL = R745DRCR + R745DRSR + R745DRNR + R745DWRC + R745DWSR + R745DWNH HITS = R745DCRH + R745DRSH + R745DNRH + R745DWCH + R745DWSH + R745DWNH DFW_INHIBIT = R745DWRC + R745DWSR + R745DWNH - R745DFWC - R745DFWS - R745DWNH CFW_BYPASS = R745DFWR DFW_BYPASS = R745DFWB	(TOTAL - HITS - DFW_INHIBIT - CFW_BYPASS - DFW_BYPASS) / R745CINT
Device Activity: Sequential Stage Rate	CADSTGS	SSID(ssid) DEVN(devn)	R745DWSR R745DWSH R745DRSR R745DRSH R745DWSR R745CINT	Sum((R745DWSR - R745DWSH) + (R745DRSR - R745DRSH) - R745DWSR + R745DFWS) / R745CINT
RAID rank read request rate	CARRRT	SSID(ssid) RRID(rrid)	R7451RRQ R745CINT	R7451RRQ / R745CINT
Average number of megabytes read with an I/O request	CARRMB	SSID(ssid) RRID(rrid)	R7451SR R7451HSS R7451RRQ	(R7451SR * R7451HSS) / R7451RRQ
Average number of megabytes read per second	CARRMBS	SSID(ssid) RRID(rrid)	R7451SR R7451HSS R745CINT	(R7451SR * R7451HSS) / R745CINT
Average response time of a read request	CARRRTIM	SSID(ssid) RRID(rrid)	R7451RRT R7451RRQ	R7451RRT / R7451RRQ
RAID rank write request rate	CARWRT	SSID(ssid) RRID(rrid)	R7451WRQ R745CINT	R7451WRQ / R745CINT
Average number of megabytes written with an I/O request	CARWMB	SSID(ssid) RRID(rrid)	R7451SW R7451HSS R7451WRQ	(R7451SW * R7451HSS) / R7451WRQ
Average number of megabytes written per second	CARWMBS	SSID(ssid) RRID(rrid)	R7451SW R7451HSS R745CINT	(R7451SW * R7451HSS) / R745CINT
Average response time of a write request	CARWRTIM	SSID(ssid) RRID(rrid)	R7451WRT R7451WRQ	R7451WRT / R7451WRQ
zHPF read I/O rate	CASZHPFR CADZHPFR	SSID(ssid) DEVN(devn)	R7451CT5 R745CINT	Sum(R7451CT5) / R745CINT R7451CT5 / R745CINT
zHPF write I/O rate	CASZHPFW CADZHPFW	SSID(ssid) DEVN(devn)	R7451CT6 R745CINT	Sum(R7451CT6) / R745CINT R7451CT6 / R745CINT
Global Mirror Collisions Sidefile count	CASGMCSF CADGMCSF	SSID(ssid) DEVN(devn)	R7451GSF	Sum(R7451GSF) R7451GSF
Global Mirror Collisions Send Synchronous count	CASGMCSS CADGMCSS	SSID(ssid) DEVN(devn)	R7451GSS	Sum(R7451GSS) R7451GSS

Table 34. Cache Activity - Conditions Based on SMF Record Type 74-5

Condition	Categorie	Condition Name	Qualifier	Source	Algorithm
Device Overview: % I/O (total)	*CACHE-OFF	CASCOIO	SSID(ssid)	TOTAL = R745DRCR + R745DRSR + R745DRNR + R745DWRC + R745DWSR + R745DWNH + R745DICL + R745DBCR COTOT = TOTAL for CACHE-OFF ALLTOT = TOTAL for ALL I/Os	COTOT * 100 / ALLTOT

Table 34. Cache Activity - Conditions Based on SMF Record Type 74-5 (continued)

Condition	Categorie	Condition Name	Qualifier	Source	Algorithm
Device Overview: I/O RATE (total)	*ALL *CACHE *CACHE-OFF	CASAT CASCT CASOT	SSID(ssid)	R745CINT TOTAL = R745DRCR + R745DRSR + R745DRNR + R745DWRC + R745DWSR + R745DWNR + R745DICL + R745DBCR	TOTAL / R745CINT
Device Overview: DASD I/O RATE STAGE	*ALL *CACHE	CASASTG CASCSTG	SSID(ssid)	R745CINT TOTAL = R745DRCR + R745DRSR + R745DRNR + R745DWRC + R745DWSR + R745DWNR HITS = R745DCRH + R745DRSH + R745DNRH + R745DWCH + R745DWSH + R745DWNH DFW_INHIBIT = R745DWRC + R745DWSR + R745DWNR - R745DFWC - R745DFWS - R745DWNR CFW_BYPASS = R745DFWR DFW_BYPASS = R745DFWB	(TOTAL - HITS - DFW_INHIBIT - CFW_BYPASS) / R745CINT
Device Overview: Sequential Stage Rate	*ALL *CACHE	CASASTGS CASCSTGS	SSID(ssid)	R745CINT R745DFWS TOTAL_SEQ = R745DWSR + R745DRSR HIT_SEQ = R745DWSH + R745DRSH	Sum(TOTAL_SEQ - HIT_SEQ - R745DWSR + R745DFWS) / R745CINT

Ficon Director Activity - SMF record type 74-7

Following qualifiers are required:

sdev Switch device number

port Port address

Table 35. Ficon Director Activity - Conditions Based on SMF Record 74-7

Condition	Condition Name	Qualifier	Source	Algorithm
Average frame pacing time	FDAFPT	SDEV(sdev) PORT(port)	R747PFPT R747PNFT	(PFPT * 2.5) / PNFT
Rate (in MBs/sec) of data received during the interval	FDMBREAD	SDEV(sdev) PORT(port)	R747PNWR	(PNWR / (250 * 1000)) / Interval
Rate (in MBs/sec) of data transmitted during the interval	FDMBWRT	SDEV(sdev) PORT(port)	R747PNWT	(PNWT / (250 * 1000)) / Interval
Number of errors	FDNERR	SDEV(sdev) PORT(port)	R747PNER	Value or comparison

Enterprise Disk Systems Statistics - SMF record type 74-8

Following qualifiers are possible:

sern serial number

said system adapter ID

xpid extent pool ID

rrid RAID rank ID

Table 36. ESS Link Statistics - Conditions Based on SMF Record 74-8

Condition	Condition Name	Qualifier	Source	Algorithm
Transfer Rate ECKD™ Read	ESTRERD	SERN(sern), SAID(said)	R748LERB R748CINT	R748LERB / R748CINT

Table 36. ESS Link Statistics - Conditions Based on SMF Record 74-8 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Transfer Rate SCSI Read	ESTRSRD	SERN(sern), SAID(said)	R748LSRB R748CINT	R748LSRB / R748CINT
Transfer Rate ECKD Write	ESTREWR	SERN(sern), SAID(said)	R748LEWB R748CINT	R748LEWB / R748CINT
Transfer Rate SCSI Write	ESTRSWR	SERN(sern), SAID(said)	R748LSWB R748CINT	R748LSWB / R748CINT
Packet Size ECKD Read	ESPSERD	SERN(sern), SAID(said)	R748LERB R748LERO	R748LERB / R748LERO
Packet Size SCSI Read	ESPSSRD	SERN(sern), SAID(said)	R748LSRB R748LSRO	R748LSRB / R748LSRO
Packet Size ECKD Write	ESPSEWR	SERN(sern), SAID(said)	R748LEWB R748LEWO	R748LEWB / R748LEWO
Packet Size SCSI Write	ESPSSWR	SERN(sern), SAID(said)	R748LSWB R748LSWO	R748LSWB / R748LSWO
Activity Rate ECKD Read	ESARERD	SERN(sern), SAID(said)	R748LERO R748CINT	R748LERO / R748CINT
Activity Rate SCSI Read	ESARSRD	SERN(sern), SAID(said)	R748LSRO R748CINT	R748LSRO / R748CINT
Activity Rate ECKD Write	ESAREWR	SERN(sern), SAID(said)	R748LEWO R748CINT	R748LEWO / R748CINT
Activity Rate SCSI Write	ESARSWR	SERN(sern), SAID(said)	R748LSWO R748CINT	R748LSWO / R748CINT
Response Time ECKD Read	ESRTERD	SERN(sern), SAID(said)	R748LERT R748LERO	R748LERT / R748LERO
Response Time SCSI Read	ESRTSRD	SERN(sern), SAID(said)	R748LSRT R748LSRO	R748LSRT / R748LSRO
Response Time ECKD Write	ESRTEWR	SERN(sern), SAID(said)	R748LEWT R748LEWO	R748LEWT / R748LEWO
Response Time SCSI Write	ESRTSWR	SERN(sern), SAID(said)	R748LSWT R748LSWO	R748LSWT / R748LSWO
I/O Intensity ECKD Read	ESIOIERD	SERN(sern), SAID(said)	R748LERT R748CINT	R748LERT / R748CINT
I/O Intensity SCSI Read	ESIOISRD	SERN(sern), SAID(said)	R748LSRT R748CINT	R748LSRT / R748CINT
I/O Intensity ECKD Write	ESIOIEWR	SERN(sern), SAID(said)	R748LEWT R748CINT	R748LEWT / R748CINT
I/O Intensity SCSI Write	ESIOISWR	SERN(sern), SAID(said)	R748LSWT R748CINT	R748LSWT / R748CINT
I/O Intensity ECKD Total	ESIOIET	SERN(sern), SAID(said)	R748LERT R748EWT R748CINT	(R748LERT + R748EWT) / R748CINT
I/O Intensity SCSI Total	ESIOIST	SERN(sern), SAID(said)	R748LSWT R748CINT	(R748LSRT + R748LSWT) / R748CINT
Transfer Rate PPRC Send	ESTRPSD	SERN(sern), SAID(said)	R748LPSPB R748CINT	R748LPSPB / R748CINT
Transfer Rate PPRC Received	ESTRPRV	SERN(sern), SAID(said)	R748LPSPB R748CINT	R748LPSPB / R748CINT
Packet Size PPRC Send	ESPSPSD	SERN(sern), SAID(said)	R748LPSPB R748LPPO	R748LPSPB / R748LPPO
Packet Size PPRC Received	ESPSPRV	SERN(sern), SAID(said)	R748LPSPB R748LPPO	R748LPSPB / R748LPPO
Activity Rate PPRC Send	ESARPSD	SERN(sern), SAID(said)	R748LPPO R748CINT	R748LPPO / R748CINT
Activity Rate PPRC Received	ESARPRV	SERN(sern), SAID(said)	R748LPPO R748CINT	R748LPPO / R748CINT
Response Time PPRC Send	ESRTPSD	SERN(sern), SAID(said)	R748LPST R748LPPO	R748LPST / R748LPPO
Response Time PPRC Received	ESRTPRV	SERN(sern), SAID(said)	R748LPST R748LPPO	R748LPST / R748LPPO

Table 36. ESS Link Statistics - Conditions Based on SMF Record 74-8 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
I/O Intensity PPRC Send	ESIOIPSD	SERN(sern), SAID(said)	R748LPST R748CINT	R748LPST / R748CINT
I/O Intensity PPRC Received	ESIOIPRV	SERN(sern), SAID(said)	R748LPRT R748CINT	R748LPRT / R748CINT
I/O Intensity PPRC Total	ESIOIPT	SERN(sern), SAID(said)	R748LPST R748LPRT SMF74INT	(R748LPST + R748LPRT) / R748CINT
Number of FC link failures	ESLFLF	SERN(sern) SAID(said)	R748LFLF	Value or comparison
Number of FC synchronization failures	ESLFLY	SERN(sern) SAID(said)	R748LFLY	Value or comparison
Number of FC signal failures	ESLFLS	SERB(sern) SAID(said)	R748LFLS	Value or comparison
Number of FC primitive sequence errors	ESLFPQ	SERB(sern) SAID(said)	R748LFPQ	Value or comparison
Number of FC invalid transmission word errors	ESLFIT	SERB(sern) SAID(said)	R748LFIT	Value or comparison
Number of FC Cyclic Redundancy Check (CRC) errors	ESLFCR	SERB(sern) SAID(said)	R748LFCR	Value or comparison
Number of FC link recovery (LR) sent	ESLFLR1	SERB(sern) SAID(said)	R748LFR1	Value or comparison
Number of FC link recovery (LR) received	ESLFLR2	SERB(sern) SAID(said)	R748LFR2	Value or comparison
Number of FC illegal frame errors	ESLFIF	SERB(sern) SAID(said)	R748LFIF	Value or comparison
Number of FC out of order data errors	ESLFOD	SERB(sern) SAID(said)	R748LFOD	Value or comparison
Number of FC out of order ACK errors	ESLFOA	SERB(sern) SAID(said)	R748LFOA	Value or comparison
Number of FC duplicate frame errors	ESLFDF	SERB(sern) SAID(said)	R748LFDF	Value or comparison
Number of FC invalid offset failures	ESLFIO	SERB(sern) SAID(said)	R748LFIO	Value or comparison
Number of FC sequence timeout errors	ESLFTC	SERB(sern) SAID(said)	R748LFTC	Value or comparison
FC bit error rate	ESLFBC	SERB(sern) SAID(said)	R748LFBC	Value or comparison

Table 37. ESS Extent Pool Statistics - Conditions Based on SMF Record 74-8

Criterion	Criterion Name	Qualifier	Source	Algorithm
Real Extent Pool Capacity	ESXRCAP	SERN(sern), XPID(xpid)	R748XRCP	Value or comparison
Number of Real Extents	ESXRNSG	SERN(sern), XPID(xpid)	R748XRNS	Value or comparison

Table 38. ESS Rank Statistics - Conditions Based on SMF Record 74-8

Criterion	Criterion Name	Qualifier	Source	Algorithm
Read Operation Rate	ESRROP	SERN(sern), RRID(rrid)	R748RRROP R748CINT	R748RRROP / R748CINT
Number of Bytes per Read Operation	ESRRBOP	SERN(sern), RRID(rrid)	R748RBYR R748RRROP	R748RBYR / R748RRROP
Average Bandwidth of Read Operation	ESRRBD	SERN(sern), RRID(rrid)	R748RBYR R748CINT	R748RBYR / R748CINT
Average Response Time of Read Operations	ESRRRT	SERN(sern), RRID(rrid)	R748RKRT R748RRROP	R748RKRT / R748RRROP
Write Operation Rate	ESRWOP	SERN(sern), RRID(rrid)	R748RWOP R748CINT	R748RWOP / R748CINT
Number of Bytes per Write Operation	ESRWBOP	SERN(sern), RRID(rrid)	R748RBYW R748RWOP	R748RBYW / R748RWOP
Average Bandwidth of Write Operations	ESRWBD	SERN(sern), RRID(rrid)	R748RBYW R748CINT	R748RBYW / R748CINT

Table 38. ESS Rank Statistics - Conditions Based on SMF Record 74-8 (continued)

Criterion	Criterion Name	Qualifier	Source	Algorithm
Average Response Time of Write Operations	ESRWRT	SERN(sern), RRID(rrid)	R748RKWT R748RWOP	R748RKWT / R748RWOP

PCIe Function Activity - SMF record type 74-9

Each Overview condition requires the specification of qualifier *pfid* that represents a PCIe function that is in status *Allocated*, *De-Allocate-Pending*, or *Re-Allocated* at the end of the reporting interval. For PCIe functions that are in another status, no overview report is generated.

pfid A four digit hexadecimal function ID for the PCIe function in the range of 0000 through FFFF. Example: (000F)

Table 39. PCIe Function Activity - Conditions Based on SMF Record 74-9

Condition	Condition Name	Qualifier	Source	Algorithm
Number of executed PCI Load operations per second	PCILOAD	pfid	R749LOOP R749ALLT	LOOP * 1000 / ALLT
Number of executed PCI Store operations per second	PCISTOR	pfid	R749STOP R749ALLT	STOP * 1000 / ALLT
Number of executed PCI Store Block operations per second	PCISTBL	pfid	R749SBOP R749ALLT	SBOP * 1000 / ALLT
Number of executed Refresh PCI Translation operations per second	PCIRPTR	pfid	R749RFOP R749ALLT	RFOP * 1000 / ALLT
Number of megabytes per second transferred by DMA reads from all defined DMA address spaces to the PCIe function (on zEC12 and zBC12 hardware only)	PCIDMAR	pfid	R749DMAR R749ALLT	DMAR / (ALLT * 1000)
Number of megabytes per second transferred by DMA writes from the PCIe function to all defined DMA address spaces (on zEC12 and zBC12 hardware only)	PCIDMAW	pfid	R749DMAW R749ALLT	DMAW / (ALLT * 1000)
Percentage of time this partition kept the hardware accelerator busy	FPGBUSY	pfid	R749FTET R749ALLT	FTET * 100 / (ALLT * 1000)
Average time in micro-seconds the hardware accelerator took to process a request	FPGRTIM	pfid	R749FTET R749FRQC	FTET / FRQC
Average queue time in microseconds spent for a request	FPGQTIM	pfid	R749FTQT R749FRQC	FTQT / FRQC
Number of megabytes per second transferred by DMA operations	FPGBYTS	pfid	R749FDRD R749FDWR R749ALLT	(FDRD + FDWR) * 256 / (ALLT * 1000)
Average request size in kilobytes	FPGBYTR	pfid	R749FDRD R749FDWR R749FRQC	(FDRD + FDWR) * 256 / (FRQC * 1000)
Number of compression requests per second	FPGCORS	pfid	R7491DCT R749ALLT	1DCT * 1000 / ALLT
Number of megabytes compressed per second	FPGCOBS	pfid	R7491DIB R749ALLT	1DIB / (ALLT * 1000)
Compression ratio	FPGCORT	pfid	R7491DIB R7491DOB	1DIB / 1DOB
Number of decompression requests per second	FPGDCRS	pfid	R7491ICT R749ALLT	1ICT * 1000 / ALLT
Number of megabytes decompressed per second	FPGDCBS	pfid	R7491IIB R749ALLT	1IIB / (ALLT * 1000)
Decompression ratio	FPGDCRT	pfid	R7491IIB R7491IOB	1IIB / 1IOB
Buffer pool size	FPGBPSZ	pfid	R7491BPS	1BPS
Average buffer pool utilization	FPGBPRT	pfid	R7491BPC R7491BPS R7491DCT R7491ICT	1BPC / ((1DCT + 1ICT) * 1BPS)

Table 39. PCIe Function Activity - Conditions Based on SMF Record 74-9 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Number of megabytes received per second (RoCE on z13 only)	PCIBYTR	pfid	R749DBYR R749ALLT	DBYR / (ALLT * 1000)
Number of megabytes transmitted per second (RoCE on z13 only)	PCIBYTT	pfid	R749DBYT R749ALLT	DBYT / (ALLT * 1000)
Number of packets received per second (RoCE on z13 only)	PCIPAKR	pfid	R749DPKR R749ALLT	DPKR * 1000 / ALLT
Number of packets transmitted per second (RoCE on z13 only)	PCIPAKT	pfid	R749DPKT R749ALLT	DPKT * 1000 / ALLT
Number of work units processed per second (zEDC on z13 only)	PCIWUP	pfid	R749DWUP R749ALLT	DWUP * 1000 / ALLT
PCI Function Utilization (zEDC on z13 only)	PCIUTIL	pfid	R749DWUP R749DWUM R749ALLT	DWUP * 100000 / (ALLT * DWUM)

Page Data Set Activity - SMF record type 75

One qualifier is possible:

pagename

The name of a page data set.

For exception processing, the data set name is optional. If this qualifier is omitted, the threshold value applies to each page data set.

Table 40. Page Data Set Activity - Conditions Based on SMF Record Type 75

Condition	Condition Name	Qualifier	Source	Algorithm
Percent busy-data set	PSBSY	pagename	SMF75USE SMF75SAM	(USE*100)/SAM
Page transfer time	PSPTT	pagename	SMF75REQ SMF75SAM SMF75PGX SMF75INT	((REQ*INT)/SAM)/PGX
Pages transfer rate	PSPT	pagename	SMF75INT SMF75PGX	PGX/INT
Page data set activity rate	PSART	pagename	SMF75INT SMF75SIO	SIO/INT
Average slots used	PSAVGSL	pagename	SMF75LVU	Value or comparison
Number of bad slots	PSBADS	pagename	SMF75BDS	Value or comparison

Enqueue Activity - SMF record type 77

One qualifier is possible:

major/minor

One- to eight-character major name of a resource, optionally followed by a comma and a one- to forty-four character minor name.

For exception processing only:

If this qualifier is omitted, the threshold value is checked for every minor name within each major name. If only the major name is specified, the threshold is checked for every minor name within the specified major name. A minor name cannot be specified without major name.

Table 41. Enqueue Activity - Conditions Based on SMF Record Type 77

Condition	Condition Name	Qualifier	Source	Algorithm
Total enqueue contention time in seconds	ENQT	major/minor	SMF77WTT	Value or comparison
Number of enqueue contention events	ENQNE	major/minor	SMF77EVT	Value or comparison

Table 41. Enqueue Activity - Conditions Based on SMF Record Type 77 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Average contention time	ENQAVG	major/minor	SMF77WTT SMF77EVT	WTT/EVT
Maximum contention time	ENQMAX	major/minor	SMF77WTX	Value or comparison

Virtual Storage Activity - SMF record type 78-2

If you use the conditions listed in Table 42 to define thresholds in EXCEPT statements, you must specify the values in Kbytes, for example, to define an SQA expansion condition with a threshold of 16384 bytes, specify:

EXCEPT(EXSQAE(SQAE,GE,16))

Table 42. Virtual Storage Activity - Conditions Based on SMF Record Type 78-2

Condition	Condition Name	Qualifier	Source	Algorithm
Maximum amount of allocated CSA below the 16-megabyte line in bytes	CSAA	none	R782CSAU	Value or comparison
Minimum size of largest free block - CSA in bytes	CSAFB	none	R782CSLF	Value or comparison
Minimum number of free CSA bytes below the 16-megabyte line	CSAFP	none	R782CSAF	Value or comparison
Maximum amount of allocated SQA below the 16-megabyte line in bytes	SQAA	none	R782SQUAU	Value or comparison
Maximum amount of SQA expansion into CSA in bytes	SQAE	none	R782SQEX	Value or comparison
Minimum size of largest free block - SQA in bytes	SQAFB	none	R782SQLF	Value or comparison
Minimum number of free SQA bytes below the 16-megabyte line	SQAFP	none	R782SQAF	Value or comparison

I/O Queuing Activity - SMF record type 78-3

One qualifier is possible:

lcuid A four-digit hexadecimal number that identifies a logical control unit.

For exception processing only:

This qualifier is optional. If it is omitted, the threshold applies to all logical control units in the SMF record.

iopid A two-digit hexadecimal number that identifies the I/O processor (IOP).

This qualifier is optional. If specified, RMF performs the exception checking just for the specified IOP identifier. If omitted, RMF loops through all IOP sections and uses the sum values for checking. In this case, the exception refers to the system wide value.

In the **Algorithm** column:

MAX Applies to exception operator GE, and specifies the sum of each channel path taken, where i represents channel path 0 to channel path 7.

MIN Applies to exception operator LE, and specifies the sum of each channel path taken, where i represents channel path 0 to channel path 7.

CHPID(i)

Is calculated as R783PB(i)/R783GSAM.

Table 43. I/O Queuing Activity - Conditions Based on SMF Record Type 78-3

Condition	Condition Name	Qualifier	Source	Algorithm
I/O processor (IOP) queue activity rate	IOPAC	iopid	R783IQCT SMF78INT	IQCT/INT
I/O processor (IOP) initiative queue average queue length	IOPQL	iopid	R783IQCT R783QSM	(IQSM - IQCT) / IQCT
Contention rate of an LCU	IOCTR	lcuid	R783QCT SMF78INT	QCT/INT
Average queue length of delayed I/O requests	IODLQ	lcuid	R783QCT R783QSM	QSM/QCT
Rate of LCU channel path taken	IOART	lcuid	R783PT SMF78INT	(PTi)/INT
Percentage of requests caused by control unit busy	IOCUB	lcuid	R783DPB R783CUB R783PT	MAX (CUBi*100)/ (PT+CUB+DPB)i MIN (CUBi*100)/ (PT+CUB+DPB)i
Percentage of requests caused by director port busy	IODPB	lcuid	R783DPB R783CUB R783PT	MAX (DPBi*100)/ (PT+DPB+CUB)i MIN (DPBi*100)/ (PT+DPB+CUB)i
Percent I/O processor busy	IOIPB	iopid	R783IIPB R783IIPi	(IIPB * 100) / (IIPB + IIPi)
Percent I/O processor idle	IOIPI	iopid	R783IIPB R783IIPi	(IIPi * 100) / (IIPB + IIPi)
Rate I/O functions started	IORIFS	iopid	R783IIFS SMF78INT	IIFS / SMF78INT
Rate processed I/O interrupts	IORPII	iopid	R783IPII SMF78INT	IPII / SMF78INT
Percent of I/O retries	IOPALB	iopid	R783ICHB R783IDPB R783ICUB R783IDVB R783IIFS	((Sum all retries) * 100) / (IIFS + Sum all retries)
Percent of I/O retries due to channel busy	IOPCHB	iopid	R783ICHB R783IDPB R783ICUB R783IDVB R783IIFS	(ICHB * 100) / (IIFS + Sum all retries)
Percent of I/O retries due to director port busy	IOPDPB	iopid	R783ICHB R783IDPB R783ICUB R783IDVB R783IIFS	(IDPB * 100) / (IIFS + Sum all retries)

Table 43. I/O Queuing Activity - Conditions Based on SMF Record Type 78-3 (continued)

Condition	Condition Name	Qualifier	Source	Algorithm
Percent of I/O retries due to control unit busy	IOPCUB	iopid	R783ICHB R783IDPB R783ICUB R783IDVB R783IIFS	$(ICUB * 100) / (IIFS + \text{Sum all retries})$
Percent of I/O retries due to device busy	IOPDVB	iopid	R783ICHB R783IDPB R783ICUB R783IDVB R783IIFS	$(IDVB * 100) / (IIFS + \text{Sum all retries})$
Number of I/O retries per SSCH	IONALB	iopid	R783ICHB R783IDPB R783ICUB R783IDVB R783IIFS	$(\text{Sum all retries}) / IIFS$
Number of I/O retries per SSCH due to channel busy	IONCHB	iopid	R783ICHB R783IIFS	$ICHB / IIFS$
Number of I/O retries per SSCH due to director port busy	IONDPB	iopid	R783IDPB R783IIFS	$IDPB / IIFS$
Number of I/O retries per SSCH due to control unit busy	IONCUB	iopid	R783ICUB R783IIFS	$ICUB / IIFS$
Number of I/O retries per SSCH due to device busy	IONDVB	iopid	R783IDVB R783IIFS	$IDVB / IIFS$
Average control unit busy delay time	IOCBT	lcuid	R783CBT R783PT	$\text{Sum}(R783CBT) / \text{Sum}(R783PT)$
Average initial command response time	IOCMR	lcuid	R783CMR R783PT	$\text{Sum}(R783CMR) / \text{Sum}(R783PT)$
Average channel subsystem delay time	IOCSS	lcuid	R783CSST R783PT	$R783CSST / \text{Sum}(R783PT)$
HyperPAV wait ratio	IOHWAIT	lcuid	R783HNAI R783HTIO	$R783HNAI / R783HTIO$
Maximum number of in-use HyperPAV aliases for the LCU	IOHMAX	lcuid	R783HAIU	Value or comparison
Maximum number of in-use HyperPAV aliases for one device	IOHDMAX	lcuid	R783HCAD	Value or comparison
The high watermark of queued I/O requests	IOHIOQC	lcuid	R783HIOQ	Value or comparison

Chapter 16. Cross platform monitoring with RMF XP

For installations running operating systems other than z/OS, RMF XP provides a solution to monitor the performance of heterogeneous environments. RMF XP supports the following operating systems:

- AIX on System p
- Linux on System x
- Linux on System z
- Windows on System x

RMF XP has been tested for the following versions or distributions of the supported operating systems:

- Red Hat RHEL 5.6 and 6.0 (64-bit)
- SUSE SLES 11 SP1 (64-bit)
- AIX 6.1.0 and 7.1.0
- Windows Server 2008 SP2

Hence, with RMF XP, you can monitor operating systems which run on an IBM zEnterprise System, including the zEnterprise Blade Center Extension (zBX).

Additionally, RMF XP does not require any proprietary agent software on the monitored endpoints. It exploits the existing Common Information Model (CIM) instrumentation for the AIX, Linux and Windows operating systems. The CIM server, as well as the metric providers, are integral parts of the supported AIX and Linux distributions, and therefore no additional software needs to be installed. You just need to ensure that the CIM servers with their metric providers are properly set up and running on the monitored endpoints.

With regard to Linux, RMF XP can communicate with two different CIM server implementations, the Open Pegasus CIM server and the Small Footprint CIM Broker (SFCB).

For the Windows operating system, the IBM Systems Director platform agent is required to be active on the Windows endpoints to enable RMF cross-platform monitoring (RMF XP) for Windows.

Please refer to IBM Systems Director v6.3.2 for more information about the ISD platform agents.

For the purpose of evaluating the functions provided by RMF XP, a no-charge version of the IBM Systems Director agent can be found at the following website: IBM Systems Director agents download. This agent is not entitled to a support contract and an account would use it on an as-is basis. IBM recommends that an appropriate support contract be active with the IBM Systems Director agent on the monitored endpoint for Windows. If the account does not have a support contract for the IBM Systems Director endpoint, the recommended approach would be IBM Systems Director Express® Edition. It can be obtained by contacting your local IBM representative and ordering appropriate entitlements for IBM Systems Director Express Edition for System x.

Additional information on IBM Systems Director Express Edition for System x is located at the following website: IBM systems management solutions for System x, Editions

The core component of RMF XP is the GPM4CIM server. Similar to the existing Distributed Data Server (DDS) for z/OS (also known as GPMSEVER), the GPM4CIM started task receives HTTP requests and sends back responses as structured XML documents. Since the GPM4CIM started task runs in the z/OS Unix environment, at least one z/OS system is needed to utilize the RMF XP component.

As shown in Figure 43, the existing RMF DDS for z/OS (GPMSEVER) remains unchanged and does not interfere with the RMF XP started task. Thus, RMF XP provides an additional RMF DDS flavour for the CIM instrumented operating systems AIX, Linux and Windows. For this reason the RMF XP started task is denoted as GPM4CIM.




Figure 43. Cross platform performance monitoring

The left part of Figure 43 shows the configuration in a z/OS environment and the right part shows the configuration in a mixed AIX, Linux and Windows environment. One instance of an RMF generic CIM client, which is a logical module of the GPM4CIM server, communicates with a CIM server residing on either a Linux, AIX or Windows system anywhere in the heterogeneous environment. In a mixed AIX, Linux and Windows environment, you need to start at least one specific GPM4CIM instance for each platform, as described in "Configuration files and parameters" on page 278.

You can exploit the RMF XP capabilities in the following ways:

- Users of the *Resource Monitoring* plug-in for the *IBM z/OS Management Facility (z/OSMF)* can define connections to GPM4CIM servers running in the z/OS Unix environment. Thus, you can display performance data from AIX, Linux and Windows systems in the same way as z/OS performance data (see also Chapter 20, “z/OS Management Facility - Resource Monitoring,” on page 369). Of course, you can also combine performance data from different platforms into a common view as required.
- Besides handling requests for z/OS performance data, the HTTP API of the DDS can serve requests for AIX, Linux and Windows performance data as well, as soon as a GPM4CIM instance is configured and active. The DDS returns the requested data as a structured XML document. See the *z/OS RMF Programmer's Guide* for the documentation of the DDS HTTP API.
- You can exploit the capability of RMF XP to monitor the long-term performance history of AIX, Linux and Windows systems. For this purpose, you request that performance data collected from the managed endpoints is written to SMF record type 104. Thus, RMF XP offers a standard method for detailed long-term performance analysis and capacity planning.

For information on how to request the collection of SMF record type 104 from the systems of all or selected supported platforms, refer to “How to use RMF XP for long-term performance analysis” on page 281.

How to set up RMF XP

To start the GPM4CIM server from the console, RMF provides procedure GPM4CIM as a member in SYS1.PROCLIB, which you must modify according to your needs:

```
//GPM4CIM PROC OS=A
//*****
//* STEP 1 - Execute GPM4CIM *
//*****/
//STEP1 EXEC PGM=BPXBATCH,TIME=NOLIMIT,REGION=0M,
// PARM='PGM /usr/lpp/gpm/bin/gpm4cim cfg=/etc/gpm/gpm4 &OS. .cfg'
//STDENV DD PATH='/etc/gpm/gpm4cim.env'
//STDOUT DD PATH='/var/gpm/logs/gpm4cim&OS..out',
// PATHOPTS=(OWRONLY,OCREAT,OTRUNC),
// PATHMODE=(SIRUSR,SIWUSR,SIRGRP)
//STDERR DD PATH='/var/gpm/logs/gpm4cim&OS..trc',
// PATHOPTS=(OWRONLY,OCREAT,OTRUNC),
// PATHMODE=(SIRUSR,SIWUSR,SIRGRP)
//*****
//* STEP 2 - Copy stdout back to joblog *
//*****/
//STEP2 EXEC PGM=BPXBATCH,
// PARM='PGM /bin/cat /var/gpm/logs/gpm4cim&OS..out'
//STDOUT DD SYSOUT=*
//STDERR DD SYSOUT=*
//*****
//* STEP 3 - Copy stderr back to joblog *
//*****/
//STEP3 EXEC PGM=BPXBATCH,
// PARM='PGM /bin/cat /var/gpm/logs/gpm4cim&OS..trc'
//STDOUT DD SYSOUT=*
//STDERR DD SYSOUT=*
//*****
// PEND
```

Figure 44. RMF XP startup member in SYS1.PROCLIB(GPM4CIM)

The procedure invokes the IBM-supplied BPXBATCH utility which passes control to the gpm4cim load module. This module is located in the HFS directory /usr/lpp/gpm/bin.

The **cfg** parameter in the PARM statement points to the GPM4CIM configuration file. Since one instance of GPM4CIM is needed per platform, no unique configuration file is used. You can supply different configuration files using the **OS** variable to denote the target platform:

A	AIX on System p (as shown in the example from Figure 44 on page 277)
X	Linux on System x
Z	Linux on System z
W	Windows on System x

For the GPM4CIM procedure, the environment variables are kept in file **gpm4cim.env** which is specified with the STDENV ddname. Log and trace output is written to the files specified with the STDOUT and STDERR ddnames. If multiple instances of GPM4CIM are running simultaneously, you can specify individual output files by altering the file names in the PATH parameter.

Configuration files and parameters

The GPM4CIM parameters are supplied with the platform specific configuration files /etc/gpm/gpm4[A|X|Z|W].cfg. This allows you to run one separate GPM4CIM instance per platform, which is required if you want to monitor AIX, Linux and Windows systems in a mixed environment.

Of course, you can also start multiple instances of GPM4CIM for the same platform. In this case, provide a dedicated copy of the configuration file per instance, for example, *gpm4a1.cfg* and *gpm4a2.cfg* and use these names in the GPM4CIM procedures (Figure 44 on page 277).


```

/*****
/*
/* Name:          gpm4A.cfg
/*
/*
/* Description: Configuration for the RMF Distributed Data Server
/*
/*          USS/CIM (GPM4CIM)
/*
/*
/*****
CACHESLOTS(3)          /* Number of timestamps in CACHE */
DEBUG_LEVEL(0)         /* No informational messages */
SERVERHOST(*)          /* Don't bind to specific IP-Address */
TIMEOUT(0)             /* No timeout */
MAXSESSIONS_HTTP(20)   /* MaxNo of concurrent HTTP requests */
HTTP_PORT(8805)        /* Port number for HTTP requests */
HTTP_ALLOW(*)          /* Mask for hosts that are allowed */
HTTP_NOAUTH()          /* No server can access without auth.*/
/*****
/* Declaration of Host Connections with CIM Instrumentation */
/* =====
/* - INTERVAL is the monitoring interval length,
/*   value is given in seconds.
/* Example: INTERVAL(900)
/* Default:  INTERVAL(300)
/*
/* - AIX_Complex is the name for the root resource.
/* Example: AIX_COMPLEX(SAPPLEX)
/*
/* - AIX_IMAGE is the host name or IP-address of a
/* CIM connection including the port number.
/* Example: AIX_IMAGE(SAP1.US.IBM.COM:5988)
/*
/*****
INTERVAL(300)          /* Data monitoring interval (seconds)*/
RECORD(YES)            /* Collect SMF record type 104 data */
AIX_COMPLEX(P6PLEX)    /* User defined name of CIM complex */
/* AIX images following here
AIX_IMAGE(P6RMF1.DE.ABC.COM:5988)
AIX_IMAGE(P6RMF2.DE.ABC.COM:5988)
AIX_IMAGE(P6RMF3.DE.ABC.COM:5988)

```

Figure 45. GPM4CIM configuration file example for AIX

The parameters CACHESLOTS through HTTP_NOAUTH are common parameters for both procedures, GPM4CIM and GPMSEVERE, as described in “DDS options” on page 28.

The following parameters are specific for GPM4CIM only:

AIX_COMPLEX | LNX_COMPLEX | LNZ_COMPLEX | WIN_COMPLEX

Name for the root resource of the GPM4CIM instance. You can use your own name. This name appears as resource name for the resource type AIX_SYSTEM_COMPLEX for AIX, XLINUX_SYSTEM_COMPLEX for Linux on System x, ZLINUX_SYSTEM_COMPLEX for Linux on System z, or WINDOWS_SYSTEM_COMPLEX for Windows on System x.

AIX_IMAGE | LNX_IMAGE | LNZ_IMAGE | WIN_IMAGE

Host name or IP address of an image running the specified operating system, including the port number. Due to different resource models, all images within the same complex must run the same operating system.

INTERVAL

Length of the monitoring interval in seconds. Specify a value between 60 and 3600 seconds. The default is 300 seconds.

The interval parameter determines the frequency, with which RMF retrieves data from the monitored endpoints. It does not necessarily match the data collection interval of the CIM providers on this systems:

- For endpoints running the AIX operating system, GPM4CIM tries to synchronize the CIM provider interval with the GPM4CIM specified interval. This function requires AIX version 6.2.
- For endpoints running the Linux or Windows operating system, it is the administrator's responsibility to synchronize the interval start time and the interval length of the CIM providers with the GPM4CIM specified interval.

RECORD(YES|NO)

Specifies whether or not measurement data is collected in SMF record type 104 as defined by the settings in the active SMFPRMxx parmlib member. The default is RECORD(NO). Therefore, if you want to collect SMF records with GPM4CIM, you need to specify RECORD(YES).

Post-installation steps

Before using RMF XP for the first time, you must run the *gpm4cim_setup.sh* script. This script is located in the HFS directory */usr/lpp/gpm/bin*. If necessary, it allocates the */etc/gpm* and */var/gpm/logs* directories and copies the provided configuration files for all operating system types to the directory */etc/gpm*.

Also, you must adapt the configuration files to your environment before starting a GPM4CIM instance.

Security setup

You can turn on or off RACF authentication for the communication between z/OSMF or other clients and the GPM4CIM server. In case authentication is required (depending on the HTTP_NOAUTH parameter), z/OSMF supplies passtickets for the logon to the GPM4CIM started task. The client's credentials are then checked against RACF.

You can apply the information provided for GPMSERVE in “Configuring PassTicket support for the Distributed Data Server” on page 19 to GPM4CIM accordingly. You need to use GPM4CIM as application name instead of GPMSERVE.

If your installation requires secure communication between GPM4CIM and the CIM servers on the monitored endpoints, you can set up encryption. GPM4CIM does not provide encryption support explicitly, but you can exploit the z/OS Application Transparent Transport Layer Security (AT-TLS) in order to set up secure server-to-server communication. In this case you need to ensure that the CIM servers on the monitored endpoints are also configured for encryption. For more information about CIM servers running on System z, refer to *z/OS Common Information Model User's Guide*. For information about CIM servers running on System p or System x, refer to the appropriate documentation.




Figure 46. GPM4CIM security concept

How to start GPM4CIM and assign userIDs

You can start the GPM4CIM server from the console as started task with the following command:

```
s gpm4cim[.identifier],os=A|X|Z|W
```

Since you can run multiple GPM4CIM instances simultaneously, it is recommended to assign an identifier which you can use for subsequent STOP or MODIFY commands.

You may already have created the userID GPMSEVERE as owner of the GPMSEVERE procedure. The GPM4CIM started task can be assigned to the same userID with the following command:

```
RDEFINE STARTED GPM4CIM.* STDATA(USER(GPMSEVERE) TRUSTED(YES))
```

The GPM4CIM server uses the RACF application name GPM4CIM. Check if the use of this application is protected by a profile in the active RACF Application class (APPL) as described in “Ensure RACF access to the Distributed Data Server (GPMSEVERE and GPM4CIM)” on page 19. If required, use the following command to grant user access:

Security Server (RACF) Example

```
RDEFINE APPL GPM4CIM UACC(READ)
```

How to use RMF XP for long-term performance analysis

To enable long-term performance analysis of AIX, Linux and Windows systems, you can turn on SMF recording for SMF record type 104. This record type provides one range of subtypes for each supported platform. One specific subtype is used to keep the data for one individual CIM metric category according to the CIM data

model on the affected platform. The CIM metric category, in return, is mapped to the resource models used by RMF XP on the supported platforms as described in *z/OS RMF Programmer's Guide*.

An overview of the available subtype ranges is presented in this topic. For detailed information about the mapping of CIM metric categories to RMF XP resource types, refer to *z/OS MVS System Management Facilities (SMF)*.

AIX on System p performance data: subtype 1-12

Subtype 1

AIX_ActiveMemoryExpansionMetrics

Subtype 2

AIX_ProcessorMetrics

Subtype 3

AIX_ComputerSystemMetrics

Subtype 4

AIX_DiskMetrics

Subtype 5

AIX_NetworkPortMetrics

Subtype 6

AIX_FileSystemMetrics

Subtype 7

AIX_MemoryMetrics

Subtype 8

AIX_OperatingSystemMetrics

Subtype 9

AIX_ProcessMetrics

Subtype 10

AIX_SharedEthernetAdapterMetrics

Subtype 11

AIX_ActiveMemorySharingMetrics

Subtype 12

AIX_VirtualTargetDeviceMetrics

Linux on System x performance data: subtype 20-31

Subtype 20

Linux_IPProtocolEndpointMetrics

Subtype 21

Linux_LocalFileSystemMetrics

Subtype 22

Linux_NetworkPortMetrics

Subtype 23

Linux_OperatingSystemMetrics

Subtype 24

Linux_ProcessorMetrics

Subtype 25

Linux_UnixProcessMetrics

Subtype 26

Linux_StorageMetrics

Subtype 30

Linux_KVMMetrics

Subtype 31

Linux_XenMetrics

Linux on System z performance data: subtype 40-53

Subtype 40	Linux_IPProtocolEndpointMetrics
Subtype 41	Linux_LocalFileSystemMetrics
Subtype 42	Linux_NetworkPortMetrics
Subtype 43	Linux_OperatingSystemMetrics
Subtype 44	Linux_ProcessorMetrics
Subtype 45	Linux_UnixProcessMetrics
Subtype 46	Linux_StorageMetrics
Subtype 50	Linux_zCECMetrics
Subtype 51	Linux_zLPARMetrics
Subtype 52	Linux_zChannelMetrics
Subtype 53	Linux_zECKDMetrics

Windows on System x performance data: subtype 60-64

Subtype 60	Windows_LocalFileSystemMetrics
Subtype 61	Windows_NetworkPortMetrics
Subtype 62	Windows_OperatingSystemMetrics
Subtype 63	Windows_ProcessorMetrics
Subtype 64	Windows_StorageMetrics

How to request SMF record type 104 collection

At least one instance of RMF XP (GPM4CIM) must be running to perform data collection for a certain supported operating system. Different instances of RMF XP for different target operating systems can run either on the same or on different z/OS system(s). The z/OS system(s) must have an active connection to the monitored endpoints with a properly configured CIM server up and running as previously described. The collected SMF records provide image control sections to allow exploiters of the SMF records to associate the performance data with the appropriate systems.

To request measurement data collection from GPM4CIM into SMF record type 104, you must perform both of the following steps:

1. Using option TYPE, define the appropriate settings for SMF record type 104 in the active SMF parmlib member SMFPRMxx, for example:

```
SYS(TYPE(104(1:12)))    /* for AIX on System p */
SYS(TYPE(104))          /* for all platforms */
```

See “Defining SMF record writing” on page 30 or *z/OS MVS Initialization and Tuning Reference* for more details.

2. In your GPM4CIM configuration file, specify RECORD(YES) to indicate that SMF 104 type is to be collected as defined in the active SMFPRMxx parmlib member. The default RECORD(NO) suppresses any SMF record writing by GPM4CIM.

You can dynamically change GPM4CIM's SMF record collection as defined in the configuration file by the following commands:

f gpm4cim,record

If record type 104 or selected subtypes are specified in the active SMF parmlib member SMFPRMxx, this command switches on the according SMF recording starting with the next interval.

f gpm4cim,norecord

This command switches off any active SMF recording starting with the next interval.

For information on how to print or view the SMF type 104 records, refer to topic **Printing SMF records** in *z/OS RMF Programmer's Guide*.

The written SMF records can also be made available via the services of the RMF Sysplex Data Server (ERBDSQRY/ERBDSQ64 and ERBDSREC/ERBDSR64) as described in *z/OS RMF Programmer's Guide*. For this purpose, you must specify SMF record type 104, optionally reduced to required subtypes, with the SMFBUF parameter in one of the following ways:

- In the cataloged RMF procedure, use the SMFBUF parameter in the PARM option, for example:

```
//IEFPROC EXEC PGM=ERBMFMFC,REGION=256M,
//          PARM='SMFBUF(SPACE(40M),RECTYPE(70:79,104))'
```
- In the START command for the RMF control session, use the SMFBUF parameter to override the specifications in the RMF procedure, for example:

```
S RMF,,,SMFBUF(RECTYPE(70:78,79(1:15),104(1:12))) /* for AIX on System p */
S RMF,,,SMFBUF(RECTYPE(70:78,104(20:31)))          /* for Linux on System x */
S RMF,,,SMFBUF(RECTYPE(70:79,104(40:53)))          /* for Linux on System z */
S RMF,,,SMFBUF(RECTYPE(70:79,104(60:64)))          /* for Windows on System x */
```
- In the MODIFY command for the RMF control session, use the SMFBUF parameter to override the specifications on the START command or in the cataloged RMF procedure.

For detailed information on how to specify the SMFBUF parameter, refer to "Controlling the SMF buffer" on page 47.

How to authorize GPM4CIM to write SMF record type 104

To write SMF record type 104, the GPM4CIM started task needs at least READ access to the BPX.SMF profile of the FACILITY class, specified with your security authorization facility (SAF) product. The following is an example for RACF, where GPMSERVE is the user ID which is assigned to the GPM4CIM started task (see also "How to start GPM4CIM and assign userIDs" on page 281):

```
RDEFINE FACILITY BPX.SMF UACC(NONE)
PERMIT BPX.SMF CLASS(FACILITY) ID(GPMSERVE) ACCESS(READ)
SETROPTS RACLIST(FACILITY) REFRESH
```

Diagnostic aids

The GPM4CIM started task writes diagnostic information into a log and into a trace file. The log file is written unconditionally and contains information about the supplied options and the GPM4CIM started task's basic activity (see Figure 47). When the GPM4CIM started task is stopped, the log file entries are copied from file `/var/gpm/logs/gpm4cim[A|X|Z|W].out` to the file specified with the STDERR ddname. Specifying a `DEBUG_LEVEL > 0` in the configuration file prints additional information about the requests served by GPM4CIM into the log file.

An example for log file entries is shown in Figure 47.

```
(01) 02/25 110544 DDSRV: RMF-DDS-Server GPM4CIM/ZOSV1R13 started.
(01) 02/25 110544 DDSRV: Functionality Level=3.000
(01) 02/25 110544 DDSRV: Reading CFG file /etc/gpm/gpm4X.cfg.
(01) 02/25 110544 DDSRV: Summary of options in effect:
(01) 02/25 110544     CACHESLOTS(6)
(01) 02/25 110544     SERVERHOST(*)
(01) 02/25 110544     DEBUG_LEVEL(3)
(01) 02/25 110544     HTTP_PORT(8806)
(01) 02/25 110544     MAXSESSIONS_HTTP(20)
(01) 02/25 110544     HTTP_ALLOW(*)
(01) 02/25 110544     HTTP_NOAUTH()
(01) 02/25 110544     LNX_COMPLEX(SAPPLEX)
(01) 02/25 110544     LNX_IMAGE(SAP1)
(01) 02/25 110544     LNX_IMAGE(SAP2)
(01) 02/25 110544     LNX_IMAGE(SAP3)
(01) 02/25 110544     INTERVAL(60)
(01) 02/25 110544     RECORD(NO)
(01) 02/25 110544 DDSRV: OS-Type='XLINUX' OS-Version='SP7.1.1'.
(01) 02/25 110544 DDSRV: User='GPM4CIM' Process/Jobname='GPM4CIM'
ProcessId=6711008.
(01) 02/25 110544 DDSRV: System Id='SAP1' Sysplex='SAPPLEX' GoalMode=1.
(01) 02/25 110544 DDSRV: TCP/IP local host name=sap1.us.ibm.com.
(01) 02/25 110544 DDSRV: TCP/IP IP Address=9.152.87.80.
(01) 02/25 110544 DDSRV: Size of first cache slot=0 (rounded to 65536) Bytes.
(01) 02/25 110544 DDSRV: Timezone='GMT-01:00:00'
(01) 02/25 110544 DDSRV: Result ctime_r()=Fri Feb 25 11:05:44 2011
(04) 02/25 110544 DDCON: GPM260I RMF Distributed Data Server for CIM ready for
commands
```

Figure 47. Log file entries during GPM4CIM processing

If you require more detailed information about GPM4CIM activity, you can activate a trace by means of the following directive:

```
ICLUI_TRACETO=STDERR
```

This directive can be specified within the environment file `/etc/gpm/gpm4cim.env`.

When you stop GPM4CIM, the trace file entries are copied from file `/var/gpm/logs/gpm4cim[A|X|Z|W].trc` to the file specified with the STDERR ddname.

Part 7. Analysis on the workstation

In addition to host-based reporting functions in RMF, there are other components available that offer reporting capabilities on the workstation.

The **RMF Postprocessor XML Toolkit** is a feature that helps you to browse the Postprocessor reports that are available in XML output format, with your internet browser using the stylesheets provided by RMF. Just download the Postprocessor reports produced in XML output format into the Postprocessor XML Toolkit directory and view them with your browser, using the XSL stylesheets provided by RMF. The required stylesheet files are available in a Postprocessor XML Toolkit subdirectory.

The **RMF Spreadsheet Reporter** assists you in converting Postprocessor listings and Overview records into spreadsheets. The spreadsheet applications are shipped with RMF, and must be downloaded to your workstation before you can start. With the Spreadsheet Reporter, you can convert Postprocessor listings and Overview records to spreadsheets. This enables you to handle RMF data using techniques familiar to every spreadsheet user. In addition, the Spreadsheet Reporter provides sample macros to help you presenting and analyzing performance data at a glance.

RMF PM allows you to monitor the performance of your z/OS host from a workstation through a TCP/IP interface to one or more sysplexes. You logon to any sysplex and you can monitor the resources in the corresponding sysplex.

RMF Client/Server Enabling (RMFCS) is a concept that makes your performance management independent of a TSO host session. It allows you to establish multiple ISPF GUI sessions with any z/OS system in your network that has a Communications Server network connection configured to your workstation. This way, RMFCS combines the advantages of a single point of control for performance management with a state-of-the-art user front end.

IBM z/OS Management Facility (z/OSMF) is a web-browser based management console for z/OS. The **z/OSMF Resource Monitoring** plug-in allows cross-sysplex performance monitoring from a single point of control. From the z/OSMF task tree, you can select the following subtasks:

- The **Sysplex Status task** provides an enterprise-wide health check of all z/OS sysplexes.
- For further analysis, the **Monitoring Desktops task** can graphically display RMF Monitor III as well as AIX, Linux, or Windows metrics by means of customizable views.

Chapter 17. How to work with Postprocessor XML reports

RMF offers a set of components and features that assist you in comfortably producing and/or viewing your Postprocessor reports that you have requested in XML output format:

- “Producing and viewing XML reports with the Spreadsheet Reporter”
- “Producing and viewing XML reports with the HTTP API of the DDS”
- “Viewing XML reports with the RMF Postprocessor XML Toolkit”

Producing and viewing XML reports with the Spreadsheet Reporter

The Spreadsheet Reporter offers an option to produce selected Postprocessor reports in XML output format. Refer to Table 45 on page 319 for supported reports. For these reports, the Spreadsheet Reporter also offers an action for viewing them in a web browser. For more information, refer to Chapter 18, “RMF Spreadsheet Reporter,” on page 293.

Producing and viewing XML reports with the HTTP API of the DDS

You can also use the HTTP API of the Distributed Data Server to request one or more of the Postprocessor reports available in XML output format. For a detailed description of the syntax of such requests, refer to *Accessing performance data using the RMF Distributed Data Server* in the *z/OS RMF Programmer's Guide*.

Viewing XML reports with the RMF Postprocessor XML Toolkit

You produce Postprocessor reports in XML output format using the appropriate ddnames (XPOVWRPT, XPRPTS, or XPXSRPTS) in your job when requesting these reports, as specified in Table 24 on page 212. Topic *Long-term overview reporting with the Postprocessor* in *z/OS RMF Report Analysis* provides a table that indicates which Postprocessor reports are available as XML reports.

The *RMF Postprocessor XML Toolkit* assists you in browsing the produced XML Postprocessor reports with your internet browser, using the RMF XSL stylesheet.

How to install the Postprocessor XML Toolkit

1. The Postprocessor XML Toolkit is part of the RMF product. The application files and installation utility of the Postprocessor XML Toolkit are provided in member ERBXMLTK of the host distribution library SERBPWSV. Download this member as binary file **erbxmltk.msi**.
2. Install the MSI package using the Windows Installer, either by double-clicking on the MSI package file or by issuing the command:

```
msiexec /package erbxmltk.msi
```

The Windows Installer guides you through the installation.
3. Specify the directory where to install the Postprocessor XML Toolkit. The default, for example for Windows XP, is:

```
C:\Documents and Settings\_userid_\Application Data  
    \RMF\RMF Postprocessor XML Toolkit
```

The Postprocessor XML Toolkit is installed into program group *IBM RMF Performance Management*.

How to use the Postprocessor XML Toolkit

To view an XML Postprocessor report, download the created data set containing the XML output into the Postprocessor XML Toolkit directory on your workstation with file extension `.xml`. The required stylesheet files are available in a subdirectory of the Postprocessor XML Toolkit. The created Postprocessor reports in XML output format contain a link to the stylesheet in this subdirectory. When you open the XML Postprocessor reports within the Postprocessor XML Toolkit with your internet browser, the RMF stylesheet transforms the report into an HTML document.

Ensure to download the data set containing the XML output of the Postprocessor reports in ASCII format to the Postprocessor XML Toolkit directory.

Tip:

Check the example subdirectory of the Postprocessor XML Toolkit. There you find more information on how to exploit RMF Postprocessor XML reports. You also find an example for an RMF Postprocessor realtime reporting setup with sysplex wide scope.

Enjoy it!

How to view XML reports in a browser

When you load a Postprocessor XML report into a browser, you see the report sections collapsed. To view all sections and subsections, use the ***Expand all sections*** button (the plus signs) in the upper right corner. The sections expand, and the plus signs change to minus signs as shown in Figure 48 on page 291. To expand a single section, click on the arrow left to the section name.

Postprocessor XML reports provide a capability to sort tabular data wherever suitable. When you move the cursor over a column header, and the column contains sortable data, then you can click on the header to obtain the data sorted in either ascending or descending order as shown in the expanded ***I/O Processor (IOP) Data*** section in the ***I/O Queuing Activity Report*** sample of Figure 48 on page 291.




Figure 48. Sortable data in Postprocessor reports in XML format

Chapter 18. RMF Spreadsheet Reporter

This information unit covers the following topics:

- “Concepts of performance analysis with the RMF Spreadsheet Reporter”
- “Installing the Spreadsheet Reporter” on page 298
- “How to use the RMF Spreadsheet Reporter” on page 299
- “How to create Working Sets in batch mode” on page 323
- “Using RMF spreadsheet macros” on page 326

Concepts of performance analysis with the RMF Spreadsheet Reporter

The RMF Spreadsheet Reporter is the powerful workstation solution for graphical presentation of RMF Postprocessor data. Use it to convert your RMF data to spreadsheet format and generate representative charts for all performance relevant areas.

The RMF Spreadsheet Reporter offers the following features:

- ease of use - manage the related resources by means of an Explorer-like GUI
- fast path to graphical presentation - prepare the SMF data in one single step
- batch mode - generate the input files for the spreadsheets without any GUI interaction

Performance data derived from SMF records is the basis for z/OS performance analysis and capacity planning. The basic idea of the RMF Spreadsheet Reporter is to exploit the graphical presentation facilities of a workstation for these purposes: it extracts performance measurements from SMF records, produces Postprocessor Report Listings and Overview Records and converts this Postprocessor output into spreadsheets. Thus, the Spreadsheet Reporter offers a complete solution of enhanced graphical presentation of RMF measurement data.

The Spreadsheet Reporter also provides several sample spreadsheet macros to help you in viewing and analyzing performance data at a glance.

Figure 49 on page 294 shows a sample chart created with the Spreadsheet Reporter. It shows the CPU utilization of selected workloads for the prime shift.




Figure 49. Spreadsheet - CPU utilization for selected workloads

DISCLAIMER OF WARRANTIES

All spreadsheet macros are sample code created by IBM. They are not part of any standard IBM product and are provided to you solely for the purpose of assisting you in the development of your applications, and to demonstrate what can be done with RMF performance data.

This includes that service and upgrades for the macros through the standard IBM service channels are not available. In addition, you should be aware that there is no guarantee that the spreadsheet macros will work on your system or with your spreadsheet application, even if you use one of the listed versions.

Nevertheless, enhancements and possible fixes for the spreadsheet macros may become available in future, but they will not be shipped through the standard IBM service channels. You should monitor the RMF homepage available at <http://www.ibm.com/systems/z/os/zos/features/rmf/>

for enhancements: You find information about the Spreadsheet Reporter on the Tools page, which is accessible from the RMF homepage.

Get more out of Postprocessor reports with the Spreadsheet Reporter

RMF Postprocessor reports are the preferred media to analyze SMF data for long-term reporting. However, these are tabular report listings, and it can be cumbersome to analyze their content or retrieve exactly the required performance information. With the Spreadsheet Reporter, you can easily convert these reports to spreadsheet format, and use spreadsheet macros for graphical presentation.

Graphical charts are a comfortable medium for performance analysis and also significantly improve the capability of long-term and trend reporting of RMF data.

Besides **Report Listings**, RMF provides two flexible mechanisms to prepare long-term Postprocessor performance data for trend analysis: **Overview Reports** and **Overview Records**. They offer the capability to retrieve **selective** metrics from your system - in contrast to Postprocessor reports, which contain a predefined set of metrics in each report.

Concepts of Overview Reports and Overview Records

As mentioned above, Overview Reports and Overview Records are used for long-term performance analysis. Both are generated from overview control statements which specify the performance metrics that you want to examine. The Spreadsheet Reporter provides several macros that generate these statements for you, so that you do not need to know about their syntax.

Overview Reports have the same layout as the Summary report. However, in deviation to the Summary report, the Overview Report does not have a fixed layout, but the overview control statements specified for the Postprocessor define which metrics should be included in the report. This gives you a high flexibility because you can retrieve nearly all fields of the single-system and sysplex reports with the help of these statements. Thus, the Overview Report provides selected performance data and summarizes system activity for each interval within the reporting period.

With Postprocessor statement `OVERVIEW(REPORT)` you specify that you want to generate an Overview Report. Overview Reports provide a readable, tabular format of performance measurement data.

Overview Records are designed for further processing with spreadsheet macros. With the statement `OVERVIEW(RECORD)` you specify that you want to generate Overview Records (instead of an Overview Report).

Also you can use Overview Records with your own applications to process RMF data. For this purpose, you find a description of the Overview Record structure in *z/OS RMF Report Analysis*.




Figure 50. Overview Report and Spreadsheet

Figure 50 shows how the same metrics retrieved by a common set of overview control statements are presented in an Overview Report as well as in a spreadsheet. For example, with the overview control statement `OVW(CPUBSY(CPUBSY))` you specify that you want to retrieve the percentage of time the logical processors were busy. This value is reported as LPAR BUSY TIME PERC in the CPU Activity report for each reported interval.

The Spreadsheet Reporter's resource-oriented concept

The Spreadsheet Reporter uses a resource-oriented concept. In the main dialog, the resources are grouped into *remote* and *local* resources.

- **Remote:**
 - SMF Dump Data
 - Report Listings
 - Overview Records
- **Local:**
 - Report Listings
 - Overview Records
 - Working Sets
 - Spreadsheets

XML Reports are available as local or remote Report Listings.

Remote resources are located on the host system, local resources are located on your workstation. All actions that you perform with resources are described in "How to work with Resources" on page 304.

With the Spreadsheet Reporter, it is possible to convert Postprocessor output (Report Listings or Overview Records) to a data format that you can feed into spreadsheet macros for graphical presentation.

1. You define a host system to the Spreadsheet Reporter to enable the data transfer between the host and the workstation.

2. On your workstation, you create a remote resource of type **SMF Dump Data**. This designates an SMF data set on the host.
3. From this resource, you create a **Working Set** on your workstation. You can complete this step with one single action, because the Spreadsheet Reporter automatically performs the involved data preparation tasks as shown in Figure 51:
 - **Collect:** To create Report Listings or Overview Records, a Postprocessor job is generated and submitted on the host system.
 - **Download:** This remote Postprocessor output is downloaded to your workstation.
 - **Generate:** To generate a Working Set from the downloaded Report Listings or Overview Records, the extraction and conversion of the relevant data is performed.




Figure 51. From SMF Dump Data to a Spreadsheet presentation

4. You select a **Spreadsheet** which you feed with the created Working Set. Thus you receive a graphical display of the performance data captured in the original SMF Dump Data.

You can also create a Working Set step by step: for example, you can initiate the Postprocessor data collection and the download at one time and generate the Working Set later.

Note:

1. The RMF Spreadsheet Reporter also provides a collection of procedures that allow you to generate Working Sets in batch mode without any GUI interaction. For more information refer to “How to create Working Sets in batch mode” on page 323.
2. If you want to process Working Sets created by Spreadsheet Reporter versions prior to 5.1.0, you can migrate these Working Sets to the current format. For this purpose, the Spreadsheet Reporter provides appropriate procedures that are described in “How to migrate Working Sets” on page 326.
3. The RMF Spreadsheet Reporter also provides a set of reports in XML format. You can view these XML reports in a web browser.

Installing the Spreadsheet Reporter

The Spreadsheet Reporter is part of the RMF product. The deliverable includes:

- the Spreadsheet Reporter application files
- Spreadsheet Macros for Microsoft Excel
- a sample RMF Report Listing

Note: Install the Spreadsheet Reporter with each new release of RMF to cover all changes in Postprocessor reports.

Prerequisites

Hardware Requirements

- A workstation with at least 150 MB of free disk space.

Software Requirements

- Operating System
 - Windows XP / Windows Vista / Windows 7
- Spreadsheet Program

In order to use the spreadsheet macros shipped with the Spreadsheet Reporter, one of the following products is required:

- Microsoft Excel (Office XP with Excel 2003/ Excel 2007)

- Level of access

To install the Spreadsheet Reporter on a workstation running Windows Vista, you need administrator rights for this workstation. To install the Spreadsheet Reporter on a workstation with an earlier Windows version, you need not be the administrator. However, if you are not the administrator, you cannot install the Spreadsheet Reporter into the default installation directory C:\Program Files\RMF\IBM RMF Spreadsheet Reporter, but you must install it into a different directory.

On a multiuser operating system, each user who wants to work with the Spreadsheet Reporter must install it separately.

Installation steps

1. The code and installation utility of the Spreadsheet Reporter is available as member ERB9R2SW of the host distribution library SERBPWSV. Download this member as binary file **erb9r2sw.msi**.
2. Install the MSI package using the Windows Installer, either by double-clicking on the MSI package file or by issuing the command:

```
msiexec /package erb9r2sw.msi
```

The Windows Installer guides you through the installation.
3. Specify the directory where to install the Spreadsheet Reporter. The default is:

C:\Program Files\RMF\RMF Spreadsheet Reporter

4. Specify the resource directory, that is the drive and folder where the Spreadsheet Reporter will place the resources, for example Report Listings, macros or Working Sets. You may accept the default or specify another folder.

Note: The specified folder cannot be moved to another drive later. Ensure that you have sufficient disk space on the target drive (see “Prerequisites” on page 298).

The Spreadsheet Reporter is installed into program group *IBM RMF Performance Management*.

If you want to process Working Sets created by Spreadsheet Reporter versions prior to z/OS V1R11, you must migrate these Working Sets (refer to “How to migrate Working Sets” on page 326).

Note: If you have installed a Spreadsheet Reporter version earlier than z/OS V1R11, it is recommended to uninstall it before installing the MSI package.


How to use the RMF Spreadsheet Reporter

The Spreadsheet Reporter uses a resource oriented concept. Starting from the main dialog, you maintain the **systems** from which you retrieve the performance measurements, and the **resources**.

A typical usage scenario for the Spreadsheet Reporter is to run a Postprocessor job on the host, download the resulting Postprocessor data set (Report Listing or Overview Records) to the workstation, and convert it into a Working Set. The Working Set now contains your individual performance measurement data. Now you start a spreadsheet macro for graphical performance analysis and you feed the macro with the generated selected Working Set. This is the fast path for graphical presentation of RMF performance data.

As you can perform most of the transitions between the resource types in all variations, you may use the Spreadsheet Reporter also as a remote Postprocessor execution and download utility. For example, you may initiate the execution of a Postprocessor job from your workstation and store the Report Listing output on the host without downloading. Using **File ---> Transfer**, you can download RMF Postprocessor data sets later.

If you use the RMF Spreadsheet Reporter for the first time, you need to define one or more systems from which you want to retrieve performance measurements. If no system is defined, the Spreadsheet Reporter issues the following message:



If you do not want to define a system now, you can still continue to work with local resources. How to define systems is described in “How to work with Systems” on page 313.

Spreadsheet Reporter resources

The Spreadsheet Reporter resources introduced in “The Spreadsheet Reporter's resource-oriented concept” on page 296 are described in the following sections.

SMF Dump Data

RMF writes performance measurement data into SMF records, from which the Postprocessor extracts the reports or metrics by producing Report Listings or Overview Records. **SMF data** from two sources can be used as input to the Postprocessor:

- SMF records from SMF dump data sets: SMF dump data is usually stored in generation data groups (GDGs). With the Spreadsheet Reporter, these data sets can be defined as remote resources of type **SMF Dump Data**.
- SMF records from the *RMF Sysplex Data Server's SMF buffer*: This is a wrap-around buffer, used by RMF to store copies of SMF records written to the SMF dump data sets.

When you create Working Sets, Report Listings, XML Reports, or Overview Records, you determine which type of SMF data the Postprocessor should use: If you do not select any **SMF Dump Data** resource, then the Postprocessor automatically extracts the requested data from the *RMF Sysplex Data Server's SMF buffer*.

Report Listings

A Report Listing is the result of a Postprocessor batch job submitted on the host system, with SMF data specified as input.

There are two methods to create a resource of type **Report Listing** from SMF data:

- Using **Create ---> Report Listing...** generates and submits a Postprocessor job. The Postprocessor collects the data from SMF and produces a Report Listing on the host (remote). You have the choice to additionally transfer it to your workstation (local). Local Report Listings are stored with file extension '.lis', for example: SYSF.D128.T135719.lis.

For the Postprocessor report types listed in Figure 64 on page 322, the Spreadsheet Reporter can generate Report Listings with a file extension .xml which you can view in a web browser. For more information see “How to specify Postprocessor report types” on page 319 and “How to view local Report Listings” on page 307.

- Using **Create ---> Working Set...** directly converts the SMF data into spreadsheet format, optionally storing the intermediate Postprocessor Report Listing as local and remote resources.

The default names for Report Listings consist of four parts:

- TSO high level qualifier for remote listings or system name for local listings
- prefix D + julian day
- prefix T + time in HHMMSS
- 'LISTING' for remote data sets or 'lis' or 'xml' for local copies

Examples:

IBMUSER.D203.T104615.LISTING (remote)
SYSF.D203.T104615.lis (local)

Overview Records

Overview Records are the result of a Postprocessor batch job submitted on the host system.

There are two methods to create a resource of type **Overview Records** from SMF data:

- Using **Create ---> Overview Record...** generates and submits a Postprocessor job. The Postprocessor collects the data from SMF and produces an Overview Record data set on the host (remote). You have the choice to additionally transfer it to your workstation (local). Local Overview Records are stored with file extension '.rec', for example: SYSF.D128.T135719.rec.
- Using **Create ---> Working Set...**, you can also create and store an intermediate data set containing Overview Records. In this case, you need to select option **Create Overview Records** in the *Options* dialog (see "How to specify processing options and report types" on page 317). This causes the Postprocessor to create a resource of type **Overview Records** instead of an Overview Report.

In both cases you need to attach a file with overview control statements to the current system (see "How to use overview control statements" on page 315).

The default names for data sets containing Overview Records consist of the following parts:

- TSO high level qualifier for remote data sets or system name for local copies
- prefix D + julian day
- prefix T + time in HHMMSS
- 'OVWREC' for remote data sets or 'rec' for local copies

Examples:

```
IBMUSER.D199.T131456.OVWREC (remote)
SYSF.D199.T131456.rec (local)
```

Working Sets

Working Sets are used as input to a spreadsheet macro. The RMF Spreadsheet Reporter can create a Working Set from:

- SMF Dump Data
- Report Listings (local and remote).
- Overview Records (local and remote).

When creating a Working Set directly from **SMF Dump Data**, you must be aware for which spreadsheet macro you want to use it, because certain spreadsheet macros require as input a Working Set derived from a Report Listing, while other spreadsheets require a Working Set derived from Overview Records. How to produce either type of a Working Set is described in "How to create Working Sets" on page 309.

The default names for Working Sets comprise the following parts:

- type indicator: 'Rpt' (created from a Report Listing) or 'Ovw' (created from Overview Records)
- system name
- prefix D + julian day
- prefix T + time in HHMMSS

Examples:

Rpt.SYSF.D203.T103840

Ovw.SYSA.D143.T144210

To use a Working Set for graphical analysis with a spreadsheet macro, you can invoke a **Spreadsheet** from the Spreadsheet Reporter dialog. The spreadsheet macro lets you then select a Working Set to be processed.

Spreadsheets

You use spreadsheet macros for the final presentation of the SMF data. The Spreadsheet Reporter extracts measurement data from SMF dump data and finally converts it into a Working Set which then contains your individual performance data. Just load and view your performance data by feeding the spreadsheet macros with a selected Working Set.

The Spreadsheet Reporter provides several sample spreadsheet macros to help you in viewing and analyzing performance data at a glance. Examples of available spreadsheets are **Workload Activity Trend Report** or **DASD Activity Report**.

Certain spreadsheet macros require as input a Working Set derived from a Report Listing (*Report Working Set*), while other spreadsheets require a Working Set derived from Overview Records (*Overview Working Set*).

"How to create Working Sets" on page 309 provides information about which spreadsheet macro requires which type of Working Set.

If you want to write your own spreadsheet macros, you can use the registry entry

HKEY_CURRENT_USER -> SOFTWARE -> VB and VBA Program Settings ->
RMF Spreadsheet Reporter -> JEnvironment -> path

to access the Spreadsheet Reporter working set directory, for example,

C:\Documents and Settings\Administrator\Application Data\RMF\
RMF Spreadsheet Reporter\WorkingSets

To view this registry entry, open the *Windows Registry Editor* by typing regedit into a command prompt.

In previous versions of the RMF Spreadsheet Reporter, the registry entry

HKEY_CURRENT_USER -> SOFTWARE -> VB and VBA Program Settings ->
RMF Spreadsheet Reporter -> Environment -> path

was used to store the working set directory.

Spreadsheet Reporter main dialog

After invoking the Spreadsheet Reporter using the program folder **IBM RMF Performance Management - RMF Spreadsheet Reporter**, you see the *main dialog* (Figure 52 on page 303).




Figure 52. Spreadsheet Reporter - main dialog

The main dialog consists of a split pane:

- the **navigation pane** (left hand side): Here you navigate to resources and systems that you want to manage.

Clicking on the **Resources** tab shows you all resource types. You can open a resource type folder containing the corresponding resources. Existing resources are then displayed on the **view pane** (right hand side).

The resources are organized as a tree containing *remote* and *local* resources.

Clicking on the **Systems** tab opens a folder **All Systems**. All defined z/OS host systems are displayed in the **view pane**.

If you want to access remote resources, you must first select the system where these resources reside. Otherwise, the list will be empty.

- the **view pane**: Here you see available resources or systems. You can select resources or systems and initiate actions (for example, create a Working Set from a Report Listing) or maintain them (for example, modify properties or add/delete systems or resources).

Example: In Figure 52, the folder for the remote resource type **SMF Dump Data** is opened. In the **view pane** you see a list of SMF data sets residing on system **SYSF**. These data sets have been defined to the Spreadsheet Reporter earlier with the *New SMF Dump Data* dialog (see “How to create SMF Dump Data” on page 305).

Note:

1. The currently selected system (in our example, **SYSF**) is shown in the title bar of the main dialog.
2. While using the RMF Spreadsheet Reporter, you may always press F1 to get help.
3. To reach the context menu for a list of resources or systems, use **Alt+A** from free space.

4. To reach the context menu for a single list element, select this element and use **Alt+A**.

Menu bar

Each menu bar item offers a pull-down choice to trigger the RMF Spreadsheet Reporter actions. The most frequent and important actions can also be started from the icons in the tool bar. Positioning the cursor over an icon displays a fly-over text indicating the provided action.

The pull-down choices of all menu bar items are described in detail in the Spreadsheet Reporter's online help. Here is an overview of the Spreadsheet Reporter actions that you can trigger:

File offers actions to print Report Listings or to transfer remote resources to your workstation.

Define

offers dialogs to define:

- a new System
- new SMF Dump Data
- a new remote Report Listing
- a new remote Overview Record

View offers choices to display a Java™ or Windows look and feel of the Spreadsheet Reporter.

Settings

offers dialogs to specify:

- general processing options
- Postprocessor report types which you want to include in the Report Listings
- start and end times of the reporting periods
- duration periods.

Create offers actions to generate Report Listings, Overview Records or Working Sets. The **Create** actions are related to the currently selected resources, so that only actions valid in a given context are selectable.

Messages

offers entry points to access all types of messages resulting from running the remote RMF Postprocessor jobs.

How to work with Resources

Before you can work with **remote** resources, you need to select a system under the **Systems** tab in the **navigation pane**. The system name is then added to the title bar. It changes if you select another system or disappears if no system is selected at all.

You can work with **local** resources without an existing system definition.

To display available resources, click on the **Resources** tab in the **navigation pane** and then click on a resource type. The existing resources appear in the **view pane**. In the example from Figure 52 on page 303, the folder of resource type **SMF Dump Data** is opened and the **view pane** shows a list with the contained data sets.

If you right-click into the **view pane's** free space, a pop-up menu with action **New** appears, as shown in Figure 52 on page 303. With this action, you add a new resource. If you right-click a selected resource from the list, the pop-up menu appears with applicable actions for this resource, for example, **Rename** or **Delete**.

How to create SMF Dump Data

Creating resources of type **SMF Dump Data** consists of adding the name of a remote SMF data set to the **view pane**. This name is just the pointer to the SMF data set which remains located on the host.

To create an **SMF Dump Data** resource, perform the following steps starting from the main dialog:

1. Select a system.
2. Click the **Resources** tab in the **navigation pane** and select resource type **SMF Dump Data**.
3. Right-click in the free space of the **view pane**.
4. In the pop-up menu, choice **New** is enabled. Clicking on this choice opens the *New SMF Dump Data* dialog, where you enter the fully qualified name (without quotes) of an existing host data set.
5. Press the **OK** button to add the new name to the list.




Figure 53. New SMF Dump Data

How to create Report Listings

You have two possibilities to create a new Report Listing:

- Create a Report Listing from **SMF Dump Data** using the appropriate *Create...* dialog. This method is described in the following.
- Add a new list element to the **view pane** for an existing remote Report Listing using the appropriate *New...* dialog (see “How to add remote resources” on page 313).

To create a Report Listing, you have two choices:

- On the **navigation pane**, open resource type **SMF Dump Data** to select one or more remote SMF data sets as input to the *Create...* dialog.
- If you do not select an SMF data set, the Postprocessor accesses the RMF Sysplex Data Server's SMF buffer on the current system.

If you now open the **Create** menu, you see the item **Report Listing...** enabled. Clicking on this item opens the *Create Report Listing* dialog. With this dialog you can generate a Postprocessor job and start it on the remote system. The created Report Listing contains the reports selected in the *Options* dialog (see “How to specify Postprocessor report types” on page 319) and covers the time range specified in the *Intervals* dialog (see Figure 65 on page 323).

In the **RMF Postprocessor Data Sets** group box from the *Create Report Listing* dialog, you specify whether to download the generated Report Listing:

- To download the Report Listing to the workstation, accept or specify a local data set name in the **Local** entry field as shown in the example dialog from Figure 54.
- If you delete the name in the **Local** entry field, the Report Listing is created only on the host system with the data set name specified in the **Remote** entry field.




Figure 54. Create Report Listing

In the example from Figure 54, with multiple SMF input data sets selected and an optional local file name specified, the **Run** button starts the following processing:

1. The data set with the name **IBMUSER.D254.T162349.LISTING** shown in the **Remote** entry field is allocated on the host system. This suggested name indicates the current date and time (see also “Report Listings” on page 300 for naming conventions). You may accept or overwrite this name.
2. A job that uses the specified SMF data sets as input is created and sent to the host. For the job creation, the Spreadsheet Reporter uses a job skeleton that is stored in the installation directory in subdirectory \Connect\rmfpp1.jcl. If you need additional parameters or records (for example, a //STEPLIB record), you can modify the skeleton according to your requirements.

Note: The Spreadsheet Reporter uses the rmfpp1.jcl job skeleton also for the creation of Overview Records, Working Sets, and XML reports.

The generated reports are stored as a Report Listing in the allocated remote data set from step 1.

3. This remote data set is transferred to your workstation with the file name suggested in the **Local** entry field (which you may overtype). Thus, you create a new Report Listing both as remote and local resource. You can use the local resource to create a Working Set later.

Use **File ---> Transfer ...** to create a local Report Listing from a remote Report Listing (that is, downloading an RMF Postprocessor data set containing a Report Listing).

If you receive a Report Listing, for example by mail attachment, you may want to import this listing locally to the RMF Spreadsheet Reporter for further processing. You must store this data set with file extension *.lis* or *.xml* in subdirectory '*...\RmfListings*' of that directory that you specified during installation as the resource directory. The default is:

```
C:\Documents and Settings\Administrator\Application Data\RMF\  
RMF Spreadsheet Reporter\RmfListings
```

Then you can convert this Report Listing to a Working Set to use it in a spreadsheet macro, or if it is an XML type listing, view it in a web browser for further analysis.

How to create XML reports: XML reports are a special type of Report Listings. If you want to receive XML reports for viewing in a web browser, you need to select option *Use XML Report Format* in the **Options** dialog (Figure 62 on page 318). Then the Spreadsheet Reporter generates a local Report Listing with extension *.xml* for those reports that you selected from the available Postprocessor report types, which are presented in Figure 64 on page 322). The generated local XML Report Listing contains Postprocessor output in XML format for the selected reports. You can view the reports in a web browser as described in "How to view local Report Listings." An appropriate XSLT style sheet which transforms the XML to HTML is provided with the Spreadsheet Reporter.

How to view local Report Listings: There are two ways to view local Report Listings. Either double-click on a local Report Listing in the **view pane** or select an entry, click the right mouse button and then select *View* from the pop-up menu:

- Report Listings with extension *.lis* are opened in a text view.
- Report Listings with extension *.xml* are opened in a web browser.

How to create Overview Records

You have two possibilities to create a new Overview Record:

- Create an Overview Record from **SMF Dump Data** using the appropriate *Create ...* dialog. This method is described in the following.

Note: You need to attach a file with overview control statements to the system on which you want to generate Overview Records from an SMF input data set. For information about this process, see "How to use overview control statements" on page 315. Also you must activate option **Create Overview Records** from the **Options** dialog (see "How to specify processing options and report types" on page 317).

- Add a new list element to the **view pane** for an existing remote Overview Record using the appropriate *New ...* dialog (see "How to add remote resources" on page 313).

To create an Overview Record, you have two choices:

- On the **view pane**, open resource type **SMF Dump Data** to select one or more remote SMF data sets as input to the *Create...* dialog.
- If you do not select an SMF data set, the Postprocessor accesses the RMF Sysplex Data Server's SMF buffer on the current system.

The created Overview Records cover the time range specified in the *Intervals* dialog (see Figure 65 on page 323).

If you now open the **Create** menu, you see the item **Overview Record...** enabled. Clicking on this item opens the *Create Overview Record* dialog. With this dialog you generate a Postprocessor job and start it on the remote system. The created Overview Records are built according to the overview control statements contained in the attached file.

In the **RMF Postprocessor Data Sets** group box from the *Create Overview Records* dialog, you specify whether to download the generated Overview Records:

- To download the Overview Records to the workstation, accept or specify a local data set name in the **Local** entry field as shown in the example dialog from Figure 55.
- If you delete the name in the **Local** entry field, the Overview Record is created only on the host system with the data set name specified in the **Remote** entry field.




Figure 55. Create Overview Record

In the example from Figure 55, with multiple SMF input data sets selected and an optional local file name specified, the **Run** button starts the following processing:

1. The data set with the name **IBMUSER.D254.T165807.OVWREC** shown in the **Remote** entry field is allocated on the host system. This suggested name indicates the current date and time (see also "Overview Records" on page 301 for naming conventions). You may accept or overtype this name.
2. A job that uses the specified SMF data sets as input is created and sent to the host. The job extracts the metrics as specified in the attached file with overview control statements. The extracted overview data is stored as Overview Records in the allocated remote data set from step 1. If you did not attach a file with overview control statements to the current system, you receive a message stating that Overview Records cannot be created.
3. This remote data set is transferred to your workstation with the file name that you specified in the **Local** entry field (which you may overtype). Thus, you create a new Overview Record file both as remote and local resource. You can use the local resource to create a Working Set later.

Use **File ---> Transfer ...** to create a local Overview Record from a remote Overview Record (that is, download an RMF Postprocessor data set containing Overview Records).

If you receive Overview Records, for example by mail attachment, you may want to import this data set locally to the RMF Spreadsheet Reporter for further processing. You must store this data set with file extension '.rec' in subdirectory '...\RmfRecords' of that directory that you specified during installation as the resource directory. The default is:

```
C:\Documents and Settings\Administrator\Application Data\RMF\
  RMF Spreadsheet Reporter\RmfRecords
```

Then you can convert this Overview Record to a Working Set to use it in a spreadsheet macro for further analysis.

How to create Working Sets

You can create a new Working Set from resources of the following types:

- SMF Dump Data
- Report Listing (local and remote, however, not from XML type listings)
- Overview Records (local and remote).

To generate a Working Set from one of these resources, first select the resource type to display the list of available resources. Then, from the list in the **view pane**, select the resource that you want to process.

If you now open the **Create** menu, you see that the choice **Working Set...** is enabled. Click that choice to open the **Create Working Set** dialog. The layout of this dialog depends on the resource type that you want to process.

The most comprehensive dialog appears if you want to create a new Working Set from **SMF Dump Data**. Then, the dialog reflects all processing steps: run the Postprocessor job on the host, transfer the Report Listing or Overview Record to the workstation, and finally generate the Working Set. During this process, the Spreadsheet Reporter creates the corresponding intermediate Postprocessor data sets. The dialog consists of three group boxes:

- the **SMF Dump Input Data** box reflects the **Collect** step: it shows the selected SMF input data sets. This box is empty if you did not select an SMF input data set, but decided to collect data from the RMF Sysplex Data Server's SMF buffer. In either case, a Postprocessor job is created and submitted that generates Report Listings or Overview Records.

- the **RMF Postprocessor Data Sets** box reflects the **Download** step: it proposes a name for the remote Report Listing or Overview Record and for its local copy after download. You may overtype both names and specify your own ones.
- the **New Working Set** group box reflects the **Generate** step: it proposes a name for the Working Set and shows the location where the Working Set is to be stored. You may overtype the proposed Working Set name. The Spreadsheet Reporter keeps name and location consistent. Therefore you cannot edit the **Location** entry field.

Create Working Set

SMF Dump Input Data

- RMF.SMFDATA.SYSF(0)
- RMF.SMFDATA.SYSF(-1)
- RMF.SMFDATA.SYSF(-2)

RMF Postprocessor Data Sets

Remote

IBMUUSER.D278.T145447.LISTING

Local

SYSF.D278.T145447.lis

New Working Set

Name

Ovw.SYSF.D278.T145447

Location

on Data\RMF\RMF Spreadsheet Reporter\WorkingSets\Ovw.SYSF.D278.T145447

Run Cancel

Figure 56. Create Working Set

If you want to create a Working Set from a remote Report Listing or a remote Overview Record, then the dialog has a simplified layout:

- The **SMF Dump Input Data** box is missing, because the **Collect** step has already been performed.
- The **RMF Postprocessor Data Sets** box and the **New Working Set** box have the same functions as described previously.

If you want to create a Working Set from SMF data, you must consider which spreadsheet macro you want to use for further analysis of this Working Set, because certain spreadsheet macros require as input a Working Set derived from a Report Listing (*Report Working Set*), while other spreadsheets require a Working Set derived from Overview Records (*Overview Working Set*).

The following spreadsheet macros require a Report Working Set as input:

- Cache Subsystem Report
- Coupling Facility Trend Report
- DASD Activity Report
- I/O Subsystem Report
- LPAR Trend Report
- Open RMF Report Spreadsheets
- Summary Report
- Tape Mount Report
- Workload Activity Trend Report
- XCF Trend Report

The following spreadsheet macros require an Overview Working Set as input:

- Cache Subsystem Overview Report
- Channel Overview Report
- Device Overview Report
- Open RMF Overview Spreadsheets
- LPAR Overview Report
- System Overview Report
- Workload Overview Report

How to create Spreadsheets

On the **navigation pane**, you select the local resource type **Spreadsheets** to display a list of available spreadsheet macros in the **view pane**. Now double-click on a spreadsheet from that list.




Figure 57. Available Excel Spreadsheet Macros

By double-clicking a spreadsheet, or by selecting *Start* from its context menu, you can open a spreadsheet that presents a graphical view of a report or Overview Record, for example, the **LPAR Overview Report** spreadsheet.

Once you have opened the spreadsheet, you are ready to feed in the data from a Working Set. A dialog lets you select the Working Set and the reporting intervals contained in this Working Set.

For information on how to use the spreadsheet dialogs, either refer to the spreadsheets' online helps or to "Spreadsheet usage examples" on page 336.

How to add remote resources

Remote Report Listings and Overview Records are automatically created as intermediate resources during Working Set creation.

Additionally, you may want to define existing remote Report Listings/XML Reports and Overview Records to the Spreadsheet Reporter. Use the same procedure as described for the creation of SMF Dump Data (see "How to create SMF Dump Data" on page 305). Either the *New Report Listing* or the *New Overview Record* dialog is invoked where you enter the fully qualified name (without quotes) of an existing host data set.

A second method to add remote resources is to select *Define ---> Remote Report Listing...* or *Define ---> Remote Overview Record...*

How to work with Systems

To manage the systems, click the **Systems** tab in the main dialog to display all defined systems in the **view pane**:





Figure 58. List of Systems in the view pane

How to create Systems

Right-click in the free space of the **view pane** with the systems currently defined. From the context menu, click **New** to start the *Define new System* dialog. After you specified all required values for the parameters, click **OK** to add the new

system to the current list.



The screenshot shows a Windows-style dialog box titled "Define new System". It contains the following fields and controls:

- System ID:
- Hostname:
- Dataset HLQ:
- User ID:
- Password:
- Account:
- Jobclass:
- Msgclass:
- FTP Security: - FTP Mode: - OVW button (with folder icon) and a text field containing "D:\RMFSR\Text\sovgoal.txt"
- Ok and Cancel buttons at the bottom.

Figure 59. Define new System

The **OVW** button opens a standard file dialog. Browse for a file containing overview control statements that you can attach to the current system. The file name is then displayed in the entry field. For information about the purpose of this file, read “How to use overview control statements” on page 315.

Note:

1. If you do not know your systems's hostname, you may retrieve this hostname and the system's TCP/IP address with the TSO command *hometest*.
2. The Spreadsheet Reporter uses an FTP connection to the host, with default port 21. To change this port, for example because of security restrictions, specify the port which you want to use with the hostname, appended with a ':'. For example, to use port 2001, specify your hostname using the following pattern: boesysf.boeblingen.de.ibm.com:2001
3. Some of the parameters that you define for a system are used for Postprocessor job creation. Another source for the job creation is a job skeleton that is stored in the installation directory in subdirectory \Connect\rmfpp1.jcl. If you need additional parameters or records (for example, a //STEPLIB record), you can modify the skeleton according to your requirements.

How to maintain Systems

If you want to maintain a system, right-click this system to display its context menu to see the following choices enabled:

- Rename
- Delete
- Properties

To view or change a system's properties, select **Properties** to display the *System Properties* dialog. This dialog is identical to the *Define new System* dialog (see Figure 59 on page 314), except that you cannot change the **System ID**.

How to use overview control statements

The dialogs *Define new System* and *System Properties* contain the **OVW** push button used to attach a file containing overview control statements to the current system.

Attaching such a file to a system is optional. If this file is attached, the contained overview control statements are processed in the following ways:

- You use action **Create ---> Report Listing...** from **SMF Dump Data**. Then the overview control statements are used to create a Report Listing containing a tabular Overview Report (suitable for reading, but not suitable for processing with a spreadsheet). This is the recommended way to create an Overview Report.




Figure 60. Processing a file containing overview control statements (1)

- You use action **Create ---> Working Set...** from **SMF Dump Data**. Then you should select general option **Create Overview Records** in the *Options* dialog. The overview control statements are used to create an Overview Records resource which is needed as input to an Overview Working Set. If option **Create Overview Records** is **not** selected, the Spreadsheet Reporter creates a Report Listing containing a tabular Overview Report according to the overview control statements. The Spreadsheet Reporter issues a message to indicate, that the Working Set could not be created from this Report Listing.




Figure 61. Processing a file containing overview control statements (2)

- You use action **Create ---> Overview Record...** from **SMF Dump Data**. This action is **not** possible without an attached file with overview control statements.

If **no** file containing overview control statements is attached to the system, and you want to create a Working Set in one step from **SMF Dump Data**, then the Spreadsheet Reporter will produce a Report Working Set containing the Postprocessor report types selected in the *Options* dialog.

If you want to create a Working Set from Report Listings or Overview Records, then the existence of an attached file is not checked, because the decision whether the Working Set should contain overview or report information has already been made.

Note: If you want to work with multiple sets of overview control statements for the same system, you can define multiple copies of the same system with a **different System ID**, a different file attachment, but with the **same Hostname**. This enables you to work with fixed attachments instead of changing the system properties all the time.

You can use the following spreadsheet macros to create a file containing overview control statements:

- Use the macro *Create Overview Control Statements* to produce overview control statements for all spreadsheet macros requiring an Overview Working Set.
- Use the *DASD Activity Report* macro to produce overview control statements for the *Device Overview Report* macro.

- Use the *Cache Subsystem Report* macro to produce overview control statements for the *Cache Subsystem Overview Report* macro.
- Use the *Workload Activity Trend Report* macro to produce overview control statements for the *Workload Overview Report* macro.

The advantage of the macros *Cache Subsystem Report* or *DASD Activity Report* is that they derive a list of identifiers from the loaded Working Set, for example, volume names when using the *DASD Activity Report*. You can select identifiers from this list, for example, the volumes with the highest activity.

How to specify settings for the RMF Spreadsheet Reporter

With the RMF Spreadsheet Reporter you define the following settings:

- processing options
- report types
- reporting periods
- duration intervals

How to specify processing options and report types

Select **Settings ---> Options** to invoke the *Options* dialog. Clicking on either the **General** or **Reports** tab, you start the following tasks:

- Specify general processing options
- Specify RMF Postprocessor report types

Note: The specified options and report types are active for all defined systems.

How to specify general processing options

Settings ---> Options ---> General starts the dialog shown in Figure 62 on page 318, where you select the options that you want to be active:




Figure 62. Processing options

The options have the following meaning:

Dialog option	Description
Create Overview Records	Select this option to create Overview Records using the overview control statements contained in a file that is attached to the current system. If a file with overview control statement is attached to the current system, but this option is not selected, then a readable Postprocessor Overview Report is generated according to the overview control statements contained in this file. However, you cannot process this report with spreadsheets. If no file containing overview control statements is attached, then this option is ignored and the Working Set is generated with the selected RMF Postprocessor report types.
Delete Postprocessor Datasets after Download	Remote Report Listings or Overview Records on the host are deleted after a successful download to the workstation.
Ignore specified Duration Period	No DINTV control statement is generated from the interval options so that no duration reports will be created (see "How to specify reporting periods and duration intervals" on page 322).
Ignore specified Interval Time	No RTOD control statement is generated from the interval options. The default from 00:00 to 24:00 is used (see "How to specify reporting periods and duration intervals" on page 322).
Save Password with System Profile	The password that you specified for a system in dialogs <i>Define new System</i> or <i>System Properties</i> is saved but not encrypted. Otherwise you are prompted for the password for all actions that require a host logon.

Dialog option	Description
Scratch Overview Records after Conversion	The Spreadsheet Reporter deletes local Overview Records (*.rec files) after Working Set generation.
Scratch Report Listings after Conversion	The Spreadsheet Reporter deletes local Report Listings (*.lis files) after Working Set generation.
Scratch extracted OVW Files after Conversion	The Spreadsheet Reporter deletes local .OVW files after generating Working Sets from Overview Records.
Scratch extracted RPT Files after Conversion	The Spreadsheet Reporter deletes local .RPT files after generating Working Sets from Report Listings.
Sort SMF Datasets	You can specify whether the SMF data should be sorted. To ensure correct reports, the records in an SMF data set must be sorted by interval start date and interval start time.
Use XML Report Format	The Spreadsheet Reporter generates local Report Listings with extension <i>.xml</i> for certain Postprocessor report types. You can view these reports in a web browser.

How to specify Postprocessor report types

The Spreadsheet Reporter supports a subset of Postprocessor reports for use in a spreadsheet (also compare with Figure 63 on page 321).

Table 44. Supported Reports for spreadsheets

Supported Reports	Postprocessor Options
Cache Subsystem Activity Report	REPORTS(CACHE(SUBSYS))
Channel Path Activity	REPORTS(CHAN)
Coupling Facility Activity	SYSRPTS(CF)
CPU Activity	REPORTS(CPU)
DASD Device Activity	REPORTS(DEVICE(DASD))
Partition Data	REPORTS(CPU)
I/O Queuing Activity	REPORTS(IOQ)
Summary Report	SUMMARY(INT)
Tape Device Activity	REPORTS(DEVICE(TAPE))
Workload Activity	SYSRPTS(WLMGL(SCPER,WGROUP,RCLASS,POLICY))
XCF Activity	REPORTS(XCF)

In addition, the Spreadsheet Reporter supports a selection of Postprocessor reports for viewing in a web browser (also compare with Figure 64 on page 322).

Table 45. Supported Reports for web browser display

Supported Reports	Postprocessor Options
Cache Subsystem Activity	REPORTS(CACHE(SUBSYS))
Channel Path Activity	REPORTS(CHAN)
CPU Activity	REPORTS(CPU)

Table 45. Supported Reports for web browser display (continued)

Supported Reports	Postprocessor Options
Crypto Hardware Activity	REPORTS(CRYPTO)
Enqueue Activity	REPORTS(ENQ)
Enterprise Disk Systems	REPORTS(ESS)
FICON Director Activity	REPORTS(FCD)
Hierarchical File System Statistics	REPORTS(HFS)
I/O Queuing Activity	REPORTS(IOQ)
OMVS Kernel Activity	REPORTS(OMVS)
Paging Activity	REPORTS(PAGING)
Partition Data	REPORTS(CPU)
Page Data Set Activity	REPORTS(PAGESP)
PCIE Activity	REPORTS(PCIE)
Serialization Delay	REPORTS(SDELAY)
Virtual Storage Activity	REPORTS(VSTOR)
Cross-System Coupling Facility Activity	REPORTS(XCF)

Settings ---> Options ---> Reports starts the dialog shown in Figure 63 on page 321 or Figure 64 on page 322, where you select Postprocessor report types supported by the Spreadsheet Reporter. A REPORTS or SYSRPTS control statement as shown in Table 44 on page 319 or Table 45 on page 319 is generated for each selected report type. At least one report type must be selected. The selected report types are considered by the RMF Spreadsheet Reporter in the following cases:

- you decide to create Report Listings,
- you decide to create a Working Set from SMF Dump Data without a file with overview control statements attached to the current system (see “How to use overview control statements” on page 315).




Figure 63. Report types available in text format

If you selected option *Use XML Report Format* in the **Options** dialog (Figure 62 on page 318), then **Settings ---> Options ---> Reports** starts the dialog shown in Figure 64 on page 322, where you select Postprocessor report types supported in XML format. A REPORTS control statement as shown in Table 45 on page 319 is generated for each checked report type. At least one report type must be selected.




Figure 64. Report types available in XML format

How to specify reporting periods and duration intervals

Settings ---> **Intervals...** opens the *Intervals* dialog. Use this dialog to specify the start and end time for data collection. The input from this dialog is converted to the corresponding DATE and RTOD control statements.

By means of the **Duration** sliders you can generate a DINTV control statement to produce duration reports.

For more information about DATE, RTOD and DINTV refer to Chapter 15, "Long-term reporting with the Postprocessor," on page 201.

Note: The specified intervals and duration periods are active for all defined systems.




Figure 65. Reporting periods and duration intervals

How to create Working Sets in batch mode

The RMF Spreadsheet Reporter provides a collection of procedures that allow you to generate Working Sets in batch mode without any GUI interaction. They are located in the installation directory. The default is:

C:\Program Files\RMF\RMF Spreadsheet Reporter

Jclgen.bat

The **Jclgen.bat** procedure generates the JCL for a job to run the Postprocessor on the host. It contains the variables listed in Table 46. Use an editor to modify their values according to your needs.

Table 46. Variables in Jclgen.bat

Variable	Meaning	Examples
in	specifies the input file for the JCL generation. The default file <code>rmfpp1.jcl</code> is a job skeleton delivered with the RMF Spreadsheet Reporter. In this skeleton, parameters specified in brackets (like <i><parameter_name></i>) are overwritten with the variable values specified in Jclgen.bat .	<code>set in=%workpath%\rmfpp1.jcl</code>
out	specifies the name of the output file holding the generated JCL.	<code>set out=%workpath%\rmfpp2.jcl</code>
acct	specifies the account ID, for example your department number.	<code>set acct=D3248</code>
class	specifies the job class.	<code>set class=A</code>
date	specifies the date of the reporting period.	<code>set date=DATE(01011990,12312050)</code>
time	specifies the time of the duration interval.	<code>set time=RTOD(0000,2400)</code>

Table 46. Variables in Jclgen.bat (continued)

Variable	Meaning	Examples
hlq	specifies the data set high level qualifier required for the allocation of data sets.	set hlq=IBMUSER set hlq=D3248.IBMUSER
user	specifies the TSO user ID of the user submitting the job.	set user=IBMUSER
ppdsn	specifies the Postprocessor output data set for the Report Listing or Overview Records.	set ppdsn=CPU.REPORT.LISTING If the high level qualifier is D3248.IBMUSER, then the output data set is D3248.IBMUSER.CPU.REPORT.LISTING
sysin	Before you run this procedure, decide whether you want to produce Overview Records or a Report Listing. This parameter specifies the required input file containing Postprocessor control statements. The RMF Spreadsheet Reporter offers two sample input files. You may take these as a template and edit them according to your needs. They are located in subdirectory ... \Work of the installation directory: • sysinovw.txt - as a sample input file for Overview Records • sysinrpt.txt - as a sample input file for Report Listings	set sysin=%workpath%\sysinrpt.txt
mfpinput	specifies the file containing the names of the required SMF data sets. Not relevant if Jclgen.bat is invoked with option <i>buffer</i> .	set mfpinput=%workpath%\mfpinput.txt

Invocation: jclgen option

where *option* can be one of the following:

sort sorts the SMF data sets. If you use this option, the *mfpinput* parameter in *Jclgen.bat* points to a file containing the names of the required SMF data sets.

nosort does not sort the SMF data sets. As with option *sort*, specify the names of the required SMF data sets with the *mfpinput* parameter.

buffer takes the SMF data from the RMF Sysplex Data Server's SMF buffer.

Example: To generate a Postprocessor job that takes the SMF data from the RMF Sysplex Data Server's SMF buffer, use the following command:

```
jclgen buffer
```

Collect.bat

The **Collect.bat** procedure performs the complete SMF data collection as well as the download to the workstation. That is, it submits the job on the host and downloads the resulting Postprocessor output (Report Listing or Overview Record) to the workstation. **Collect.bat** contains the same variables as **Jclgen.bat**, plus the following ones:

Table 47. Additional variables in *Collect.bat*

Variable	Meaning	Examples
log	specifies the file containing messages from the JES subsystem.	set log=%workpath%\jes.joblog
ppfile	specifies the file containing the Postprocessor output after the download to the workstation (either a Report Listing or Overview Records). This file is input either to CreateRptWset.bat (with parameter <i>listing</i>) or to CreateOvwWset.bat (with parameter <i>ovwrec</i>).	set ppfile=%workpath%\sysname.date.time.ppdata
msg	specifies the file containing messages from the Postprocessor.	set msg=%workpath%\rmfpp.msg
passive	specifies the FTP connection type used by the Spreadsheet Reporter. For a passive FTP connection, specify -p. The default is an active FTP connection.	set passive=-p
ssl	specifies the FTP security mode. To use an FTP with SSL encryption, specify -s. The default is an FTP without explicit security.	set ssl=-s

Invocation: collect hostname password type

where:

hostname

is the name of the host where you want to connect to

password

is the password for the TSO user specified with variable *user*

type specifies the type of Postprocessor output: use **-r** if you want to collect data for a Report Listing or **-o** for an Overview Record.

Example: To produce a Report Listing, use the following command:

```
collect myhost mypasswd -r
```

If FTP errors occurred during file transmission, the corresponding messages are written into the file ...\\Work\\ftp.log in the installation directory.

CreateRptWSet.bat

This procedure generates a Working Set from an existing local Report Listing:

Invocation: CreateRptWSet listing dir name

where:

listing is the path and filename of the Report Listing

dir is the Working Set directory

name is the Working Set name.

Example: To convert a Report Listing called *C:\DASD.lis* into a Working Set called *My Dasd Working Set*, and store the Working Set in directory *C:\WSets\DasdWS*, use the following command:

```
CreateRptWSet C:\DASD.lis C:\WSets\DasdWS "My Dasd Working Set"
```

CreateOvwWSet.bat

The **CreateOvwWSet.bat** procedure generates a Working Set from an existing Overview Record data set on the workstation.

Invocation: `CreateOvwWSet ovwrec dir name`

where:

ovwrec

is the path and filename of the Overview Record file

dir

is the Working Set directory

name

is the Working Set name in double quotes.

Example: To convert a file with Overview Records called *C:\DASD.rec* into a Working Set called *My Dasd Working Set*, and store the Working Set in directory *C:\WSets\DasdWS*, use the following command:

```
CreateOvwWSet C:\DASD.rec C:\WSets\DasdWS "My Dasd Working Set"
```

How to migrate Working Sets

If you want to process Working Sets created by Spreadsheet Reporter versions prior to z/OS V1R11, you must migrate these Working Sets.

For this purpose, the Spreadsheet Reporter provides the **MigrateWSetsToXLS.bat** procedure to perform the migration in batch mode. It is located in the installation directory. The following invocation example uses the default installation directory.

Invocation:

```
C:\Program Files\RMF\RMF Spreadsheet Reporter\MigrateWSetsToXLS
```

Using RMF spreadsheet macros

The previous sections describe the steps that you need to perform before you can feed a Working Set into a spreadsheet macro. The purpose of this section is to assist you in using spreadsheet macros to process converted Postprocessor reports and Overview Records.

You can start a spreadsheet from the Spreadsheet Reporter dialog by opening the appropriate spreadsheet macro.

Available RMF spreadsheet macros

Read a short description of the available spreadsheet macros in the following subtopics:

- "Macros for Report Working Sets" on page 327
- "Macros for Overview Working Sets" on page 328

Macros for Report Working Sets

Table 48 lists the available macros and specifies which RMF reports are processed by them. These reports must be contained in the Working Sets that you select as input for a macro. If you want to obtain the required reports via JCL, use the Postprocessor options listed in Table 44 on page 319.

Table 48. Macros Based on Reports

Macro	Excel
Open RMF Report Spreadsheets This macro displays reports from a Report Working Set.	RMFR9OPN
Filter DASD or Cache Reports This macro filters devices from large DASD Activity and Cache Activity reports. Use this macro to focus on important, frequently used DASDs or cache subsystems. You can also use this macro to reduce the amount of data if you receive a message that the data exceeds the limit that can be processed by the spreadsheet macro.	DASDCONV
Summary Report This macro processes a Summary report and creates analysis summaries and graphics from its data.	RMFN9SUM
DASD Activity Report This macro analyzes a DASD Activity report and provides summaries for the most frequently used LCUs and DASDs in your installation.	RMFR9DAS (*)
Workload Activity Trend Report This macro provides performance reports and analyzes your system's behavior. It also supports zAAP and zIIP workload analysis and workload projection.	RMFR9WLM
Coupling Facility Trend Report This macro provides reports about activities in your coupling facilities.	RMFR9CF
Cache Subsystem Report This macro provides reports about activities in your cache subsystems.	RMFR9CAC (*)
I/O Subsystem Report This macro analyzes DASD Activity reports from several systems and provides summaries for the most frequently used LCUs and DASDs in your sysplex.	RMFR9MDV
LPAR Trend Report This macro analyzes Partition Data reports and provides information about the active partitions in your PR/SM environment.	RMFR9LP
Tape Mount Report This macro displays the tape mounts and the tape activities for one or several systems.	RMFR9TAP
XCF Trend Report This macro processes XCF Activity reports from one or multiple systems and provides XCF performance trend reports.	RMFR9XCF

All macros marked by (*) offer the capability to generate control statements to create Overview Records. You can attach a file containing overview control statements to a system as described in “How to use overview control statements” on page 315.

Macros for Overview Working Sets

Table 49. Macros Based on Overview Records

Macro	Excel
Open RMF Overview Spreadsheets This macro processes any Overview Working Set. It provides interval charts, day charts and trend charts based on the data of the selected Overview Working Set.	RMFOVW
LPAR Overview Report This macro creates a long-term overview about CPU consumption for selected partitions. It expects Overview Records as described in “LPAR Overview Report” on page 329.	RMFX9CPC
System Overview Report You can create a summary for one week, by a specified shift, for each hour and every day contained in the data. This allows you to examine data for one week at a glance. The macro expects Overview Records as described in “System Overview Report” on page 330.	RMFY9OVW
Workload Overview Report This macro creates summaries and graphics for a set of selected service classes and workloads of your installation. It expects Overview Records as described in “Workload Overview Report” on page 331. Note: This macro cannot process more than 27 workloads.	RMFY9WKL
Device Overview Report This macro creates a trend report for selected devices of your installation. It expects Overview Records as described in “Device Overview Report” on page 333. Note: This macro cannot process more than 30 devices.	RMFX9DEV
Cache Subsystem Overview Report This macro creates a trend report for selected cache subsystems of your installation. It expects Overview Records as described in “Cache Subsystem Overview Report” on page 334. Note: This macro cannot process more than 18 control units or devices.	RMFX9CAC
Channel Overview Report This macro creates a channel report for selected channels of your installation. It expects Overview Records as described in “Channel Overview Report” on page 335.	RMFX9CHN
Create Overview Control Statements This macro creates OVW and EXCEPT statements that can be used to generate data for the above described macros for Overview Working Sets.	RMFX9MAK

Behavior of RMF spreadsheet macros

All macros, both for Lotus and Excel, have a similar behavior. They provide **dialogs** to ask for your input. Those macros that process data from Working Sets provide dialogs to specify the input Working Set, the system ID and the reporting periods.

The macros are organized by worksheets (or tabs) which contain various types of information. For example, the Summary Report spreadsheet consists of the following tabs:

Main This is the first sheet in the macro.

It contains buttons to select a report and to create a final report. When you create a final report, this sheet is removed.

Info This is always the second sheet in the macro.

Its main purpose is to provide a very brief summary of the data, and it is used as main sheet for the final report.

Summaries

These are tabular reports which summarize and analyze the processed RMF data.

Charts These are the following sheets: *SumChart*, *DayChart*, *JobChart*. They contain graphics based on the RMF data used to create the report. In many cases, they also contain buttons to modify the view on the data.

Helps Some spreadsheet macros contain help information about how to use them and how to interpret their contents.

About The last sheet for general usage in the spreadsheet. It contains the spreadsheet version and ways how you can contact the RMF team.

Others

All following sheets contain range names, dialog definitions and macro definitions. These sheets are only of interest if you want to write your own spreadsheet macros.

Sets of overview control statements

For nearly all fields in Postprocessor reports you can manually define overview conditions using overview control statements. This gives you an enormous flexibility to define your own customized reports.

A second possibility to create overview control statements is to use the appropriate spreadsheet macros. For more information refer to “How to use overview control statements” on page 315.

As a third source, this section presents some examples of meaningful sets of overview control statements. You can copy and edit the sets according to your needs and use it with the related overview spreadsheet macros.

The difference between the System Overview Report and the other reports is that the first one is based on a fixed set of control statements while the other reports require customized statements containing information about your favorite service classes, DASD devices, or cache controllers, for example.

LPAR Overview Report

The purpose of this report is to observe selected LPARs of interest in your installation. For example, you can monitor an LPAR's defined capacity limit in MSUs per hour, its actual CPU consumption, or the percentage of time when WLM capped the partition.

The template consists of a variable part which must be repeated for each LPAR included in the Overview Record and a fixed part with three statements for the LPAR *PHYSICAL*.

Template:

```
----- Variable Part, replace vv with
----- the LPAR name
OVW(DEFvv(WDEFL(vv))) statements for LPAR
OVW(ACTvv(WACTL(vv)))
OVW(MINvv(WMINL(vv)))
OVW(MAXvv(WMAXL(vv)))
OVW(NLDvv(NLDEFL(vv)))
OVW(NLAvv(NLACTL(vv)))
OVW(LBUvv(LBUSYL(vv)))
OVW(PBUvv(PBUSYL(vv)))
OVW(LDMvv(LDEFMSU(vv)))
OVW(LDAvv(LACTMSU(vv)))
OVW(CAPvv(WCAPPER(vv)))
-----
OVW(NLDPHY(NLDEFL(PHYSICAL))) This part is fixed and must be added
OVW(LBUPHY(LBUSYL(PHYSICAL))) at the end of the report
OVW(PBUPHY(PBUSYL(PHYSICAL)))
```

If you want to use the template to monitor the LPARs SYSA and SYSF, you can create the following overview control statements:

Example:

```
OVW(DEFSYSYSA(WDEFL(SYSA)))
OVW(ACTSYSYSA(WACTL(SYSA)))
OVW(MINSYSYSA(WMINL(SYSA)))
OVW(MAXSYSYSA(WMAXL(SYSA)))
OVW(NLDSYSYSA(NLDEFL(SYSA)))
OVW(NLASYSYSA(NLACTL(SYSA)))
OVW(LBUSYSYSA(LBUSYL(SYSA)))
OVW(PBUSYSYSA(PBUSYL(SYSA)))
OVW(LDMSYSYSA(LDEFMSU(SYSA)))
OVW(LDASYSYSA(LACTMSU(SYSA)))
OVW(CAPSYSYSA(WCAPPER(SYSA)))
OVW(DEFSYSYF(WDEFL(SYSF)))
OVW(ACTSYSYF(WACTL(SYSF)))
OVW(MINSYSYF(WMINL(SYSF)))
OVW(MAXSYSYF(WMAXL(SYSF)))
OVW(NLDSYSYF(NLDEFL(SYSF)))
OVW(NLASYSYF(NLACTL(SYSF)))
OVW(LBUSYSYF(LBUSYL(SYSF)))
OVW(PBUSYSYF(PBUSYL(SYSF)))
OVW(LDMSYSYF(LDEFMSU(SYSF)))
OVW(LDASYSYF(LACTMSU(SYSF)))
OVW(CAPSYSYF(WCAPPER(SYSF)))
OVW(NLDPHY(NLDEFL(PHYSICAL)))
OVW(LBUPHY(LBUSYL(PHYSICAL)))
OVW(PBUPHY(PBUSYL(PHYSICAL)))
```

The LPAR Overview Report macro processes converted Overview Records which have been created based on this control set.

System Overview Report

The report is based on data which can be used without any installation-dependent qualification. This means that you can use the described set of control statements directly without modification.

The idea of the report is to give a comprehensive overview of the system resources CPU and storage in context with the total workload activity. Therefore, the data is derived from SMF record types 70-1, 71, and 72-3. You can find most of the data in the CPU, Paging, and Workload Activity reports.

Control Statements	SMF Record	Description
OVW(CPUBUSY(CPUBSY))	70.1	CPU utilization data
OVW(MVSBUSY(MVSBYS))		
OVW(APPLPER(APPLPER(POLICY)))	72.3	V Workload utilization
OVW(NUMPROC(NUMPROC))	70.1	----
OVW(EXCP(EXCP(POLICY)))	72.3	Total EXCPs
OVW(EXCPRT(EXCPRT(POLICY)))		EXCP rate
OVW(OCPU1(OCPU1))	70.1	CPU contention
OVW(OCPU2(OCPU2))		V
OVW(OCPU3(OCPU3))		----
OVW(INREADY(AVGIRDY))	70.1	SRM queues
OVW(OUTREADY(AVGORDY))		
OVW(OUTWAIT(AVGUOWT))		
OVW(LOGRDY(AVGULRDY))		
OVW(LOGWAIT(AVGULWT))		V
OVW(INUSER(AVGUIN))		----
OVW(AVGSTC(AVGSTC))		Average number of
OVW(AVGBATCH(AVGBATCH))		address spaces
OVW(AVGTSO(AVGTSO))		by class
OVW(AVGAPPC(AVGASCH))		V
OVW(AVGOMVS(AVGOMVS))		----
OVW(STOTOT(STOTOT(POLICY)))	72.3	Workload storage
OVW(STOCEN(STOCEN(POLICY)))		and paging
OVW(STOEXP(STOEXP(POLICY)))		for all
OVW(SINGLE(SINGLE(POLICY)))		system workloads
OVW(BLOCK(BLOCK(POLICY)))		
OVW(EXPSNGL(EXPSNGL(POLICY)))		V
OVW(EXPBLK(EXPBLK(POLICY)))		----
OVW(TPAGRT(TPAGRT))	71	System paging
OVW(FAULTS(PAGERT))		
OVW(DEMAND(DPAGRT))		
OVW(PGMOVERT(PGMVRT))		
OVW(TOEXPAND(PTES))		
OVW(MIGTOAUX(ESMR))		
OVW(AVGHUIC(AVGHUIC))		V
OVW(SWAPS(SWART))		----
OVW(AVGCSA(AVGCSAT))	71	Common storage utilization
OVW(AVGSQA(AVGSQA))		

You find the set of overview control statements for the report in file
 <resource_root>\Text\%OVWGOAL.TXT.

<resource_root> denotes the drive and folder that you specified during installation as the location where the Spreadsheet Reporter should place the Report Listings, macros and Working Sets.

The macro **System Overview Report** processes converted Overview Records which have been created based on this control set.

Workload Overview Report

The idea for the Workload Overview report is to monitor the CPU utilization for your most important service classes.

The template for the Overview Records contains the overview control statements for workload utilization (APPLPER) and EXCP rate (EXCPRT) for one service class. You have to define these statements for each selected service class. In addition, the total workload utilization and EXCP rate, as well as the CPU utilization and the

number of processors is added at the end of the template. The EXCP rate and the total numbers are used to calculate and distribute the uncaptured CPU time across the service classes.

Template:

```

-----
OVW(CPUclass(APPLPER(S.class.p)))      This part must be repeated for each service
OVW(EXPclass(EXCPRT(S.class.p)))      class; replace class by a service class name

OVW(MPLclass(TRANSAVG(S.class.p)))    These conditions can be added optionally
OVW(TOTclass(TRANS(S.class.p)))      and are exploited only with Excel
OVW(RTMclass(RTIME(S.class.p)))
OVW(EVLclass(EXVEL(S.class.p)))
-----
OVW(NUMPROC(NUMPROC))                This part is fixed and must be added
OVW(CPUBUSY(CPUBSY))                at the end of the report
OVW(APPLPER(APPLPER(POLICY)))
OVW(EXCPRT(EXCPRT(POLICY)))

```

Note: Please ensure that the OVW control-statement names (for example, CPUclass) have a maximal length of 8 characters.

The basic control set with the two conditions APPLPER and EXCPRT can be enhanced by up to four conditions for each service class. The corresponding Excel spreadsheet macros can process the Overview Records with or without the additional information. If you do not specify all information, some of the reports in the macro show that no information is available for display. Otherwise the functionality of the macros is not restricted.

In order to use the additional control statements you must adhere to the following rules:

- Whenever you add additional control statements, you must add them for **all** service classes contained in the control set.
- It is not necessary to add all four statements, you can add and combine them in the following three groups:

1. Information about transactions:

TRANSAVG

Average number of active transactions

TRANSMPL

Average number of resident transactions

2. You must always specify the following two exceptions in common:

TRANS Number of ended transactions per second

RTIME Transaction execution time

TRANSTOT

Total number of ended transactions

RTIMETOT

Total transaction time (execution and queue time)

3. Execution velocity

EXVEL Execution Velocity

This means that you can add them **all** for all service classes, or **any combination of the three groups** for all service classes. Only keep in mind that you must always add the **same conditions** for all service classes.

- When you add the additional definitions, it is required to add them in the sequence listed above. Also make sure that the prefix of the column header is always MPL, TRX, RTM and EVL.
- You can use the macro **Create Overview Control Statements** to generate the control statements for your favorite service classes and workloads. You can also specify report classes instead of service classes.

To specify a report class TSOR, use:

```
OVW(CPUTSOA1(APPLPER(R.TSOR.1)))
```

To specify a service class TSOS, use:

```
OVW(CPUTSOA1(APPLPER(S.TSOS.1)))
```

If you want to report service class *CICS* and the first and second period of service class *TSO*, you have to define the records as follows:

Example:

```
OVW(CPUCICS(APPLPER(S.CICS.1)))
OVW(EXPCICS(EXCPRT(S.CICS.1)))
OVW(MPLCICS(TRANSAVG(S.CICS.1)))
OVW(TOTCICS(TRANS(S.CICS.1)))
OVW(RTMCICS(RTIME(S.CICS.1)))
OVW(CPUTSO1(APPLPER(S.TSO.1)))
OVW(EXPTS01(EXCPRT(S.TSO.1)))
OVW(MPLTS01(TRANSAVG(S.TSO.1)))
OVW(TOTTS01(TRANS(S.TSO.1)))
OVW(RTMTS01(RTIME(S.TSO.1)))
OVW(CPUTSO2(APPLPER(S.TSO.2)))
OVW(EXPTS02(EXCPRT(S.TSO.2)))
OVW(MPLTS02(TRANSAVG(S.TSO.2)))
OVW(TOTTS02(TRANS(S.TSO.2)))
OVW(RTMTS02(RTIME(S.TSO.2)))
OVW(NUMPROC(NUMPROC))
OVW(CPUBUSY(CPUBSY))
OVW(APPLPER(APPLPER(POLICY)))
OVW(EXCPRT(EXCPRT(POLICY)))
```

The macro **Workload Overview Report** processes converted Overview Records which have been created based on this control set.

Device Overview Report

The purpose of this report is to observe the utilization and response time for devices of interest in your installation.

The template consists of a variable part which must be repeated for each device included in the Overview Record and a fixed part with one statement reporting the processor utilization.

Template:

```
----- Variable Part, replace vvvvvv with
----- the VolSer of the device
OVW(QTvvvvvv(DQTAVG('vvvvvv'))))      Average IOS Queue Time
OVW(PTvvvvvv(DPTAVG('vvvvvv'))))      Average Device Pending Time
OVW(DTvvvvvv(DDTAVG('vvvvvv'))))      Average Device Disconnect Time
OVW(CTvvvvvv(DCTAVG('vvvvvv'))))      Average Connect Time
OVW(ARvvvvvv(DART('vvvvvv'))))        Device Activity Rate
-----
OVW(CPUBUSY(CPUBSY))                  should be added at the end
```

You can either use the macro **Create Overview Control Statements** to generate control statements for any device in your installation, or you can use the function **Create Exceptions** which is part of the macro **DASD Activity Report**, pointing to your most important volumes.

If you want to use the template to monitor the two devices MVS215 and MVS217, you can create the following overview control statements:

Example:

```
OVW(QTMVS215(DQTAVG('MVS215'))))
OVW(PTMVS215(DPTAVG('MVS215'))))
OVW(DTMVS215(DDTAVG('MVS215'))))
OVW(CTMVS215(DCTAVG('MVS215'))))
OVW(ARMVS215(DART('MVS215'))))
OVW(QTMVS217(DQTAVG('MVS217'))))
OVW(PTMVS217(DPTAVG('MVS217'))))
OVW(DTMVS217(DDTAVG('MVS217'))))
OVW(CTMVS217(DCTAVG('MVS217'))))
OVW(ARMVS217(DART('MVS217'))))
OVW(CPUBUSY(CPUBSY))
```

The macro **Device Overview Report** processes converted Overview Records which have been created based on this control set.

Cache Subsystem Overview Report

The purpose of this report is to show key characteristics (cache hit rates and cache miss rates) for the cache subsystems of interest in your installation.

The template consists of a variable part which must be repeated for each control unit included in the Overview Record, and a fixed part, which is necessary for the spreadsheet macro to determine the last control unit included in the report.

Template:

```
----- Variable Part, replace cu with
----- the address of the control unit
OVW(RHTcu(CASRHT(SSID(cu))))      Total Read hits rate
OVW(WHTcu(CASWHT(SSID(cu))))      Total Write hits rate
OVW(MRNCu(CASMRN(SSID(cu))))      Normal Read miss rate
OVW(MRScu(CASMRS(SSID(cu))))      Sequential Read miss rate
OVW(MRCcu(CASMRC(SSID(cu))))      CWF Read miss rate
OVW(MWNCu(CASMWN(SSID(cu))))      Normal Write miss rate
OVW(MWScu(CASMWS(SSID(cu))))      Sequential Write miss rate
OVW(MWCcu(CASMWC(SSID(cu))))      CWF Write miss rate
OVW(DFBcu(CASDFWB(SSID(cu))))     DFW bypass rate
OVW(CFBcu(CASCFWB(SSID(cu))))     CFW bypass rate
OVW(DFIcu(CASDFWI(SSID(cu))))     DFW inhibit rate
OVW(NCIcu(CASNCICL(SSID(cu))))     Non-cache I/O ICL rate
OVW(NCBcu(CASNCB(SSID(cu))))       Non-cache I/O bypass rate
OVW(ASYcu(CASASYNC(SSID(cu))))     Async rate
-----
OVW(LASTRNG(CASRHT(SSID(cu))))     should be added at the end
```

You can either use the macro **Create Overview Control Statements** to generate control statements for any cache control unit in your installation, or you can use the function **Create Exceptions** which is part of the macro **Cache Subsystem Report**, pointing to your most important control units (Excel only).

If you want to monitor two control units 0050 and 0068, you can create the following overview control statements:

Example:

```

OVW(RHT0050(CASRHT(SSID(0050))))
OVW(WHT0050(CASWHT(SSID(0050))))
OVW(MRN0050(CASMRN(SSID(0050))))
OVW(MRS0050(CASMRS(SSID(0050))))
OVW(MRC0050(CASMRC(SSID(0050))))
OVW(MWN0050(CASMWN(SSID(0050))))
OVW(MWS0050(CASMWS(SSID(0050))))
OVW(MWC0050(CASMWC(SSID(0050))))
OVW(DFB0050(CASDFWB(SSID(0050))))
OVW(CFB0050(CASCFWB(SSID(0050))))
OVW(DFI0050(CASDFWI(SSID(0050))))
OVW(NCI0050(CASNCICL(SSID(0050))))
OVW(NCB0050(CASNCB(SSID(0050))))
OVW(ASY0050(CASASYNC(SSID(0050))))
OVW(RHT0068(CASRHT(SSID(0068))))
OVW(WHT0068(CASWHT(SSID(0068))))
OVW(MRN0068(CASMRN(SSID(0068))))
OVW(MRS0068(CASMRS(SSID(0068))))
OVW(MRC0068(CASMRC(SSID(0068))))
OVW(MWN0068(CASMWN(SSID(0068))))
OVW(MWS0068(CASMWS(SSID(0068))))
OVW(MWC0068(CASMWC(SSID(0068))))
OVW(DFB0068(CASDFWB(SSID(0068))))
OVW(CFB0068(CASCFWB(SSID(0068))))
OVW(DFI0068(CASDFWI(SSID(0068))))
OVW(NCI0068(CASNCICL(SSID(0068))))
OVW(NCB0068(CASNCB(SSID(0068))))
OVW(ASY0068(CASASYNC(SSID(0068))))
OVW(LASTRNG(CASRHT(SSID(0068))))

```

The macro **Cache Subsystem Overview Report** processes converted Overview Records which have been created based on this control set.

Channel Overview Report

The purpose of this report is to observe selected channels of interest in your installation. The template consists of a variable part which must be repeated for each channel included in the Overview Record and a fixed part with one statement.

Template:

```

----- Variable Part, replace vv with
----- the Channel Path ID
OVW(TBSYvv(CHTBSY(vv))) statements for channels
OVW(LBSYvv(CHLBSY(vv)))

OVW(NUMPROC(NUMPROC)) This part is fixed and must be added
OVW(CPUBUSY(CPUBSY)) at the end of the report
OVW(APPLPER(APPLPER(POLICY)))
OVW(EXCPRT(EXCPRT(POLICY)))

```

If you want to use the template to monitor the two channels 85 and 86, you can generate the following exception control statements:

Example:

```

OVW(TBSY85(CHTBSY(85)))
OVW(TBSY86(CHTBSY(86)))
OVW(LBSY85(CHLBSY(85)))
OVW(LBSY86(CHLBSY(86)))

```



```
OVW(NUMPROC(NUMPROC))
OVW(CPUBUSY(CPUBSY))
OVW(APPLPER(APPLPER(POLICY)))
OVW(EXCPRT(EXCPRT(POLICY)))
```

The macro **Channel Overview Report** processes converted Overview Records which have been created based on this control set.

Spreadsheet usage examples

To follow the instructions in this demonstration, please use the sample Report Listing *Sample.lis* which is located in subdirectory <resource_root>\RmfListings\Sample.lis where <resource_root> is the location of the resources used by the Spreadsheet Reporter, that you specified during the installation process. This sample Report Listing contains one report of each of the following types:

- Cache Subsystem Activity
- Channel Path Activity
- CPU Activity
- DASD Device Activity
- I/O Queuing Activity
- Paging Activity
- Partition Data
- Summary
- Workload Activity

Note: This usage example guides you from the Working Set creation to the graphical presentation with a spreadsheet macro. To transfer all of the above report types contained in the sample Report Listing into the Working Set, ensure that you selected these report types on the **Reports** page in the **Options** dialog (“How to specify Postprocessor report types” on page 319).

1. Start the Spreadsheet Reporter from your desktop. If you use this program for the first time, with no system defined yet, you will receive a message asking whether to want to define a system now. For this usage example, you can answer **No**, because you want to work with an existing local report listing.
2. Open the folder with local **Report Listings**. This will display the *Sample.lis* Report Listing in the **view pane**.
3. Select the *Sample.lis* resource and apply action *Create Working Set*. This invokes the *Create Working Set* dialog where you press the **Run** button to start the generation of a Working Set. Accept the suggested name *Rpt.Sample* for the new Working Set. After successful generation, press the **OK** button. Now you can find this new Working Set in the local **Working Sets** folder.

For the following steps, we assume:

- You have installed Microsoft Excel on your workstation.
 - You have installed the RMF Spreadsheet Macros for Microsoft Excel.
4. Open the *DASD Activity Report* spreadsheet from the Spreadsheet Reporter dialog. Enable the macros for execution with the Spreadsheet Reporter as described in “Setting the security level for Excel macros” on page 339.




Figure 66. DASD Activity Report

5. This spreadsheet contains a button *Select Working Set and process data*. Pressing this button starts a dialog called *Select Report Working Set* which lists the available Working Sets and lets you select your *Rpt.Sample* that you created from the *Sample.lis* Report Listing.
6. Select this Working Set and press **OK**. This starts the following selection dialogs:
 - a. *Select System* which lets you select one of the systems contained in the report. In our example, select **UIG1** and press **OK**.
 - b. *Select Interval* which lets you select a reporting interval. In our example, select **UIG1-10/02/2003-09.05.11-DASD Device Activity** and press **OK**.
7. Now the *DASD Activity Report* macro processes the data from the report contained in the Working Set. On your screen, take a look at the dialog shown in Figure 66 and wait until the *Ready* message appears in the lower left corner. Now you can scroll through the tabs shown at the bottom and view the various charts by using the arrows on the left hand side. For example, click on the *Top10Act* tab to view the chart called *Activity of Top 10 Volumes*, as shown in Figure 67 on page 338.




Figure 67. Activity of Top 10 Volumes

In Figure 67, you see the 10 volumes with the highest activity. Pressing the **Sort Lists** button in the upper right corner, you can display these volumes sorted either by their *I/O intensity*, *Service time intensity* or *Path intensity*.

Intensity expresses the usage of a device: it is the average number of milliseconds a device was active, during one second. This metric is measured in milliseconds per second.

I/O Intensity: The product of activity rate and response time. The average number of milliseconds a device was busy with responding during one second.

Service Time Intensity: The product of activity rate and service time (that is the sum of disconnect time and connect time of a device). The average number of milliseconds a device was busy with service during one second.

Path Intensity: The product of activity rate and connect time of a device. The average number of milliseconds a device was connected during one second.

- When you have finished with the first try, you can close Excel. To leave the Spreadsheet Reporter, click on **File** ---> **Exit** in the menu bar.

Now you are ready to process SMF data from your own installation:

- Define a system from your environment.
- Specify appropriate report types and reporting periods/duration intervals.
- Create an **SMF Dump Data** resource.
- From this resource, create a Working Set.
- Open a spreadsheet to obtain graphical views of the SMF data contained in the Working Set.

Open RMF Overview Spreadsheets usage example

This topic explains how to use the *Open RMF Overview Spreadsheets* macro to create customized graphical reports about certain aspects of system activity.

1. Create a file \$OVWxxxx.txt containing selected overview control statements and attach it to the system. For the example shown in Figure 68, this file would contain only one line: OVW(CPUBUSY(CPUBSY)).
2. Define the SMF dump data containing the data that you want to process.
3. Create an Overview Working Set from this data.
4. Open the *Open RMF Overview Spreadsheets* macro and select the Overview Working Set created in step 3.
5. Select a system or sysplex.
6. Select an interval.
7. Start processing the data.
8. Select a chart type to view the graphical presentation of your data.

Figure 68 shows the Spreadsheet Reporter output format for the overview control statement OVW(CPUBUSY(CPUBSY)).

You can apply a title of your choice to the report and label the x-/y-axis using the *Setup* tab in the spreadsheet. In the example from Figure 68, *CPU Busy* has been specified as the report title.




Figure 68. Spreadsheet Reporter Output — Open RMF Overview Spreadsheets macro

Setting the security level for Excel macros

The Excel macros shipped with the Spreadsheet Reporter are signed. When you modify and save the Excel macros, the signature is discarded.

Within Excel you can specify the security level to enable or disable macro execution. To use the Excel macros with the Spreadsheet Reporter, you must enable these macros for execution. Therefore, you have to set the security level to *Medium* or *Low*.

- *Medium* means that each time a macro is loaded, Excel asks if you want to execute it or not.
- *Low* means that macros are automatically executed when loaded.

Excel 2000:

For Excel 2000, the default security level is *Medium*. To change the security level, from a macro's menu bar use **Tools → Macro → Security**.

Excel 2002:

For Excel 2002, the default security level is *High*. To change the security level, from a macro's menu bar use **Tools → Options → Security → Macro Security**.

Excel 2007:

To change the security level, select **Excel Options → Trust Center → Trust Center Settings... → Macro Settings....**

Chapter 19. RMF Performance Monitoring

Items discussed in this information unit are:

- Performance monitoring - Overview
- Getting started - Installation and setup
- RMF PM - Let's go
- Working with sysplexes
- Working with PerfDesks and DataViews
- Working with Series
 - Series Definition Dialog
 - Plot/Save Series Dialog
 - Filter Dialog
 - Work Scope Dialog
 - Analysis Dialog
- What is monitored?
 - Sysplexes
 - Resources
 - Metrics
 - Work Scopes
- Messages

Performance monitoring - Overview

RMF Performance Monitoring (RMF PM) allows you to monitor the performance of your z/OS host from a workstation through a TCP/IP interface to one or more z/OS sysplexes. You logon to any sysplex and you can monitor the resources in the corresponding sysplex.

RMF PM takes its input data from a single data server on one system in the sysplex, which gathers the data from the RMF Monitor III on each MVS image. Therefore, this is called the **Distributed Data Server (DDS)**. If you want to monitor several sysplexes, each one needs to have an active DDS.

RMF PM provides a selected subset of the information provided by the RMF Monitor III gatherer: general performance data, performance data for jobs, and workload-related performance data like:

- WLM workloads
- WLM service classes
- WLM service class periods
- WLM report classes

You have the flexibility to create unique scenarios that monitor the performance of your system. You can sample real-time data as bar charts, and combine data from different resources together. Once you have created these scenarios, you can save them as **PerfDesks**. With PerfDesks, you create a set of **DataViews** customized to your monitored system(s). DataViews sample performance data into one or more **Series** displayed as bar charts. You can reuse the DataViews any time. You can

simply open the PerfDesk and start it whenever you want to view performance data in your monitored system again from the same angle.

Note: A saved PerfDesk does not contain any performance data. The PerfDesk samples new performance data each time it is opened and started. However, you can save the sampled data in spreadsheet files.

When you open RMF PM, you will find a list of **Sysplexes** in the resource view on the left side. Each sysplex comprises a hierarchy of **Resources** of your monitored system(s). Each Resource has a set of **Metrics**, which specify what is measured, for example, how the Resource is used or loaded with work.

Getting help in RMF PM

All help information will be displayed using the default browser of your installation.

Putting it all together

- You select Resources and Metrics, add them as Series definitions to a DataView.
- You save one or more DataViews as a PerfDesk.
- You start a DataView or PerfDesk to sample performance data on monitored system(s) and display them in the Series of DataView(s).
- You save the Series to a spreadsheet file.
- You open PerfDesks any time later to restart performance data sampling.

Getting started - Installation and setup

Here are the steps, to get **RMF PM Java Edition** working.

- Step 1: Check the Prerequisites
- Step 2: Install the client

Prerequisites

Prerequisites for the client

- Windows XP, Windows Vista, or Windows 7
- Microsoft Internet Explorer 7.0 or later or Mozilla Firefox 3.0 or later.

Level of access:

For Windows XP or later, to install and use the **RMF PM**, you must have *Computer Administrator* account type or higher (see the Windows user management facilities). If you relate to the *Limited* account type, you will encounter problems during the installation and use of the **RMF PM**.

Prerequisites for the z/OS sysplex

RMF PM uses a server on the z/OS sysplex to retrieve its performance data. This server is called the **RMF Distributed Data Server**.

This host component is installed automatically. Please, refer to "Setting up the Distributed Data Server for z/OS" on page 26 with the description how to customize and to start the RMF Distributed Data Server.

Also make sure, that the following prerequisites are met on your z/OS host:

- Unix System Services must be configured in "full function mode".

- TCP/IP under Unix System Services must be configured and active.

Client installation

1. The code and installation utility of the RMF PM is available as member GPMWINV2 of the host distribution library SERBPWSV. Download this member as binary file **gpmwinv2.msi**.
2. Install the MSI package using the Windows Installer, either by double-clicking on the MSI package file or by issuing the command:

```
msiexec /package gpmwinv2.msi
```

The Windows Installer guides you through the installation.
3. Specify the directory where to install the RMF PM. The default is:

```
C:\Program Files\RMF\RMF Performance Monitoring
```

Note: If you have installed an RMF PM version earlier than z/OS V1R11, it is recommended to uninstall it before installing the MSI package.

RMF PM - Let's go

After invoking RMF PM from the program folder **IBM RMF Performance Management - RMF Performance Monitoring**, you see the main window (Figure 69).




Figure 69. RMF PM — main window

There are two views:

- The view on the right is the **Opened PerfDesks** view
- The view on the left with two tabs contains the **PerfDesks** view and the **Resources** view

In the initial setup of RMF PM, the PerfDesk Folder **Samples** containing the PerfDesk **Sysplex-Overview** has been defined to be started automatically when starting RMF PM. This PerfDesk is connected to the sysplex that you have defined during installation. Therefore, the **Sysplex Logon** dialog will now be started.

Sysplex logon dialog




Figure 70. Sysplex logon dialog

Logon to a sysplex requires a **TSO user ID** and a **Password**.

Note: If you protect your Monitor III data as described in “Controlling access to RMF data for the sysplex data services” on page 20, ensure that the user ID is authorized to use the data service protected by this profile. Otherwise you will get an error message.

The Logon dialog will pop up for each sysplex that you have defined with an Automatic Startup PerfDesk Folder. Selecting **Ok** submits a logon request to the host. The request may take a considerable amount of time depending on how far away the sysplex to be monitored is, and how much it is loaded.

The user ID is **preset** in the Logon dialog but can be overwritten. You can change the **preset** user ID in the **Change Sysplex** dialog.

Note: If a logon fails several times, an expired password may be the reason. You should logon as TSO user to the host session, and change the password there before returning to this dialog. You **cannot** change the password from RMF PM. If a logon fails a certain amount of times, the user ID is revoked.

When the logon request has been completed successfully, the initial PerfDesk is loaded and started immediately. Unless you have changed the setup, you will see four DataViews:




Figure 71. Initial DataViews

Accessibility

You can access the RMF PM graphical interface using your keyboard, for example, you use Tab+F6 to toggle between the PerfDesks/Resource view and the Open PerfDesks view. Some examples of how you can access the controls in the PerfDesks using your keyboard are:

Alt-S: start
 Alt-N: next sample
 Alt-T: stop
 Alt-M: Sample...
 Alt-E: Save
 Alt-C: Close
 Alt-U: Startup check-box

Working with sysplexes

With RMF PM, you can monitor the performance in one or more z/OS sysplexes. This section describes how to define a sysplex for monitoring and how to open and close a sysplex.

Sysplex dialog

This dialog will pop up as a result of selecting **File - New - Sysplex...**

Or if you want to change the settings of a sysplex use **File - Change Settings - Sysplex...**




Figure 72. Sysplex Dialog

The areas in this dialog are:

- Sysplex
- Host Name
- Classical RMF PM Port Number
- HTTP Port Number
- User Id
- Communication timeout [sec]
- Date/Time at Sysplex
- GMT-Offset of Time-zone

Sysplex

Here, you can enter a description of the sysplex.

Host Name

The TCP/IP name of the host, which can be either:

- A symbolic name which can be resolved via a name-server or your *etc/hosts* file.
- An IP address like 9.164.182.251.

Classical RMF PM Port number

Here, you should enter the port number of the RMF DDS Server on the host, 8801 is the default value. Your system administrator should be able to give you the number.

HTTP Port number

Type in the associated HTTP port number. Please note that this is normally port number 8803.

User ID

Here, you define your TSO user ID.

Communication timeout [sec]

You may leave the default value of 180 seconds.

Date/Time at sysplex

Date/Time at sysplex must show the current wall clock time at the sysplex. If it is incorrect (for the remote sysplex), select a time-zone with a different GMT-Offset so that the correct **Date/Time at sysplex** is shown.

GMT-Offset of time-zone

This field allows you to select a GMT-Offset using the path **View - Options - Set Time-zone....** If you never changed the selection, the **local** GMT-Offset of your workstation should be shown selected.

Therefore, you need to change it **only** if the host associated with the sysplex resides in a different **Time-Zone**. The GMT-Offset affects the time stamps in your DataViews.

Open and close a sysplex

An open sysplex is characterized by the indicator **[+]** in front of its name. This indicator can be used to expand the sysplex.

Without this indicator, the sysplex is closed and needs to be opened to access its resources. To open a sysplex either double click on it, or select **Open** from its context menu. Then, the **Sysplex Logon** dialog will pop up.

To close a sysplex, select **Close** from its context menu. All contained resources will not be accessible until the sysplex is opened again.

You can also use the path **File - Open/Close - Sysplex...** to open or close a sysplex. This leads you to the **Sysplex Selection** dialog where you have to specify the name of the sysplex.

Expand a sysplex

The **[+]** indicator before the sysplex means that you can expand it. Just click on the **[+]** and the resources should be listed indented below the sysplex.

```
[+]My Sysplex - PRODPLEX, Sysplex
    [+]SYS1, Image
    [+]SYS2, Image
    [+]CF1, Coupling Facility
    [+]CF2, Coupling Facility
```

The **[-]** before 'My Sysplex' means it is expanded. You can collapse it by clicking on the **[-]**.

Note: This will not close the sysplex, it only affects the graphical representation of the resource structures.

Expand a Resource

The **[+]** before a Resource means that you can expand it. If you click on the **[+]**, the contained Resources should be listed indented below the expanded Resource.

```
[+]My Sysplex - PRODPLEX, Sysplex
    [-] SYS1,Image
        [+] SYS1,I/O-Subsystem
            SYS1,Processor
        [+] SYS1,Storage
            SYS1,Enqueue
            SYS1,Operator
        [+] SYS1,S/W-Subsystem
    [+] SYS2,Image
    [+] CF1,Coupling Facility
    [+] CF2,Coupling Facility
```

A Resource without a **[+]** before its name means it does not have contained Resources.

Creating a DataView from a Resource

To create a DataView from a Resource, select **New DataView...** from its context menu. This leads you to the **Creating a DataView** dialog.

Working with PerfDesks and DataViews

A PerfDesk is a set of DataViews that can be created, saved, opened and started altogether. A PerfDesk Folder is a container for one or more PerfDesks.

The **PerfDesks** notebook page allows you to manage your PerfDesks and DataViews from a central point:

- Creating a PerfDesk Folder
- Creating a PerfDesk
- Opening a PerfDesk
- Expanding the PerfDesk Tree into DataViews and Series
- Creating a DataView
- Exporting and importing a PerfDesk

Creating a PerfDesk folder

Using **File - New - PerfDesk Folder...**, you can add a new PerfDesk folder.

Creating a PerfDesk

Using **File - New - PerfDesk...**, you start the **Save PerfDesk** dialog, which displays a list of all existing PerfDesk folders. Now, you can select one of these folders as container for the new PerfDesk, or you can create a new folder. In addition, you will specify the name of the new PerfDesk by overwriting the preset name. With clicking on **Save**, you complete this process.

Opening a PerfDesk

You can open a PerfDesk by double-clicking or through its context menu. This will add a new tab to the right part of the window, called **Open PerfDesks**, and you will see all DataViews belonging to the PerfDesk. Using **File - Open - PerfDesk...** is another way for this task.

Expanding the tree

All containers in the **PerfDesks Overview** which are not empty, can be expanded very easily:

- PerfDesk Folder → PerfDesks
- PerfDesk → DataView
- DataView → Series

Example:

If a PerfDesk contains DataViews, you will see a **[+]** in front of its name:

[+]Sysplex-Overview

If you click on the **+**, the DataViews of the PerfDesk will be shown:

[-]Sysplex-Overview
 [+]Processor utilization of systems in Sysplex
 [+]Performance index of most important workloads
 [+]I/O: Volumes with highest activity in Sysplex
 [+]Users per image

In an **Open PerfDesk**:

- You have buttons:
 - Start and Stop all DataViews of the PerfDesk
 - Save a PerfDesk
 - Close a PerfDesk
- Additionally, using the context menu, you can:
 - Export a PerfDesk
 - Add a DataView to a PerfDesk
 - Change the Name of a PerfDesk
 - Paste a DataView to a PerfDesk from Clipboard

Start and stop a PerfDesk

Clicking on **Start** will start data sampling in all DataViews belonging to the PerfDesk, this is indicated by the Run Light of all DataViews which will turn to green. Clicking on **Stop** stops data collection.

Of course, you can also start or stop each DataView individually by clicking on the appropriate Run Light.

Save a PerfDesk

Before you save a PerfDesk, make sure the name of the PerfDesk and its DataViews reflect what is sampled in its Series. Neither PerfDesk, nor one of its DataViews should be named 'New'. Clicking on **Save...** opens the **Save PerfDesk** dialog, here you have the option to select a PerfDesk Folder, where you want to save the PerfDesk.

Close a PerfDesk

During **Close** you will be prompted whether you want to save any changes that you have performed to the PerfDesk.

Stepping through the history of DataViews

You can step forward and backward through the history of DataViews in a PerfDesk.

Backward in History

Step all DataViews of this PerfDesk back one sample by pressing the < button. If the data needed is not available, the affected DataViews are stopping.

Forward in History

Step all DataViews of this PerfDesk forward one sample by pressing the > button. If the data needed is not available, the affected DataViews are stopping.

Sample

Open the following dialog by pressing the **Sample...** button:




Figure 73. Collecting historical data in RMF PM

You can specify a start date/time and a stop date/time. After you've pressed the **Ok** button, all available performance samples for all DataViews of this PerfDesk for the specified time range are automatically collected.

Note that RMF PM gets its data from the z/OS host, meaning that all performance data has to be calculated by the z/OS RMF Monitor III. So if you're calculating the data for many different times, you can use some amount of CPU on the host.

Synchronize all DataViews

All DataViews of this PerfDesk are synchronized to the time of the selected DataView.

Startup

If this checkbox is checked, this PerfDesk is automatically started each time you are starting RMF PM. Please note that you do not need to save the PerfDesk after changing the Startup attribute of a PerfDesk; it is directly written to the disk.

Add a DataView to a PerfDesk

Context item **New DataView...** is one way for Creating a DataView.

Change the name of a PerfDesk

The context item **Change Description...** offers the function to change the name of the PerfDesk.

PerfDesk or DataView should be given names so that their purpose can be easily identified. PerfDesk names should sum up the purpose of its DataViews and DataView names should reflect what is sampled in their Series.

Names can contain alphanumeric characters and spaces.

You can also use the path **File - Change Settings - PerfDesk** to rename a PerfDesk.

The corresponding paths are also available for changing the name of a PerfDesk folder.

Paste DataView to a PerfDesk

If a DataView has been previously copied, it can be pasted to the selected PerfDesk.

Exporting and Importing a PerfDesk

If you want to transfer PerfDesks to another user or another PC, you may use the export functionality of RMF PM.

Invoke this function via **File - Export PerfDesk...** to export a **single** PerfDesk. RMF PM asks you for a file name where the exported PerfDesk description will be saved.

Alternatively you may use the **Export ... option** from the context menu of a PerfDesk which exports this current PerfDesk.

In the same way, you may import previously exported PerfDesk descriptions into a PerfDesk folder by using **File - Import PerfDesks...** which lets you select multiple files to be imported in one step. Importing PerfDesk descriptions from explicitly selected files into a PerfDesk folder can also be done via the **Import ... option** from the context menu of a **PerfDesk folder**.

To export **all** PerfDesks from an existing PerfDesk folder at once, use the **Export Folder ... option** from the PerfDesk folder context menu. This function exports all PerfDesks from this folder into files in a target directory. The file names are derived from the PerfDesk names.

If you use the **Import Folder ... option** from the PerfDesk folder context menu, you must select a directory containing PerfDesk descriptions, all of which are then imported into the PerfDesk folder.

DataView




Figure 74. Sample DataView

A DataView displays the performance data currently sampled as so called Series.

Creating a DataView

For adding a DataView to a PerfDesk, you select **New DataView...** from the context menu either of an open PerfDesk in the **PerfDesks** notebook page or of a PerfDesk in the **Open PerfDesks** window. This will start the **New DataView - Properties** dialog, which offers you an input field to specify the **Title** of the DataView.

Furthermore:

- You can select whether you want to have a fixed scale in the DataView, or whether you prefer rescaling of the DataView depending on the values of the measurement data (check **Rescale automatically**).
- You can specify whether you want the Series to be displayed as vertical or horizontal bars.

Clicking **Ok** leads you to the **Series Definition** dialog where you can define the metrics of the Series that you want to add to the DataView.

Using **File - New - DataView...** will also lead you to this process.

Note: For defining a Series, it is necessary that you have selected a resource in the **Resources** notebook page. If this is not the case, you will get the message **Select a Resource of the correct type**. Now, you have to **Close** the dialog and to select the resource of your interest in the **Resources** notebook page.

Having in mind the above note, you also can use a resource as starting point of the above described process:

- You select a resource in the **Resources** notebook page and get the context menu for defining a DataView.
- If the resource of your interest is currently not displayed, you can use the path **Actions - Find I/O Resource...** which leads you to the **Find I/O Resource** dialog where you can specify the resource. This might be of special interest if you are looking for a specific volume which normally is not displayed in the **Resources** notebook page due to the large number of volumes in a sysplex. Therefore, this alternate path could be the faster way to find a volume for creating a new DataView with performance data for this specific volume.

Another function of this dialog is to display the properties of a specific resource.

Both ways will lead you to the **New DataView - Properties** dialog.

Working with DataViews

A DataView contains three areas:

- Bar Chart
- Legend
- Control Panel

The DataView has a context menu for actions applicable to the entire DataView as opposed to a single Series. This menu will pop up, if the right mouse button is selected in the background of the bar chart.

- Properties
- Add Series
- Plot/Save Series
- Remove Values
- Copy
- Print

DataView bar chart

When creating a DataView, you specify whether you want to create a DataView with horizontal or vertical bars. All following descriptions refer to vertical bar charts.

The bars represent the values measured, they are colored so they can be identified in the legends. There are two types of a bar chart:

- Single-Value Chart
- Value-List Chart

Note: It is **not** possible to display a Single-Value and Value-list chart in the **same** DataView.

Single-Value chart: A single-value chart displays Series with **one value per time-stamp**, for example, **% delay** or **# active users**.

Value-list chart: A value-list chart displays Series having for one metric (for example **# active users**) a list of values reported per time-stamp: **# active users by MVS image**.

Bar chart context menu: Several actions can be applied to a bar in the chart:

- Analysis
- Find highest
- Find lowest
- Series Settings
- Remove Series
- Color Chooser

Analysis: If you see a value that could be an indicator for a performance problem, analysis might be useful. Selecting **Analysis** leads to the Analysis dialog which offers you other DataViews with Series related to the value which you want to analyze.

Find highest: Selecting **Find highest** will search for the highest value in the Series and will display it.

Find lowest: Selecting **Find lowest** will search for the lowest value in the Series and will display it.

Series settings: Selecting **Series Settings** leads to the **Change Series** dialog. Depending on the type of the Series, it will have up to three pages:

- **Display** - displays the Resource and the Metrics of the Series.
- **Work Scope** - is the Work Scope dialog to specify reporting characteristics for jobs running in a goal mode environment.
- **Filter** - is the Filter dialog to define specific filter criteria for selecting Resources to be shown in the reports.

Clicking the Color Chooser button provides the capability for changing the color of the bar.

Remove Series: Will remove the Series from the DataView. Unless the PerfDesk is saved after the removal, the Series will be present the next time the PerfDesk is loaded.

Color chooser: Here, you can select the color for the Series that you would like to have.

DataView legend

A **Legend** is a description of the Series. It consists of four elements:

- The little color square in front, to identify the bar associated with it
- The Resource for which the Series is sampled
- The Metric for which the Series is sampled
- An extension like filter or workscope

Legend context-menu: Except for 'Analysis', the same actions as in the Bar Chart Context Menu can be applied.

DataView control panel

The control panel of a DataView allows you to control the sampling and display of the Series:

- Control Series Position
- Start/Stop
- Plot/Save Button

Control Series Position

At the bottom of the DataView, you find a **Position Indicator** and a **Time-Slider**.

The position indicator shows two values:

Sample

The position of the first value displayed.

Total Total number of values sampled so far.

Example: Sample: 34 Total: 105

The position of the first value in the DataView is 34 out of 105 values.

By moving the time-slider to the left or to the right, you will change the time frame for which the series will be displayed.

Start/Stop - Run Light

In the left lower corner of the panel, you see the **Run Light** which is also the **Start/Stop Button**. It indicates if the DataView is sampling:

Green The DataView is sampling

Red The DataView does not sample

By clicking on this button, you can start or stop data sampling.

Plot/Save Button

Between the Start/Stop button and the time-slider you see the **Plot/Save Button** which leads you to the Plot/Save Series dialog with the capability to plot a Series or part of it, and to save the data to a .WK1 spreadsheet file.

DataView context-menu

Properties of the DataView

You know the first page from this dialog, when you have specified the values while creating a new DataView. Now, the dialog has a second page Sampling. You

Initially, the spinbutton shows the current time minus one hour, or what was set previously.

Sample To

Selecting the checkbox, you can adjust to any time in the past or future, when sampling should be **stopped**. If the **Sample From** checkbox is not selected, it means "start sampling now" and stop when the Sample-To time is reached.

Initially, the spinbutton shows the current time, or what was set previously.

Note: The time-stamp is affected by the setting of the Time-zone. So, make sure the time-zone is set correctly.

Wrap-around buffer size

This field shows the size of the wrap-around buffer, this is the maximal number of samples that can be available to be displayed.

Set Time-zone

Using the path **View - Options - Set Time-zone...**, you get a panel where you can specify the GMT-Offset and Time-zone.

Series

A Series comprises:

- The definition of what Metric is sampled on a Resource
- The Series of time-stamps with their values or list of name-value pairs.

Series time-stamps

The time-stamps are in the form HH:MM:SS, with **HH** in the 24-hour format. The date in MM/DD format is shown separately in front of the scale below the x-axis in the DataView.

Note: The time-stamp is affected by the setting of the Time-zone. So, make sure the time-zone is set correctly.

Series definition dialog

This dialog will pop up when you:

- Create a DataView...
- Add a Series to a DataView...
- Use the path **File - New - Series...**

It basically allows you to add Series to a DataView. Here you specify all details which are required to sample data on the monitored system like the Sysplex, the Resource, or the Metric, to form a Series definition.

If you have selected a Resource in the **Resources** notebook page, all you need to select is a Metric, because the Resource is preselected. Otherwise, you will be prompted to select a Resource.

The areas in the dialog from top to bottom are:

- Sysplex and Resource: You cannot change directly these fields by overtyping them. If you need another selection, you can specify this in the **Resources** notebook page.
- Metric Types radio buttons
- Metrics list
- Buttons: Add - Done - Metric Help - Help

Metric types

Selection affects the type of Metrics shown in the Metrics list.

- **All** - Shows Single-Value and Value-List Metrics
- **Single** - Shows Single-Value Metrics
- **List** - Shows Value-List Metrics

See also Single-Value Chart and Value-List Chart.

Metrics list

Contains the types of Metrics available for the Resource. Context item **Description** on a selected Metric or button **Metric Help** will pop up a description.

Push Buttons: Add - Close - Metric Help - Help

If a Metric is selected, clicking on **Add** causes the resulting Series to be added to the DataView. The dialog is not ended, thus allowing further selections.

The **Close** button ends the dialog.

Clicking on **Metric help** leads to an explanation of the selected Metric.

Clicking on **Help** displays general information about this dialog.

Plot/Save Series dialog

This dialog will pop up as a result of selecting Plot/Save Series... in the DataView's context menu, or by clicking the Plot/Save button. It allows you to view Series plots and save Series of a DataView into a CSV file.

You can select:

- One Series from a list of Series
- Multiple Series if the Series contain Single-Values
- A range to be zoomed (from one time-stamp to another time-stamp)
- Series for one or more value names (on Value-Lists only) selectable from two Value Name lists , sorted by highest maximum values and by highest average value in descending order.

The areas in the dialog from top to bottom are:

- Series Plot Area
- Control Panel
- Dialog Button Area

Series plot area

In this area, the values of Series are plotted over time. The parts in this area from top to bottom are:

- Line plots of one or more values plotted over time. Triangles pointing upwards indicate maximum values, and those pointing downwards indicate minimum values. The plotted values belong to one or more Series, described by legends below the x-axis.
- A horizontal **Range-Bar** below the x-axis indicates which range of values can be saved or zoomed.

By changing the time-stamps in the **From** and **To** spinbuttons of the Control Panel, the range can be limited. Therefore, only the values lining up above the remaining dark-gray area will be saved or zoomed.

- Below the Range-Bar, two time-stamps are shown. **If not zoomed**, these are the time-stamps of
 - the first sample taken
 - the current last sample taken.

If zoomed, these are the time-stamps of:

- the first sample to be saved
- the last sample to be saved.

- Below the time-stamps, the Series legends are shown.
 - For Series containing Value-lists only one common legend is shown, because only one Series can be selected for one or more value names. The Series for each value name is represented by a different color.

When the dialog pops up, the first Series is shown, but any other Series can be selected in the Series listbox of Control Panel.

- For Series containing single values, more than one Series can be selected.

When the dialog pops up, all Series are shown, but any of them can be eliminated from the plot by deselecting the legend in the Series listbox of Control Panel.

Control Panel

The Control Panel allows you to:

- Select the range of the Series to be saved or zoomed
- Select one or more (if single values) Series to be plotted.
- Select (for value-lists) the value name for which the Series should be plotted.

The areas from top to bottom are:

- The range selection, composed of:
 - The **From** spinbutton, which allows to adjust to the time-stamp of the first sample in the Series be saved or zoomed. Changing the time-stamp will change the start point in the Range-Bar and the number of samples.
 - The **# Samples** in the selected range.
 - The **To** spinbutton, which allows to adjust to the time-stamp of the last sample of the Series to be saved or zoomed. Changing the time-stamp will change the ending point in the Range-Bar and the number of samples.
- The **Zoom** button (if there are at least 11 values in the range), used to expand the dark part of the Range-Bar into the entire plot area. Note, that when pressed, its label changes to **Total** to indicate that by pressing the button again, the total Series can be shown.
- The Series listbox. Note, that this listbox will appear only if the DataView has more than one Series.

The effect of selecting a Series is that it is shown in the Plot Area.

If the Series are Value-lists, then initially the first Series is selected. Selecting another Value-list Series causes the value name with the highest maximum value to be shown. See Value Name lists below for how to select other value names.

If the Series has Single-Values (not Value-lists) then all Series are selected initially. Deselect Series you do **not** wish to see in the plot.

- The **Value Name** lists. Note, that these list boxes will only appear if the DataView contains Series with Value-lists.

Both lists contain **value names**. For example, if the selected Series legend is % **delay by job**, then the names in the list boxes are **job names**.

The first list box contains value names sorted by **maximum values** in descending order. For example, list item **BWSC(30)** means that the maximum delay detected for job BWSC was 30%.

The second list box contains value names sorted by **average values** in descending order. For example, list item **BWSC(2.5)** means that job BWSC experienced a delay of 2.5% on average.

If you select a name from either list, the values for that name are plotted over time. Deselecting withdraws the plot for the name. For example, selecting BWSC will show the % **delay for job BWSC** over the entire time range. The maximum (30%) will be marked by a triangle pointing upwards.

Dialog button area

Besides the usual Cancel and Help buttons, there are the buttons **Save...** and **Print...**

Clicking on **Save...** leads to a **Save Series** dialog. A file dialog will pop up to specify a file which is stored in .CSV format. RMF PM uses the following CSV format: The semicolon (";") is used as list separator. You can specify the decimal separator using *View ---> Options ---> Set Decimal Separator...* from the menu.

Clicking on **Print...** lets you print the plot on the printer of your choice.

Filter dialog




Figure 75. Filter Dialog

Resource and Metric information

This area on the top of the dialog informs you about the Resource and the Metric for which a filter can be specified in this panel.

A multi-valued metric consists of a list of name-value pairs, for example, "Response Time by Volume". A filter is provided to reduce the number of name-value pairs of a multi-valued metric, and to sort them by name or value.

Filtering is performed stepwise in the following sequence:

1. Name-pattern Matching
2. Value-bound Comparison
3. Ordering
4. List-length Reduction

Name pattern

Optionally, one or more name patterns in the form of a simple expression to be matched against the names in the list of name-value pairs of the multi-valued metric can be entered in this field.

The following rules apply to this definition:

- ? - one character
- * - zero, one, or several characters
- If the name contains an '*' (for example, *MASTER*), each * must be preceded by a back slash

Example: You want to define a filter that will accept the following job names:

- XJSMITH1
- *MASTER*
- All BAxx, where xx is any two characters
- All starting with CIC
- All with HOT somewhere in the name

You have to specify in the entry field:

```
XJSMITH1|\*MASTER\*|BA??|CIC*|*HOT*
```

Note: The '\' in '*' means: take * as a character, and not as a character string of any length (wildcard).

A list of valid names to be used as patterns is provided. It may take some time for the program to bring up a long list of these names for the first time. When the selection list is available, the entry field can be filled in by selecting items from the list. Selected items are concatenated with | in the entry field. The entry field can be edited.

Clicking **Refresh** will receive a refreshed selection list from the host.

Value bounds

Optionally, an upper bound and a lower bound to be compared against the values in the list of name-value pairs of the multi-valued Metric can be entered. If both upper and lower bounds are specified, the upper bound must be greater than or equal to the lower bound.

All list elements with values higher than the specified upper bound, and all list elements with values lower than the specified lower bound, are discarded.

Sort order

For ordering the values in the list of name-value pairs of the multi-valued metric, you must select one of the choices:

- Value ascending
- Value descending
- Name ascending
- Name descending

List length

For restricting the length of the list of name-value pairs of the multi-valued metric, you must select one of the choices:

- Highest values
- Lowest values

Additionally, the maximum number of list elements must be specified.

Work Scope dialog




Figure 76. Work Scope Dialog

Resource and Metric information

This area informs you about the Resource and the Metric, for which a Work Scope name can be specified in this panel.

Work Scope type

You have to select a work scope type:

- Total
- Job
- WLM Workload
- WLM Service Class
- WLM Service Class Period
- WLM Report Class

Only those radio buttons are enabled that are available for the selected combination of resource and metric.

Work Scope name

For **Work Scope Type = Total** nothing has to be entered here.

For other work scope types, a list of valid names is provided. It may take some time to create a long list of these names for the first time. When the selection list has been filled, the entry field may be filled in by selecting an item from the list.

RMF PM will remember the selection list and will display it again the next time. **Refresh** can be used to get a refreshed selection list.

Analysis dialog

The **Title** shows the sysplex for which analysis is intended.

Resource, Work Scope, Metric, Value, Name, and Sample Time

These fields inform you about the source of the analysis which is the context of the data point in the DataView that you clicked on. You can either have analysis based on

- the time when data was collected, or
- **current** data, which analyzes the next Metric sample.

If the work scope field is empty, then the Metric is not for a specific work scope rather than global.

Analysis Type

This listbox allows you to select another PerfDesk as the next step of the analysis. Each alternative is shown with the resource and its work scope (if required) as the target for the analysis.

Close Previous PerfDesk(s)

The PerfDesks that were previously shown during the analysis process can be closed to avoid an overloaded screen with too many DataViews. The selected PerfDesks will be closed if the checkbox is selected.

What is monitored?

The following sections describe the objects involved in monitoring with RMF PM.

A metric summary documentation is provided on the RMF web site in the document What you can monitor - the metrics available at <http://www.ibm.com/systems/z/os/zos/features/rmf/product/rmfhtmls/pmweb/pmweb.html>.

You can monitor these metrics using one of the following monitoring applications:

- RMF PM
- IBM Tivoli Business Systems Manager for z/OS
- IBM Tivoli Monitor (Tivoli DM)

Some of the metrics monitored with RMF PM are also available in RMF Monitor III and are described in chapter *Interactive Performance Analysis with Monitor III* in *z/OS RMF Report Analysis*.

Sysplexes

In general, RMF PM can monitor one or more sysplexes, where each sysplex provides access to system management information for a certain set of resources.

Resources

- Sysplex
 - MVS Image

- I/O-Subsystem
 - Storage Subsystem (SSID)
 - Logical Control Unit (LCU)
 - Channel
 - Volume
 - zFS Aggregate
 - zFS File System
- Processor
- Memory (Processor Storage)
 - Central Storage
 - CSA
 - ECSA
 - SQA
 - ESQA
 - Auxiliary Storage
- Enqueue
- Operator
- SW-Subsystem
 - JES
 - HSM
 - XCF
- CPC
 - LPAR
- Coupling Facility
 - Coupling Facility Structure

A resource is any facility of a computing system or an operating system required by a job or task. This includes storage, processor, channels, volumes, or software subsystems. Resources are named according to the following naming conventions:

- A resource is unique in the sysplex

SYS1,Image
CF1A,Coupling Facility

- A resource has several instances in the sysplex

SYS1,0020,Storage Subsystem
SYS1,0030,Storage Subsystem

SYS1,TSO001,Volume
SYS1,TSO002,Volume

Sysplex

A sysplex consists of MVS images and coupling facilities.

You can either expand the sysplex to these images, or you can select a new DataView. The available data is either

- total data for the sysplex or
- workload-related information provided for

- WLM Report Classes
- WLM Service Classes
- WLM Service Class Periods
- WLM Workloads

In addition, details for the sysplex (policy definition) are available.

MVS Image

An MVS image consists of several resources to which it can be expanded:

- I/O Subsystem
- Processor
- Storage
- Enqueue
- Operator
- SW Subsystem

Some of these resources can be expanded, or you can select a new DataView. The available data is either total data or data with job-related information.

I/O-Subsystem

You can either expand this resource to see the details of the I/O configuration, or you can select a new DataView with information belonging to the I/O subsystem.

- Storage Subsystem (SSID)
- Logical Control Units (LCU)
- Channels
- Volumes

Storage Subsystem (SSID)

You can either expand this resource to see SSIDs, or you can select a new DataView with information about cache hits and misses by SSID.

In addition, details for an SSID (type, model, storage size, NVS) are available.

Logical Control Unit

You can either expand this resource to see all LCUs, or you can select a new DataView with information belonging to LCUs.

Channel

You can either expand this resource to see all channels, or you can select a new DataView with information about the utilization of channel paths in the system.

In addition, details for the channel (channel type) are available.

Volume

Due to the fact that typically the number of volumes in an installation is very high, not all volumes will be shown when expanding this resource level, but ranges of volumes, which can be expanded to single volumes with a second step.

You can select a new DataView with information belonging to specific volumes.

In addition, details for the volume (device type, device address, CU type) are available.

zFS Aggregate

You can select a new DataView with information about zFS activity related to specific zFS aggregates or related to one or more file systems contained in an aggregate.

Processor

You can select a new DataView with information about the usage of the processor and about delays for jobs waiting for the processor.

In addition, details for the processor (model and version) are available.

Memory (Processor Storage)

You can select a new DataView with information about the usage of storage and about delays for jobs waiting for storage. The information is available in some more detail for the different types of storage:

- Central Storage
- Auxiliary Storage

Central Storage

You can select a new DataView with information about the usage of central storage and about delays for jobs waiting for storage. The information is available in some more detail for the different areas of central storage:

- CSA - Common System Area
- ECSA - Extended Common System Area
- SQA - System Queue Area
- ESQA - Extended System Queue Area

Auxiliary Storage

You can select a new DataView with information about the usage of auxiliary storage slots.

Enqueue

You can select a new DataView with information about delays in the system caused by usage of serially reusable resources.

Operator

You can select a new DataView with information about delays in the system caused by jobs waiting for the operator to reply to a message or mount a tape, or by address spaces that are quiesced by the operator.

SW-Subsystem

You can select a new DataView with information about delays in the system caused by jobs waiting for service from

- JES - Job Entry Subsystem
- HSM - Hierarchical Storage Manager
- XCF - Cross-System Coupling Facility

CPC

You can select a new DataView to monitor values that are relevant for software pricing. You can also get information about partition related processor activities.

Coupling Facility

You can select a new DataView with information about performance and usage of Coupling Facilities installed in your Parallel Sysplex.

Coupling Facility Structure

You can select a new DataView with information and usage of the Coupling Facility Structures on your Coupling Facility.

Metrics

RMF PM has two formats for presenting performance data:

- **Single-Value Metrics**, for example
 - % **utilization** (of a processor, of a channel, ...)
 - **i/o activity rate** (of a logical control unit, ...)
- **Value-List Metrics**, for example
 - % **utilization by job**
 - # **delayed jobs for i/o by mvs image**

The unique indicator in the name of a Value-List Metric is the keyword **by**.

Work Scopes

A work scope is the specification of an entity of work. RMF PM supports the following work scopes:

- Job (representing the work performed in an address space)
- WLM workload
- WLM service class
- WLM service class period
- WLM report class

Metrics with values for work scope instances are available in two ways:

- As a single-valued Metric, where the corresponding work scope name has been specified
- As a multi-valued Metrics (ordered lists), where each list element belongs to one instance of a work scope

Work scopes are not modeled as resources showing up in a PM configuration view, because frequently changing instances of jobs would flood the system with configuration updates.

Message Browser

Whenever an important message has to be brought to the user's attention, the Message Browser will pop up (if not opened yet), and a message will be displayed emphasized by a short beep.

Each message has the following format:

```
YYYY/MM/DD HH:MM:SS  
GPMxxxxI Message.....  
.....
```

The following actions are possible on the Message Browser:

- Save Messages to a File,
- Delete Messages,
- Query Help for a Message,
- Close the Message Browser

Save messages to a file

Using the path **File - Save Messages ...**, you can save the messages to a file. By convention all message files have a file extension of **.msg**.

Delete messages

To delete messages, simply select the part you would like to remove, and press the **Delete** key.

Query help for a message

To query help for a messages, select it (at least the message-identifier) and press F1. You can also use the path **Help - Message Help** after having selected the message.

Close the Message Browser

Use the File menu or double click on the Message Browser icon to close it. All messages will be lost. If a new message needs to be displayed, the Message Browser will pop up again.

Chapter 20. z/OS Management Facility - Resource Monitoring

IBM z/OS Management Facility (z/OSMF) provides a framework for managing various aspects of a z/OS system through a Web browser interface. By streamlining some traditional tasks and automating others, *z/OSMF* can help to simplify some areas of system management and reduce the level of expertise needed for managing a system.

Most functions of *z/OSMF* area are provided through plug-ins, which you can choose to enable when you configure the product. Choosing the ***Resource Monitoring*** plug-in adds the Resource Monitoring and the System Status tasks to the **Performance** category.

The ***Resource Monitoring*** task allows you to monitor the performance of the z/OS sysplexes, AIX system complexes (System p), Linux system complexes (System z and System x), or Windows system complexes (System x) in your environment. You can monitor the supported metrics, create and save custom views of the metrics, and display real-time data as bar charts.

For z/OS sysplexes, the ***Resource Monitoring*** task takes its input from the RMF Distributed Data Server (GPMSEVER). To monitor several sysplexes, ensure that each sysplex has an active Distributed Data Server.

Similarly for AIX, Linux or Windows system complexes, the ***Resource Monitoring*** task takes its input from the RMF XP GPM4CIM started task, which is a cross platform Distributed Data Server on a z/OS system that gathers data via connected CIM servers from the systems to be monitored.

Before you can use the ***Resource Monitoring*** task, you must define the sysplexes, system complexes, or images that you want to monitor. For this purpose, you use the ***System Status task***. This task provides a comfortable way to assess the health status of all systems in your installation at a glance. Figure 77 on page 370 presents a sample list of z/OS sysplexes. You can add all target sysplexes, system complexes, or images you want to monitor to this list.




Figure 77. System Status task

The *Performance Index Status* column provides a red, yellow, or green indicator for the health of the sysplexes, system complexes, or images. If all service class periods on the systems are meeting their goals (that is, they have a performance index of less than or equal 1), the status is green. If at least one service class with importance 1 or 2 misses the WLM goal, the status indication is red. In this case, you can figure out the reasons by looking at the details on the respective sysplex, system complex, or image, using the *Resource Monitoring task*.

A *Monitoring Dashboard* is a customizable view containing a set of performance metrics. You can create and save your own dashboards or open and modify the predefined dashboards shipped with the *z/OSMF Resource Monitoring* plug-in.

Figure 78 on page 371 shows how you can use the *Resource Monitoring task* of *z/OSMF* to monitor the common storage activity of a sysplex, system complex, or image.




Figure 78. Resource Monitoring task

A comprehensive online help is provided for each panel. Tooltips, descriptive texts as well as error, warning, and information indications will guide you through the panels and ensure an intuitive user experience. For more information about *z/OSMF*, visit the *z/OSMF* home page available at <http://www.ibm.com/systems/z/os/zos/zosmf/> or refer to the *IBM z/OS Management Facility Configuration Guide*.

Chapter 21. RMF Client/Server enabling (RMFCS)

RMF Client/Server Enabling (RMFCS) uses the client/server concept to make your performance management independent of a TSO/E session on the host system you are managing.

This information unit covers the following topics:

- Overview of RMFCS
- RMFCS scenarios
- Installation and startup of RMFCS
- RMFCS usage considerations
- RMFCS component overview
- RMFCS procedures and EXECs

What is RMF Client/Server enabling ?

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.




Figure 79. RMFCS Performance Data View - Example

The example shows an RMF-PWS client with

- A Monitor III SYSINFO report session connected to z/OS *System_1*
 - A Monitor II DEVICE report session connected to z/OS *System_n*
- running concurrently.

With RMFCS, you can establish as many sessions as you want with any z/OS systems in your network that have an APPC or TCP/IP connection configured to your PWS.

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.

Hitherto, one or more 3270 TSO sessions were used for online monitoring of z/OS performance data. The new concept of RMFCS uses a workstation as the single point of control for multiple z/OS systems.

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

The fact that both APPC and TCP/IP can be configured as communication vehicles enhances the availability of the RMF performance data.

RMFCS supports event-driven monitoring. That is, predefined events on the z/OS hosts can be configured to initiate performance monitoring. These events may be either specific system messages, or selected performance data counters that exceed predefined Monitor III exception thresholds.

RMFCS monitoring scenarios

To get an idea of the different possibilities of RMFCS, let us look at three scenarios, illustrating how monitoring can be initiated by:

- Messages
- Exceptions
- Commands

Scenario I: Message-initiated monitoring




Figure 80. RMFCS Scenario I: Message-Initiated Monitoring

In this scenario, the MPF parmlib member is used for event handling and further processing of system alerts. It is assumed that the special emergency events that will trigger the monitoring task are also producing specific console messages.

Sample MPFLSTxx Entries

```
/*-----*/
/* MESSAGES THAT RESULT IN AN ACTIVATION OF AN RMFCS SESSION */
/*-----*/
```



```

-----,
IEA995I,SUP(NO),USEREXIT(ERBCSACT),AUTO(ERBCSGUI) AUTOSTART RMFCS
IRA100E,SUP(NO),USEREXIT(ERBCSACT),AUTO(ERBCSGUI) AUTOSTART RMFCS
-----,

```

On the basis of this example, the following happens:

1. User exit ERBCSACT gets control if the supervisor produces symptom dump output (message IEA995I), or if the system resource manager has recognized an SQA storage shortage (message IRA100E).
2. Module ERBCSACT now issues a MODIFY command for the started task RMFCSC (Client Server Control) which may have been started automatically during system IPL.
3. RMFCSC then receives the name of a REXX EXEC, passed as token to the user exit through the AUTO parameter.
4. The REXX EXEC (here ERBCSGUI) is then executed unconditionally in the RMFCSC address space.
5. Immediately, ERBCSGUI requests a connection to the listening workstation.
This actual bind can be performed in batch mode by exploiting the ISPF GUI feature with the following command:

```
ISPSTART PANEL(ISR@PRIM) NEWAPPL(ISP) GUI(LU:NET_id.LU_id) +
        TITLE(RMFCS_cvtsname) GUISCRW(121) GUISCRD(32)
```
6. The GUI session pops up immediately, and the affected system identifies itself on panel ISR@PRIM and is also displayed as part of the window title bar.
7. From there, the user has unlimited access to all RMF Monitor III and Monitor II reports needed to analyze the critical situation.
In case of message IRA100E, the Monitor III STORC/STORCR reports would immediately provide detailed SQA storage information.
8. Afterwards, the session can be stopped by simply closing the GUI window.

Simultaneous client sessions with simultaneous connections to different z/OS hosts are supported.

Scenario II: Exception-initiated monitoring




Figure 81. RMFCS Scenario II: Exception-Initiated Monitoring

The concept of scenario I can easily be adapted to support exception-initiated monitoring.

This requires a Monitor III reporter address space to be running in batch mode. The sample job RMFM3B is provided to achieve this.

Whenever a new Monitor III report is produced, the ERB3RPH3 procedure can now check whether the actual performance data values exceed the thresholds.

If they do, ERB3RPH3 can activate the PWS connection by producing a predefined message, and further processing continues as in Scenario I.

Scenario III: Command-initiated monitoring

This concept is a subset of Scenario I. It covers the situation in which a monitoring session is required on the workstation, but none of the events described in scenarios I or II has occurred on the z/OS system.

You can simply force a GUI connection by issuing the `MODIFY` console command, which passes a PWS target address directly to the started task `RMFCSC`.

In other words, the function of module ERBCSACT in the event-driven scenarios has simply been replaced in Scenario III by a direct intervention.

Note: For the command-initiated monitoring, it is assumed that z/OS system commands can be either issued directly or transmitted to the affected system.

Installation and startup of RMFCS components

Before you start client/server monitoring, you must check that your system fulfills the prerequisites, and carry out installation and customization of RMFCS.

Prerequisites

The following software and hardware are required for installation and usage of RMFCS:

- Host Software: a z/OS Communications Server network connection from the workstation to the host. If APPC is used, the connection must be capable of supporting parallel LU 6.2 sessions.
- Workstation Software: all operating systems that support the z/OS ISPF Client/Server.
- Workstation Hardware
 - Processor: There are no specific requirements.
 - Display: XGA graphic card is recommended (or compatible graphics with 1024 * 768 resolution).

For details related to ISPF, please refer to *z/OS ISPF Planning and Customizing*.

Installation

During SMP/E installation of RMFCS, the following parts will be copied to the appropriate libraries:

- RMFCS JCL procedures to SYS1.PROCLIB
 - RMFCSC
 - RMFM3B
- RMFCS modules to SYS1.SERBLINK:
 - ERBCSCTL
 - ERBCSACT
 - ERBCSWTO
- RMFCS REXX procedures to the RMF CLIST library SYS1.SERBCLS:
 - ERBCSGUI
 - ERBCSINI
 - ERBM3B
 - ERBM3BWXX
 - ERB3RPH3
 - ERB3RP3I
 - ERBR3SYS
 - ERBR3WFX

Verify or adapt the library names in the JCL procedures RMFCSC and RMFM3B (&RMF, &ISPF) according to your environment.

Customization

RMFCS is designed to allow several users to monitor the z/OS system individually. Each user who wants to run this function just has to initialize the personal environment by taking the following steps:

1. Customize ISPF C/S Session
 - Install ISPFCS code on your workstation (see description under ISPF 3.7 on your host system)
 - Start the WSA.EXE on your workstation

or

Copy the WSA.EXE to your startup folder for permanent use

- Verify the correct APPC or TCP/IP connection through a workstation connection of your ISPF session (under *ISPF Settings / Workstation / Workstation Connection*)

2. Customize RMFCS Procedures

- Create &HLQ.RMFCS.CLIST, ensuring that you have consistent data-set attributes for the SYSPROC concatenation in the RMFCSC procedure.
- Copy REXX procedure ERBCSGUI into this data set and specify the address of your workstation:

For **APPC**

```
home_lu = "LU:NET_id.LU_id"          /* Provide your default LU here */
```

For **TCP/IP**

```
home_ip = "IP:IP_address"           /* Provide your default IP here */
```

If you have both an APPC and a TCP/IP connection you can specify both addresses. By default, the APPC address will be chosen first. If the connection cannot be established, ERBCSGUI tries to establish the TCP/IP connection.

If you do not have an APPC address, you should define home_lu = "", then the TCP/IP address is chosen.

RMFCSC is an ISPF background session, and needs a profile data set and a log data set.

- Create the ISPPROF library &HLQ.ISPFCS.ISPPROF in the same format as your private userid.ISPF.ISPPROF (DSORG=PO, RECFM=FB, LRECL=80, BLKSIZE=3120)
- Create the ISPLOG library &HLQ.ISPFCS.ISPLOG (DSORG=PS, RECFM=VBA, LRECL=125, BLKSIZE=3120)

3. Ensure RACF Authorization

Ensure the appropriate RACF authorization for the started tasks.

- Procedures RMFCSC and RMFM3B are defined to run as started tasks

Note: Due to internal dependencies, these names of these tasks cannot be changed.

- These tasks need access authority to the data sets that have been defined in the step *Customize RMFCS Procedures*. This can be gained, for example, by the following commands:

```
RDEF STARTED RMF*. * STDATA(USER(hlq) GROUP(hlqgrp))  
SETR REFRESH GENCMD(*) GENERIC(*) RACLIST(STARTED)
```

4. Initialize Message-Initiated Monitoring

- Define your MPFLSTxx member(s), for example:

```
IEA995I,SUP(NO),USEREXIT(ERBCSACT),AUTO(ERBCSGUI) AUTOSTART RMFCS
```

5. Initialize Exception-Initiated Monitoring

This type of monitoring requires a Monitor III Reporter session running as batch job. Without special preparation, this job will monitor the system on the basis of exceptions that are generated by the Monitor III WFEX automatic customization.

If you want to define other exceptions, you have to create a new data set with ISPF tables by calling procedure ERBM3BW. This procedure performs similar steps to those in the following example. It assumes that:

- You are working with TSO user ID TSO1

- You have selected qualifier BAT1 for your RMFCS data sets
- a. Rename your current Monitor III table data set:
ren RMFZR21.isptable rmftmp.isptable
- b. Start an RMF session. This results in the creation of a new table data set:
rmf ERB0TABL dataset 'TS01.RMFZR21.ISPTABLE' has been created.
- c. Start the Monitor III session and call the Workflow/Exceptions report WFEX, and you get the following report with the standard exceptions:

```

RMF V2R1 Workflow/Exceptions Line 1 of 19
Command ==> _ Scroll ==> HALF

Samples: 100 System: L96S Date: 04/04/13 Time: 15.31.40 Range: 100 Sec

----- Speed (Workflow) -----
Speed of 100 = Maximum, 0 = Stopped Average CPU Util: 37 %
Name Users Active Speed Name Users Active Speed
*SYSTEM 222 16 16 *DEV 24 1 74
ALL TSO 57 2 93 *MASTER* 1 0 80
ALL STC 134 2 23 TSOPROD 57 2 93
ALL BATCH 30 12 2 BTCHPROD 30 12 2
ALL ASCH 1 0 No work
ALL OMVS Not avail
*PROC 96 2 99

----- Exceptions -----
Name Reason Critical val. Possible cause or action
ALL BATCH OPER-Message 12.1 users Awaiting reply to operator request 37.
*SYSTEM OPER-Message 13.9 users Awaiting reply to operator request 90.
BBRU#488 OPER-Message 98.0 % delay Awaiting reply to operator request 78.
BCCSNET OPER-Message 100.0 % delay Awaiting reply to operator request 15.
BEBR#489 OPER-Message 92.0 % delay Awaiting reply to operator request 37.
BGFI#48A OPER-Message 76.0 % delay Awaiting reply to operator request 12.
BJHA#48C OPER-Message 98.0 % delay Awaiting reply to operator request 60.
BJMO#977 OPER-Message 100.0 % delay Awaiting reply to operator request 95.
BJOE#970 OPER-Message 68.0 % delay Awaiting reply to operator request 29.
BPSM#975 OPER-Message 100.0 % delay Awaiting reply to operator request 99.
BRUG#484 OPER-Message 98.0 % delay Awaiting reply to operator request 25.
BUAB#974 OPER-Message 83.0 % delay Awaiting reply to operator request 86.

```

- d. After entering the command **RO**, you get the Report Options panel:

```

RMF WFEX Report Options: Action Panel Line 1 of 23
Command ==> _ Scroll ==> HALF

Enter Action Code in the Action Column. To exit press END.
Action Codes: Select (S) Copy (C) Move (M) Before (B)
               Add (AD) Delete (D) Move Block (MM) After (A)

Action Class Qualifier Indicator Label Row Position
-----
Only Add (AD) and After (A) are valid on this line.
--- SYSTEM WF 1 1
--- TSO WF 1 2
--- STC WF 1 3
--- BATCH ALL WF 1 4
--- ASCH WF 1 5
--- OMVS WF 1 6
--- PROC WF 1 7
--- DEV ALL WF 2 1
--- JOB *MASTER* WF 2 2
--- SRVCLS TSOPROD WF 2 3
--- SRVCLS BTCHPROD WF 2 4
--- SYSTEM EX-ANY
--- JOB EX-UNAVAIL
--- STOR EX-AVG
--- STOR EX-AVG
--- STOR EX-AVG *STOR

```

- e. Now, you can delete the exceptions you do not need (command **D**), and you can define new exceptions (command **AD**). This leads you to this definition panel:

RMF WFEX Report Options: Definition and Criteria

Command ==> _

Enter or edit information below. To view a list of criteria name values, place the cursor in a blank "Name" field and press ENTER.
Exception will be displayed if all criteria of one color in a set are met.

Class	==> _____	For example: SYSTEM, BATCH, JOB, DEV, STC, SRVCLS
Qualifier	==> _____	For example: Jobname, volume serial, job class
Indicator	==> _____	WF, EX-ANY, EX-AVG, EX-GROUP or EX-UNAVAIL
Label	==> _____	Label for workflow monitor or exception line
Alert	==> _____	Alerting signal: BLINK, BEEP, BOTH, NONE
Text	==> _____	Leave blank for default

Criteria set 1					Criteria set 2					Criteria set 3			
Name	<>	Yel	Red		Name	<>	Yel	Red		Name	<>	Yel	Red
_____	_____	_____	_____	or	_____	_____	_____	_____	or	_____	_____	_____	_____
_____	_____	_____	_____		_____	_____	_____	_____		_____	_____	_____	_____
_____	_____	_____	_____		_____	_____	_____	_____		_____	_____	_____	_____
_____	_____	_____	_____		_____	_____	_____	_____		_____	_____	_____	_____
_____	_____	_____	_____		_____	_____	_____	_____		_____	_____	_____	_____
_____	_____	_____	_____		_____	_____	_____	_____		_____	_____	_____	_____

- f. When you have completed all definitions, you can leave the RMF monitoring session, and rename the data set correctly:

```
ren RMFZR21.isptable 'bat1.rmfm3b.isptable'
ren rmftmp.isptable RMFZR21.isptable
```

This ensures that your Monitor III batch session can run with the definitions in data set BAT1.RMF3B.ISPTABLE.

Please note that the exceptions have to be defined very carefully, to ensure that only an really severe condition will produce an exception line. Care is essential, because even one single exception line will initiate the GUI connection.

All exception handling of the WFEX report will be done under control of the WFEX exception handler ERBR3WFX. A sample of this procedure is part of the RMFCS package (see "REXX procedure ERBR3WFX" on page 389). By default, it issues the WTO message:

```
RMF101I
```

for MPF processing.

Startup

Startup the following sessions:

- "RMFCS control session"
- "ISPF C/S session" on page 382
- Figure 79 on page 374

RMFCS control session

Either start procedure RMFCS by commands shown below, or add the commands to the appropriate Parmlib member COMMNDxx to have the task started automatically during IPL of the system:

```
S RMFCS,HLQ=USER#1
S RMFCS,HLQ=USER#2
S RMFCS,HLQ=USER#3
```

Depending on whether message-initiated or exception-initiated monitoring is in effect, the MODIFY command:

```
F RMFCSC,EX:ERBCSGUI
```

is issued, and each of the RMFCSC tasks will request a connection to its specific target at the same time.

ISPF C/S session

Now, with everything ready to run, the only remaining step is:

- Start WSA.EXE on your workstation

RMF Monitor III batch session

The batch session is required for exception-initiated monitoring and can be started in the same way as the RMFCS control session for each user who wants to exception-initiated monitoring:

```
S RMFM3B.USER#1,HLQ=USER#1
S RMFM3B.USER#2,HLQ=USER#2
S RMFM3B.USER#3,HLQ=USER#3
```

Each RMFM3B session works with a personalized ISPF table, &HLQ.RMFM3B.ISPTABLE, so each user can define his or her own WFEX exceptions. To prevent unsuccessful attempts to connect to workstations, these sessions should only be started for those users who have initialized monitoring by an active ISPF C/S session.

Note: The RMFM3B control module does not listen for a STOP command event, so you have to issue the CANCEL command if you want to deactivate your RMFM3B tasks:

```
cancel user#1
```

Setup and usage considerations

Multiple PWS connections to multiple systems

The RMFCS concept can be implemented on more than one z/OS system.

Because the ISPFCS client can handle multiple connections from different origins, one workstation can act as single point of control for several systems within a network.

Prevention of duplicate connections

When a connection is active, it is inconvenient if a second connection is established to the same target. This can happen when an MPF or WFEX condition is met several times within a short time frame.

For this reason, only one RMFCSC MODIFY command can be stacked. Another MODIFY command during an active GUI session will have no effect, as reflected in the following message for the related task:

```
IEE342I MODIFY    REJECTED-TASK BUSY
```

Routing different events to different workstations

Your active MPFLSTxx member may now look as follows:

```
/*-----*/
/* MESSAGES THAT RESULT IN ACTIVATION OF AN RMFCS SESSION */
/*-----*/
```

```

-----,
IEA995I,SUP(NO),USEREXIT(ERBCSACT),AUTO(IEA995I) AUTOSTART RMFCS
IRA100E,SUP(NO),USEREXIT(ERBCSACT),AUTO(IRA100E) AUTOSTART RMFCS
-----,

```

The REXX procedures IEA995I and IRA100E are just copies of ERBCSGUI, and may contain different destination addresses. Thus, the dump event can be routed to a target other than the SQA storage shortage.

In addition, multiple instances of the RMFCS task mean that each user can decide for him or herself what kinds of event to register for. A user can provide multiple copies of ERBCSGUI with member names identical to the messages he is interested in. Then he will receive only connections for the “member instantiated” messages. A missing member will just cause a “command not found” condition for the related task.

GUI session comes up with WFEX

Especially in the context of exception-initiated monitoring, it is often useful to start the GUI session directly with the WFEX report. If an exception criterion on the host is met, the user on the workstation immediately gets the WFEX report that gives the reason for the notification.

To achieve this, just edit the procedure ERBCSGUI and overwrite the statement:

```
stdparm = m0parm                                /* ISPF primary selection menu */
```

with:

```
stdparm = wfparm                                /* RMF monitor III wfex */
```

Exception handling for all Monitor III report data

This powerful function is inherent in the design of the RMFM3B procedure and the exit ERB3RPH3. For any desired report, an RMFM3B instance can be activated at the same time:

```
s rmfm3b.u1_wf,hlq=user_1,report=wfex
s rmfm3b.u1_si,hlq=user_1,report=sysinfo
```

This makes all report data available for processing by the individual phase 3 exits, for example, ERBR3SYS for SYSINFO report. This can be an efficient solution, especially in two cases:

- A specific counter that is not implemented in the Monitor III workflow exceptions should be tracked and should cause an exception
- A threshold for the counter has been set by the WFEX options, but this active threshold value does not fit the current needs and should be temporarily deactivated in favor of another value.

Example

Your WFEX option threshold for the critical TCB+SRB time is set to 90 %. For a specific reason a GUI connection should be initiated when the value exceeds 80 %, but you want to achieve this without editing the WFEX options.

Action

Set the tcbsrb_limit value in procedure ERBR3SYS to 80%, and start an RMFM3B instance with the parameter report=sysinfo.

Exception-initiated monitoring without MPFLSTxx functions

If it is inconvenient to use MPFLSTxx events to trigger your GUI connections, an alternate “fast path” method can be implemented with little effort:

In procedure ERBR3WFX, modify the statement:

```
SELECT PGM(ERBCSWT0) PARM(wtomsg)
```

with the name of your own module (for example, FRMFCSC):

```
SELECT PGM(FRMFCSC) PARM(ERBCSGUI)
```

Module FRMFCSC builds up an internal buffer with the command string:

```
F RMFCSC,EX:ERBCSGUI
```

Afterwards, the command string is passed to the system command interface SVC34 (MGCR macro) for execution. In doing so, module FRMFCSC (instead of ERBCSACT) forces the RMFCSC tasks to initiate the PWS connections. Thus, Scenario II can easily be adapted to work without involvement of MPFLST functions and members.

The automated approach

You can easily improve the efficiency of RMFCS monitoring by combining some of the recommendations above.

Also you can combine some of the RMFCS features with your existing automated operations environment.

A suitable candidate for this is the RMFM3B task in conjunction with the ERBR3WFX Wfex_Handler_3 procedure.

If you enhance this procedure to a table-driven function (this table can be a triplet of MSGID-EXCEPTION-THRESHOLD items), you can keep track of all exception situations that are considered to force an intervention.

Obvious as it might seem, this offers the possibility of making your systems management environment more powerful with remarkably little effort.

Components of RMFCS enabling

RMFCS Enabling consists of a number of:

- JCL procedures
- Modules
- REXX procedures
- CLISTs

These are described in this section.

JCL procedures

A number of JCL procedures mentioned in the provided subtopics are provided to help you enable RMFCS.

RMFCSC - RMF Client/Server control task

This procedure is the focal point for the activation of the PWS connections. It can be started automatically with the IPL of the system. RMFCSC listens for the following commands:

```
F RMFCSC,LU:lu_name
F RMFCSC,IP:ip_address
F RMFCSC,EX:tso_command
P RMFCSC
```

If LU or IP is used, the parameters are simply passed to the REXX procedure ERBCSGUI, which sets up the final parameter string and issues the ISPSTART command for the GUI connection.

For greater flexibility, RMFCSC also accepts the EX:tso_command parameter, executing any valid command and its optional parameters at once.

In other words, the following two commands will have the same effect:

- F RMFCSC,LU:net_name.lu_name
- F RMFCSC,EX:ERBCSGUI LU:net_name.lu_name

RMFM3B - RMF Monitor III BATCH reporter

This is the procedure that runs the RMF Monitor III reports in batch mode. It is required for exception-initiated monitoring.

Scenario II describes the structure and the setup in more detail.

Modules

A number of modules mentioned in the provided subtopics are provided to help you enable RMFCS.

ERBCSCTL - RMF Client/Server control

This module performs the functions of the RMFCSC procedure. It listens for a MODIFY or STOP command, and establishes PWS connections on request.

ERBCSACT - RMF Client/Server activation

This is the user exit module for the MPF processing. It receives the token from the MPF AUTO() parameter. This token should be the name of the command or REXX EXEC that contains the ISPSTART request for the GUI session.

ERBCSWTO - RMF Client/Server WTO support

This support module allows WTOs to be issued from a REXX Procedure. It is used by Procedure ERBR3WFX to trigger the MPF processing and the subsequent GUI connection when one or more Monitor III workflow exceptions have occurred.

REXX Procedures / CLISTS

Use the REXX procedures or CLISTS described in the following subtopics that help you to enable RMFCS.

ERBCSINI - RMF Client/Server Initialization

The initialization procedure sets the prefix according to the HLQ input parameter and calls the RMFCS control module ERBCSCTL.

ERBCSGUI - RMF Client/Server GUI Connection Setup

This procedure builds the GUI command string from a given parameter or from a predefined default. It then issues the ISPSTART request for the GUI connection. It

also retrieves the system name from the CVT. When the connection has been made, this system name appears in the title bar of the window.

ERBM3B - Monitor III Reporter Batch Control

This is the Monitor III background control procedure. It sets the prefix according to the HLQ input parameter, calls procedure ERB3RP3I for the phase driver table setup and passes control to the RMF Monitor III reporter initialization module ERBCSCTL.

ERBM3BWX - Monitor III RMFM3B Table Switch

This procedure allocates the RMFM3B Monitor III table data set and calls the Monitor III reporter to define the WFEX options.

ERB3RPH3 - Monitor III Reporter Generic Phase 3 Exit

The generic Monitor III reporter phase 3 exit checks the available report type, and calls the corresponding report exit handler for further processing of the Monitor III reporter data tables.

ERBR3WFX - Monitor III Reporter WFEX Phase 3 Sample Exit

This procedure acts as phase 3 exit of the RMF Monitor III workflow exception report.

- Procedure Wfex_Handler_1

Whenever a new report is produced, this function checks whether exceptions have occurred or not. If they have, it calls module ERBCSWTO and issues a predefined message. If the current active MPF member is listening for this message, a GUI connection is initiated.

- Procedure Wfex_Handler_2

This function loops through the WFEX exception table and scans for the exception reasons "OPER-Message" and "Not avail". These are considered to be of minor severity, and are discounted.

If exceptions remain, the Wfex_Handler_2 generates a WTO which contains the exception name, the reason, and the actual counter information.

- Procedure Wfex_Handler_3

This might be the preferred method, because of its flexibility.

It tracks specific, predefined exceptions, and evaluates the worth of issuing a WTO, thereby initiating a GUI session.

In the given example, either a CPU utilization of > 80% or an ESQA storage usage of > 60% will cause two messages with different message IDs.

This allows you to tailor the further MPF processing according to the specific needs and task distribution within your installation. (See also "Routing different events to different workstations" on page 382.)

ERBR3SYS - Monitor III Reporter SYSINFO Phase 3 Sample Exit

This is another RMF Monitor III phase 3 exit sample. It processes the data tables when the SYSINFO report has been requested, for example by:

```
start RMFM3B.si,report=sysinfo
```

- Procedure Sysinfo_Handler_1:

This sample illustrates the access to the header data of the RMF Monitor III reports. All values are easily available through the VGET service. Depending on an internally defined threshold, a WTO will be generated.

In this case, the threshold is TCB+SRB > 90 %.

- Procedure Sysinfo_Handler_2:

The second sample opens and scans the SYSINFO data table ERBSYST3. If a specific instance is found and its threshold is exceeded, a WTO is issued here, too.

ERB3RP3I - Monitor III Reporter Phase 3 installer

This procedure installs the generic phase 3 exit ERB3RPH3 automatically in the phase driver table. It reads the standard phase driver table from the RMF library, sets up the entries for ERB3RPH3 and copies the modified phase driver table to the RMF Monitor III user table library.

Listings of RMFCS procedures

Find a description of the following procedure listings:

- "REXX procedure ERBCSGUI"
- "REXX procedure ERB3RPH3" on page 388
- "REXX procedure ERBR3WFX" on page 389
- "REXX procedure ERBR3SYS" on page 393

REXX procedure ERBCSGUI

```

/* REXX *****/
/*
/*01* MODULE-NAME: ERBCSGUI
/*
/*01* DESCRIPTIVE-NAME: Setup for RMFCS GUI connection
/*
/*01* FUNCTION:
/*      ERBCSGUI sets up the GUI connection and issues the
/*      connection request
/*
/*01* NOTES:
/*      None.
/*
/*01* OPERATION:
/*      1. retrieves the system name from CVT
/*      2. builds the GUI command string from the input
/*          parameter or from default
/*      3. issues the ISPSTART command for the GUI connection
/*
/*01* RECOVERY-OPERATION: None
/*
/*01* DEPENDENCIES: ISPF 4.2.0 environment (or higher)
/*
/*01* INVOCATION:
/*      1. ERBCSGUI LU:NET_id.LU_id
/*      2. ERBCSGUI IP:IP_address
/*      3. ERBCSGUI
/*
/*01* CALLER: ERBCSCTL
/*
/*****/
Trace 0
Parse Upper Arg guiaddr .

home_lu = "LU:NET_id.LU_id"      /* Provide your default LU here */
home_ip = "IP:IP_address"      /* Provide your default IP here */
error_rc = 985                  /* Invalid connection attempt */
/*-----*/
/* Selections for GUI session entry
/*-----*/
m0parm = "PANEL(ISR@PRIM)"
m1parm = "CMD(RMF)"

```

```

m2parm = "CMD(RMF 2)"
m3parm = "CMD(RMF 3)"
wfparm = "CMD(ERBRMF MON3 PARM(WFEX))"
siparm = "CMD(ERBRMF MON3 PARM(SYSINFO))"

stdparm = m0parm                                /* ISPF primary selection menu */
/*-----*/
/* Use internal default, if input parameter is empty */
/*-----*/
If guiaddr = "" Then
Do
    If home_lu \= "" Then guiparm = "GUI("home_lu")"
    Else guiparm = "GUI("home_ip")"
End
    Else guiparm = "GUI("guiaddr")"
/*-----*/
/* Setup the GUI request string and issue the connection request */
/*-----*/
cvt = c2x(storage('10',4))
cvtsname = storage(d2x(x2d(cvt)+x2d('154')),8)
title = "TITLE(RMFCS "cvtsname")"
guiscrw = "GUISCW(121)"
guiscrd = "GUISCARD(32)"
newappl = "NEWAPPL(ISR)"

"ISPSTART" stdparm guiparm title guiscrw guiscrd newappl
/*-----*/
/* Try TCP/IP alternatively, if first try was unsuccessful */
/*-----*/
If (rc = error_rc) &,
    (home_ip \= "") &,
    (home_lu \= "") Then
Do
    guiparm = "GUI("home_ip")"
    "ISPSTART" stdparm guiparm title guiscrw guiscrd newappl
End

Exit rc

```

REXX procedure ERB3RPH3

```

/* REXX ***** */
/*
/*01* MODULE-NAME: ERB3RPH3
/*
/*01* DESCRIPTIVE-NAME: RMF Monitor III phase 3 exit sample
/*
/*01* FUNCTION:
/*      ERB3RPH3 is the generic RMF Monitor III phase 3 exit
/*      for all report types
/*
/*01* NOTES:
/*      None.
/*
/*01* OPERATION:
/*      1. checks the report context (WFEX or SYSINFO)
/*      2. calls the specific report handler
/*      3. prints hardcopy to SYSOUT if the handler
/*          return code is 1
/*
/*01* RECOVERY-OPERATION: None
/*
/*01* DEPENDENCIES: RMF Monitor III Reporter phase 3 context
/*
/*01* INVOCATION:
/*      ISPEXEC SELECT CMD(ERB3RPH3)
/*

```

```

/*01* CALLER: ERB3RDPC */
/* */
/*****
Trace 0

wfex = "WFEX"
sysinfo = "SYSINFO"

ADDRESS ISPEXEC

rc = 0
"VGET (erbrepc) SHARED"          /* Obtain report type */

Select
  When erbrepc = wfex Then
    Do
      "SELECT CMD(ERBR3WFX)"      /* Process WFEX data table */
      If rc = 1
        Then "SELECT PGM(ERB3RDSP)"
      End
    When erbrepc = sysinfo Then
      Do
        "SELECT CMD(ERBR3SYS)"    /* Process SYSINFO data table */
        If rc = 1
          Then "SELECT PGM(ERB3RDSP)"
        End
      End
    /*-----*/
    /* If not WFEX or SYSINFO, just print the report to SYSOUT */
    /*-----*/
    Otherwise "SELECT PGM(ERB3RDSP)"
  End

Exit 0

```

REXX procedure ERBR3WFX

```

/* REXX ***** */
/* */
/*01* MODULE-NAME: ERBR3WFX */
/* */
/*01* DESCRIPTIVE-NAME: WFEX Report Handler Samples */
/* */
/*01* FUNCTION: */
/*   ERBR3WFX provides samples to process the RMF */
/*   Monitor III WFEX report data */
/* */
/*01* NOTES: */
/*   None. */
/* */
/*01* OPERATION: */
/*   Calls the specific WFEX handler subroutine */
/*   depending on the input parameter (default = 1) */
/* */
/*01* RECOVERY-OPERATION: None */
/* */
/*01* DEPENDENCIES: RMF Monitor III Reporter phase 3 context */
/* */
/*01* INVOCATION: */
/*   1. ISPEXEC SELECT CMD(ERBR3WFX 1) */
/*   2. ISPEXEC SELECT CMD(ERBR3WFX 2) */
/*   3. ISPEXEC SELECT CMD(ERBR3WFX 3) */
/*   4. ISPEXEC SELECT CMD(ERBR3WFX) */
/* */
/*01* CALLER: ERB3RPH3 */
/* */
/*****
Trace 0

```

```

Arg handler .

ADDRESS ISPEXEC
CONTROL ERRORS RETURN

hc = 0
msgid = "RMF100I 3B:"
name = "Name      "
reason = "Reason      "
delay = "Critical val."
process = "Processing WFX Report..."
wtomsg1 = msgid process
/*-----*/
/* Header lines for samples 2+3 */
/*-----*/
wtomsg2 = msgid name reason delay
wtomsg3 = msgid SUBSTR("-",1,38,"-")

"SELECT PGM(ERBCSWTO) PARM("wtomsg1")"

Select
  When handler = '1' Then
    Do
      rc = Wfex_Handler_1()
    End
  When handler = '2' Then
    Do
      "SELECT PGM(ERBCSWTO) PARM("wtomsg2")"
      "SELECT PGM(ERBCSWTO) PARM("wtomsg3")"
      rc = Wfex_Handler_2()
    End
  When handler = '3' Then
    Do
      "SELECT PGM(ERBCSWTO) PARM("wtomsg2")"
      "SELECT PGM(ERBCSWTO) PARM("wtomsg3")"
      rc = Wfex_Handler_3()
    End
  Otherwise
    Do
      rc = Wfex_Handler_1()
    End
End

Exit rc
/*****
/*
/*01* SUBROUTINE-NAME: Wfex_Handler_1
/*
/*01* DESCRIPTIVE-NAME: WFX Report Handler - Sample 1
/*
/*01* FUNCTION:
/*      This subroutine provides a sample for a general WTO
/*      notification in case of Monitor III exceptions
/*
/*01* OPERATION:
/*      1. checks if one or more WFX exception lines exist
/*      2. if yes, issues WTO message RMF101I and sets
/*          the hardcopy request to 1
/*
*****/
Wfex_Handler_1: Procedure

tabnam = "ERBWFX3"
msgid = "RMF101I 3B:"
msgtext = "WFX Exception(s) Encountered"

```

```

hc = 0
excpnum = 0
"TBQUERY" tabnam "ROWNUM(excpnum)"
If RC = 0 Then
Do
  If excpnum > 0 Then
    Do
      wtomsg = msgid excpnum msgtext
      "SELECT PGM(ERBCSWTO) PARM("wtomsg")"
      hc = 1
    End
  End
End
Else return 12

return hc
/*****
/*
/*01* SUBROUTINE-NAME: Wfex_Handler_2
/*
/*01* DESCRIPTIVE-NAME: WFEX Report Handler - Sample 2
/*
/*01* FUNCTION:
/*      This subroutine provides a sample for a general WTO
/*      transformation of Monitor III exception lines with
/*      an additional filter for slight exceptions.
/*
/*01* OPERATION:
/*      1. loops through the WFEX exception data table
/*      2. scans for slight exceptions and skips it
/*          (here OPER-Message and Not avail)
/*      3. if one or more exception lines are remaining:
/*          issues WTO message RMF102I and sets the
/*          hardcopy request to 1
/*
/*
*****/
Wfex_Handler_2: Procedure

tabnam = "ERBWFX3"
msgid = "RMF102I 3B:"

oper_message = "OPER-Message"      /* Sets the filter 1      */
not_avail = "Not avail"            /* Sets the filter 2      */

hc = 0
excpnum = 0
"TBQUERY" tabnam "ROWNUM(excpnum)"
If RC = 0 Then
  "TBTOP" tabnam
If RC = 0 Then
  "TBSKIP" tabnam

If RC = 0 Then
Do While (RC = 0)                  /* Loops through the table */

  If (SUBSTR(wfxreasn,1,12) \= oper_message) &,
    (SUBSTR(wfxreasn,1,9) \= not_avail) Then
    Do
      msgtext = SUBSTR(wfxname,1,11) ||
                SUBSTR(wfxreasn,1,14) ||
                SUBSTR(wfxdelay,1,14)
      wtomsg = msgid msgtext
      "SELECT PGM(ERBCSWTO) PARM("wtomsg")"
      hc = 1
    End
    "TBSKIP" tabnam
  End
End
Else return 12

```



```

return hc
/*****
/*
/*01* SUBROUTINE-NAME: Wfex_Handler_3
/*
/*01* DESCRIPTIVE-NAME: WFEX Report Handler - Sample 3
/*
/*01* FUNCTION:
/*      This subroutine provides a sample for the search of
/*      specific Monitor III exception lines and subsequent
/*      threshold handling.
/*
/*01* OPERATION:
/*      1. loops through the WFEX exception data table
/*      2. scans for CPU-utilization and Storage-ECSA-usage
/*      exceptions
/*      3. if the internally defined thresholds are exceeded:
/*      issues WTO messages with individual message ids
/*      (RMF103I, RMF104I) and sets the hardcopy request
/*      to 1
/*
*****/
Wfex_Handler_3: Procedure

tabnam = "ERBWFX3"

proc = "*PROC"
cpus = "CPUS%"
cpus_limit = "90"
cpus_msgid = "RMF103I 3B:"

ecsa = "*ECSA*"
secs = "SECS%"
secs_limit = "60"
secs_msgid = "RMF104I 3B:"

hc = 0
excpnum = 0
"TBQUERY" tabnam "ROWNUM(excpnum)"
If RC = 0 Then
    "TBTOP" tabnam
If RC = 0 Then
    "TBSKIP" tabnam

If RC = 0 Then
Do While (RC = 0)
    If (SUBSTR(wfxname,1,5) = proc) &,
        (SUBSTR(wfxreasn,1,5) = cpus) &,
        (SUBSTR(wfxdelay,2,2) >= cpus_limit) Then
        Do
            msgtext = SUBSTR(wfxname,1,11) ||
                SUBSTR(wfxreasn,1,14) ||
                SUBSTR(wfxdelay,1,14)
            wtomsg = cpus_msgid msgtext
            "SELECT PGM(ERBCSWTO) PARM("wtomsg")"
            hc = 1
        End
        If (SUBSTR(wfxname,1,6) = ecsa) &,
            (SUBSTR(wfxreasn,1,5) = secs) &,
            (SUBSTR(wfxdelay,2,2) >= secs_limit) Then
        Do
            msgtext = SUBSTR(wfxname,1,11) ||
                SUBSTR(wfxreasn,1,14) ||
                SUBSTR(wfxdelay,1,14)
            wtomsg = secs_msgid msgtext
            "SELECT PGM(ERBCSWTO) PARM("wtomsg")"

```

```

        hc = 1
    End
    "TBSKIP" tabnam
End
Else return 12

return hc

```

REXX procedure ERBR3SYS

```

/* REXX *****/
/*
/*01* MODULE-NAME: ERBR3SYS
/*
/*01* DESCRIPTIVE-NAME: SYSINFO Report Handler Samples
/*
/*01* FUNCTION:
/*      ERBR3SYS provides samples to process the RMF
/*      Monitor III SYSINFO report data
/*
/*01* NOTES:
/*      None.
/*
/*01* OPERATION:
/*      Calls the specific SYSINFO handler subroutine
/*      depending on the input parameter (default = 1)
/*
/*01* RECOVERY-OPERATION: None
/*
/*01* DEPENDENCIES: RMF Monitor III Reporter phase 3 context
/*
/*01* INVOCATION:
/*      1. ISPEXEC SELECT CMD(ERBR3SYS 1)
/*      2. ISPEXEC SELECT CMD(ERBR3SYS 2)
/*      3. ISPEXEC SELECT CMD(ERBR3SYS)
/*
/*01* CALLER: ERB3RPH3
/*
/*****/
Trace 0

```

Arg handler .

```

ADDRESS ISPEXEC
CONTROL ERRORS RETURN

```

```

hc = 0
msgid = "RMF200I 3B:"
process = "Processing SYSINFO Report..."
wtomsg1 = msgid process

```

```
"SELECT PGM(ERBCSWTO) PARM("wtomsg1")"
```

```

Select
  When handler = '1' Then
  Do
    rc = Sysinfo_Handler_1()
  End
  When handler = '2' Then
  Do
    rc = Sysinfo_Handler_2()
  End
  Otherwise
  Do
    rc = Sysinfo_Handler_1()
  End
End

```

```

Exit rc
/*****
/*
/*01* SUBROUTINE-NAME: Sysinfo_Handler_1
/*
/*01* DESCRIPTIVE-NAME: SYSINFO Report Handler - Sample 1
/*
/*01* FUNCTION:
/*      This subroutine provides a sample to process the
/*      SYSINFO report header data
/*
/*01* OPERATION:
/*      1. checks if TCB+SRB is higher than 90%
/*      2. if yes, issues WTO message RMF201I and sets
/*          the hardcopy request to 1
/*
*****/
Sysinfo_Handler_1: Procedure

tcbsrb_text = "Average TCB+SRB:"
tcbsrb_limit = " 90"          /* Set the threshold */
tcbsrb_msgid = "RMF201I 3B:"
limit = " Limit:"

hc = 0
"VGET (sysvsvvc) SHARED"      /* Obtain actual value */

If SUBSTR(sysvsvvc,1,3) > tcbsrb_limit Then
Do                               /* Threshold exceeded ? */

    msgtext = tcbsrb_text sysvsvvc||"%" limit tcbsrb_limit||"%"
    wtomsg = tcbsrb_msgid msgtext
    "SELECT PGM(ERBCTSWO) PARM("wtomsg")"
    hc = 1
End

return hc
/*****
/*
/*01* SUBROUTINE-NAME: Sysinfo_Handler_2
/*
/*01* DESCRIPTIVE-NAME: SYSINFO Report Handler - Sample 2
/*
/*01* FUNCTION:
/*      This subroutine provides a sample to process the
/*      SYSINFO report table data
/*
/*01* OPERATION:
/*      1. loops through the SYSINFO report table
/*      2. checks if Response Time for DMN002 (TSO short)
/*          is > 1.00 s
/*      3. if yes, issues WTO message RMF202I and sets
/*          the hardcopy request to 1
/*
*****/
Sysinfo_Handler_2: Procedure

tabnam = "ERBSYST3"

name_1 = "DMN002"              /* Set the search argument */
respt_text = "Response Time:"
respt_limit = "1.00"           /* Set the threshold */
respt_msgid = "RMF202I 3B:"
limit = " Limit:"

hc = 0

```

```

found = 0
"TBQUERY" tabnam "ROWNUM(excpnum)"
If rc = 0 Then
    "TBTOP" tabnam
If rc = 0 Then
    "TBSKIP" tabnam

If rc = 0 Then
Do While (rc = 0 & found = 0)          /* Loops through table rows      */

    If (SUBSTR(sysnamvc,1,6) = name_1) Then
    Do                                  /* Argument found                  */
        found = 1
        If (SUBSTR(sysrspvc,1,4) >= respt_limit) Then
        Do                            /* Threshold exceeded ?          */

            msgtext = name_1 respt_text sysrspvc limit respt_limit
            wtomsg = respt_msgid msgtext
            "SELECT PGM(ERBCSWTO) PARM("wtomsg")"
            hc = 1
        End
    End
    "TBSKIP" tabnam
End
Else return 12

return hc

```

Part 8. Appendixes

Appendix. Accessibility

Accessible publications for this product are offered through the z/OS Information Center.

If you experience difficulty with the accessibility of any z/OS information, please send a detailed message to mhvrdfs@us.ibm.com or to the following mailing address:

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

Accessibility features

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully. The major accessibility features in z/OS enable users to:

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

Using assistive technologies

Assistive technology products, such as screen readers, function with the user interfaces found in z/OS. Consult the assistive technology documentation for specific information when using such products to access z/OS interfaces.

Keyboard navigation of the user interface

Users can access z/OS user interfaces using TSO/E or ISPF. Refer to *z/OS TSO/E Primer*, *z/OS TSO/E User's Guide*, and *z/OS ISPF User's Guide Vol I* for information about accessing TSO/E and ISPF interfaces. These guides describe how to use TSO/E and ISPF, including the use of keyboard shortcuts or function keys (PF keys). Each guide includes the default settings for the PF keys and explains how to modify their functions.

Dotted decimal syntax diagrams

Syntax diagrams are provided in dotted decimal format for users accessing the z/OS Information Center using 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 can be considered as 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 your 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, you know that 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 can be used 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 giving 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, this indicates a reference that is defined elsewhere. The string following the % symbol is the name of a syntax fragment rather than a literal. For example, the line 2.1 %OP1 means that you should refer to separate syntax fragment OP1.

The following words and symbols are used next to the dotted decimal numbers:

- ? means an optional syntax element. A dotted decimal number followed by the ? 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 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.
- ! means a default syntax element. A dotted decimal number followed by the ! symbol and a syntax element indicates 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 same dotted decimal number can specify a ! 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 this example, if you include the FILE keyword but do not specify an option, default option KEEP will be applied. A default option also applies to the next higher dotted decimal number. In this example, if the FILE keyword is omitted, 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 only applies 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.

- * means a syntax element that can be repeated 0 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, and 3 STATE, you know that you can include HOST, STATE, both together, or nothing.

Note:

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 could write HOST STATE, but you could not write HOST HOST.
 3. The * symbol is equivalent to a loop-back line in a railroad syntax diagram.
- + means a syntax element that must be included one or more times. A dotted decimal number followed by the + symbol indicates that this 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 only repeat a particular item if it is the only item with that dotted decimal number. The + symbol, like the * symbol, is equivalent to a loop-back line in a railroad syntax diagram.

Notices

This information was developed for products and services offered in the U.S.A. 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 give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte 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.

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

Site Counsel
IBM Corporation
2455 South Road
Poughkeepsie, NY 12601-5400
USA

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

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

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.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

COPYRIGHT LICENSE:

This information might contain 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.

Policy for unsupported hardware

Various z/OS elements, such as DFSMS, HCD, 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 (<http://www.ibm.com/software/support/systemsz/lifecycle/>)
- For information about currently-supported IBM hardware, contact your IBM representative.

Exploitation of the Flash Express feature

IBM intends to provide exploitation of the Flash Express feature on IBM zEnterprise EC12 (zEC12) and zBC12 servers with CFLEVEL 19 for certain coupling facility list structures in the first half of 2014. This new function is designed to allow list structure data to migrate to Flash Express memory as needed, when the consumers of data do not keep pace with its creators for some reason, and migrate it back to real memory to be processed. When your installation uses WebSphere® MQ for z/OS Version 7 (5655-R36), this new capability is expected to provide significant buffering against enterprise messaging workload spikes and provide support for storing large amounts of data in shared queue structures, potentially allowing several hours' data to be stored without causing interruptions in processing. In addition, z/OS V2R1 Resource Measurement Facility (RMF) is planned to provide measurement data and reporting capabilities for Flash Express when it is used with coupling facilities. Information about externals and interfaces that are related to this planned capability are being made available in z/OS V2R1 for early planning and development purposes only.

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](http://www.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 www.ibm.com/legal/copytrade.shtml (<http://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.

Microsoft, 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

ES Expanded storage

ESCON multiple image facility (EMIF)

A facility that allows channels to be shared among PR/SM logical partitions in an ESCON environment.

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.

expanded storage (ES)

An extension of processor storage.

Optional high-speed storage that transfers 4KB pages to and from central storage.

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 System z Application Assist Processor (zAAP)
 A special purpose processor configured for running Java programming on selected zSeries machines.

IBM System 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 *System 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 System z Application Assist Processor.
- zIIP** see IBM System z Integrated Information Processor.

Index

A

- accessibility 345, 399
 - contact IBM 399
 - features 399
- ACTIVE (active non-interactive sessions)
 - parameter
 - DISPLAY command 55
- ADD|DEL suboption
 - Monitor III data gatherer session 121
 - Monitor III DATASET option 121
- address space activity
 - selection conditions 224
- address space resource data
 - ARD/NOARD 187
 - Monitor II background session (ARD/NOARD) 189
 - Monitor II display session (ARD) 189
 - Postprocessor
 - report generation 224
- address space resource data by jobname
 - Monitor II background session (ARDJ/NOARDJ) 190
 - syntax summary 187
 - Monitor II display session (ARDJ) 190
 - syntax summary 187
 - Postprocessor
 - report generation 224
- address space SRM data
 - Monitor II background session (ASRM/NOASRM) 190
 - syntax summary 188
 - Monitor II display session (ASRM) 190
 - syntax summary 188
 - Postprocessor
 - report generation 225
- address space SRM data by jobname
 - Monitor II background session (ASRMJ/NOASRMJ) 190
 - syntax summary 188
 - Monitor II display session (ASRM) 190
 - syntax summary 188
 - Monitor II display session (ASRMJ) 190
 - syntax summary 188
 - Postprocessor
 - report generation 225
- address space state data
 - Monitor II background session (ASD/NOASD) 190
 - syntax summary 188
 - Monitor II display session (ASD) 190
 - syntax summary 188
 - Postprocessor
 - report generation 224
- address space state data by jobname
 - Monitor II background session (ASDJ/NOASDJ) 190
 - syntax summary 188
 - Monitor II display session (ASDJ) 190
- address space state data by jobname (*continued*)
 - syntax summary 188
 - Postprocessor
 - report generation 225
- administration of RMF 13
- AIX systems
 - monitoring with RMF XP 8
- ALL
 - Postprocessor
 - report generation 224
- ALL (all non-interactive sessions)
 - parameter
 - DISPLAY command 56
- Analysis
 - on DataView bar chart 353
- analysis dialog 363
- Application Transparent Transport Layer Security 280
- applications
 - providing data for 7
- ARD (address space resource data)
 - Monitor II display session 189
- ARD/NOARD (address space resource data) 187
 - Monitor II background session 189
- ARD|NOARD (address space resource data)
 - Postprocessor
 - report generation 224
- ARDJ (address space resource data by jobname)
 - Monitor II display session 190
 - syntax summary 187
- ARDJ/NOARDJ (address space resource data by jobname)
 - Monitor II background session 190
 - syntax summary 187
- ARDJ|NOARDJ (address space resource data by jobname)
 - Postprocessor
 - report generation 224
- ASCH address spaces 204
- ASD (address space state data)
 - Monitor II display session 190
 - syntax summary 188
- ASD/NOASD (address space state data)
 - Monitor II background session 190
 - syntax summary 188
- ASD|NOASD (address space state data)
 - Postprocessor
 - report generation 224
- ASDJ (address space state data by jobname)
 - Monitor II display session 190
 - syntax summary 188
- ASDJ/NOASDJ (address space state data by jobname)
 - Monitor II background session 190
 - syntax summary 188

- ASDJ|NOASDJ (address space state data by jobname)
 - Postprocessor
 - report generation 225
- ASMERRS (bad slots on local page data set)
 - Trace Activity report variable 97
- ASMIORQC (I/O requests completed and returned)
 - Trace Activity report variable 97
- ASMIORQR (I/O request received)
 - Trace Activity report variable 97
- ASMNVC (total local slots for non-VIO private area pages)
 - Trace Activity report variable 97
- ASMSLOTS (total slots in open local page data sets)
 - Trace Activity report variable 98
- ASMVSC (total local slots allocated for VIO private area pages)
 - Trace Activity report variable 98
- ASRM (address space SRM data)
 - Monitor II display session 190
 - syntax summary 188
- ASRM/NOASRM (address space SRM data)
 - Monitor II background session 190
 - syntax summary 188
- ASRM|NOASRM (address space SRM data)
 - Postprocessor
 - report generation 225
- ASRMJ (address space SRM data by jobname)
 - Monitor II display session 190
 - syntax summary 188
- ASRMJ/NOASRMJ (address space SRM data by jobname)
 - Monitor II background session 190
 - syntax summary 188
- ASRMJ|NOASRMJ (address space SRM data by jobname)
 - Postprocessor
 - report generation 225
- assistive technologies 399
- AT-TLS 280
- attention (PA1) key
 - Monitor III reporter session
 - STOP mode 152
- automatic customization
 - Monitor III reporter session
 - description 162
 - option set/system ID relation 162
- automatic update 184

B

- background session 172
 - starting Monitor II 61
- bar chart
 - DataView 353

- bar chart context menu 353
- batch session
 - Monitor III 382
- behavior of RMF spreadsheet
 - macros 328
- bottlenecks and exceptions
 - observing 75
- BPX.DAEMON FACILITY resource 18
- BPX.SMF
 - READ access for GPM4CIM 284
- BPX.WLMSEVER 18
- BREF (backward reference) command
 - Monitor III reporter session
 - description 147
 - parameter exceptions 149
 - RANGE value 150
 - syntax 148
 - without parameters 150

C

- CACHDET (cache details) 164
- CACHE
 - Monitor III data gatherer session
 - description 111
- cache activity
 - Monitor I session
 - (CACHE/NOCACHE) 85
 - summary 83
 - Postprocessor
 - report generation 225
 - SMF record type 74-5
 - overview/exception control statements 261
- CACHE/NOCACHE (cache activity)
 - Monitor I session 85
 - summary 83
- CACHE|NOCACHE (cache activity)
 - Postprocessor
 - report generation 225
- CACHSUM (cache summary) 164
- CANCEL command
 - Monitor II ISPF session 179
 - Monitor III reporter session
 - description 151
 - resetting entries 155
 - system command 47
- cataloged procedure
 - Monitor III data gatherer session
 - RMFGAT 53
 - Postprocessor 204
 - RMF (Resource Measurement Facility)
 - starting non-TSO/E session 22
- CCVCPUCT (number of online CPUs)
 - Trace Activity report variable 98
- CCVEJST (low priority user job step time)
 - Trace Activity report variable 98
- CCVENQCT (number of users
 - non-swappable for enqueue reasons)
 - Trace Activity report variable 98
- CCVRBSTD (recent base time of day)
 - Trace Activity report variable 98
- CCVRBSWT (recent base system wait time)
 - Trace Activity report variable 98
- CCVUTILP (system CPU utilization)
 - Trace Activity report variable 98
- central processor unit (CPU) activity
 - Monitor I session (CPU/NOCPU)
 - summary 83
 - Postprocessor
 - overview/exception control statements 238
 - report generation 225
 - SMF record type 70-1 238
- central storage/processor/SRM activity
 - Monitor II background session
 - (SRCS/NOSRCS) 198
 - syntax summary 188
 - Monitor II display session
 - (SRCS) 198
 - syntax summary 188
 - Postprocessor
 - report generation 229
- CF|NOCF (coupling facility activity)
 - Postprocessor
 - report generation 233
- CFACT (coupling facility activity) 164
- CFDETAIL (coupling facility)
 - Monitor III data gatherer session
 - description 111
 - effect on system performance 111
- CFLEVEL 19
 - Flash Express 405
- CFOVER (coupling facility
 - overview) 164
- CFSYS (coupling facility systems) 164
- CHAN/NOCHAN (channel path activity)
 - Monitor I session 85
 - summary 83
- CHAN|NOCHAN (channel path activity)
 - Postprocessor
 - report generation 225
- CHANNEL (channel path activity) 164
 - Monitor II display session 191
- channel path activity 169
 - Monitor I session (CHAN/NOCHAN)
 - summary 83
 - Monitor II background session
 - (CHANNEL/NOCHANNEL) 191
 - syntax summary 188
 - Monitor II display session
 - (CHANNEL) 191
 - syntax summary 188
 - Postprocessor
 - report generation 225
 - SMF record
 - overview/exception control statements 256
- CHANNEL/NOCHANNEL (channel
 - path activity)
 - Monitor II background session 191
 - syntax summary 188
- CHANNEL|NOCHANNEL (channel
 - path activity)
 - Postprocessor
 - report generation 225
- Client/Server control task 385
- Client/Server enabling 373
- CLISTS
 - setting up 39
- Collect.bat 324
- COLOR (color graphic option panel)
 - command
 - Monitor III reporter session
 - panel display 158
- command
 - syntax diagrams xv
- commands
 - Monitor II session
 - report commands 187
 - session commands 175
 - Monitor III reporter session
 - session commands 146
- COMMANDS command
 - Monitor II ISPF session 179
- common storage
 - detail report
 - Monitor I VSTOR/NOVSTOR
 - option 103
 - summary report
 - Monitor I VSTOR/NOVSTOR
 - option 103
- configuration file
 - gpm4cim 278
- console output from OPTIONS
 - Monitor III data gatherer session 116
- control panel 358
- control statement
 - Postprocessor
 - examples 234
- control statements
 - Postprocessor 213
- coupling facility activity
 - Postprocessor
 - report generation 233
 - SMF record type 74-4
 - overview control statements 259
- CPC (CPC capacity) 164
- CPU activity
 - Monitor I session (CPU/NOCPU) 86
- CPU/NOCPU (central processor unit
 - activity)
 - Monitor I session 86
 - summary 83
- CPU|NOCPU (central processor unit
 - activity)
 - Postprocessor
 - report generation 225
- create report from current session 182
- CreateOvwWSet.bat 326
- CreateRptWSet.bat 325
- cross platform monitoring
 - RMF XP 275
- cross-system coupling facility (XCF) 166
 - Postprocessor
 - EXRPTS control statement 230
- CRYPTO/NOCRYPTO (cryptographic
 - hardware activity)
 - Monitor I session 86
 - summary 83
- CRYPTO|NOCRYPTO (cryptographic
 - hardware activity)
 - Postprocessor
 - report generation 225
- cryptographic hardware activity
 - Monitor I session
 - (CRYPTO/NOCRYPTO) 86

- cryptographic hardware activity
(*continued*)
 - Postprocessor
 - overview/exception control statements 244
 - report generation 225
 - SMF record type 70-2 244
- CURRENT (current range data) command
 - Monitor III reporter session description 151
- current session
 - create report from 182
- cursor-sensitive control
 - Monitor III reporter session description 163
- customizing
 - system environment 15
- CYCLE (cycle length)
 - Monitor I session 86
 - overhead reduction 105
 - summary 84
 - undesirable results 104
 - Monitor III data gatherer session description 112
 - effect on system performance 112

D

- D (delta mode) command
 - Monitor II display session 182
- data collection
 - Monitor II background session
 - DELTA/NODELTA option 176, 182
 - Monitor II display session
 - DELTA option 182
 - skipped interval
 - Monitor I session 91
 - statistical sampling
 - Monitor I session 104
- data gatherer
 - Monitor III session
 - modifying 64
- data gatherer options
 - Monitor I 85
- data gathering capabilities 81
- data index (DI) screen
 - field descriptions 142, 143
 - report limitations 142
 - screen display
 - with preallocated data sets 142
 - without preallocated data sets 141
- data set
 - ddname
 - Monitor I session 24
 - Monitor II session 24
 - preallocation
 - Monitor I session 23
 - Monitor II session 23
- data set support 109
- daylight saving time 125
- data sets
 - Postprocessor output 208
- DATASET (record samples) 120

- DATASET (record samples) (*continued*)
 - ADD|DEL suboption
 - Monitor III data gatherer session 121
 - Monitor III data gatherer session description 112
 - START|STOP suboption
 - Monitor III data gatherer session 121
 - SWITCH|NOSWITCH suboption
 - Monitor III data gatherer session 121
 - WHOLD suboption
 - Monitor III data gatherer session 122
- DataView
 - bar chart 353
 - creating 352
 - creating from a Resource 348
 - metrics list 357
 - paste to a PerfDesk 350
 - remove Series from 353
 - working with 348
- DataView context-menu 354
- DATE (report period) control statement
 - Postprocessor 213, 214
- daylight saving time 54, 125
- DCB (data set control block)
 - characteristics
 - message and report data set 24
- ddname (data definition name)
 - Monitor I session 24
 - Monitor II session 24
 - Postprocessor 210
- ddnames
 - XML output 212
- DDS
 - daylight saving time 54
 - HTTP API to retrieve Postprocessor data 20
 - Postprocessor API 20
 - selective deactivation of Monitor III reports 30
 - sysplex-wide management 53
- DDS (Distributed Data Server)
 - customization 28
 - START command 53
- deactivation of Monitor III reports 30
- default session options 84
- defining SMF synchronization 31
- DELAY 164
- delay monitoring 11
- DELAYJ (delays for specified job) 164
- DELTA
 - Monitor II display session 182
- delta mode 182
- delta mode (D command)
 - Monitor II display session 182
- delta report mode 174
- DELTA/NODELTA
 - Monitor II background session 182
 - syntax summary 176
- DELTA|NODELTA 213
 - Postprocessor 215
- DEV (device activity table report)
 - Monitor II display session 191
 - syntax summary 188

- DEV (device delays) 164
- DEV/NODEV (device activity table report)
 - Monitor II background session 191
 - syntax summary 188
- DEV|NODEV (device activity table report)
 - Postprocessor
 - report generation 226
- device activity
 - Monitor II background session row report (DEVV/NODEVV)
 - syntax summary 188
 - Monitor II background session table report (DEV/NODEV)
 - syntax summary 188
 - Monitor II display session row report (DEVV)
 - syntax summary 188
 - Monitor II display session table report (DEV)
 - syntax summary 188
 - SMF record type 74-1
 - overview/exception control statements 258
- device number (NMBR/NONMBR)
 - sub-option
 - Monitor I session device activity option 87
- DEVICE/NODEVICE (I/O device activity)
 - classes sub-option 87
 - Monitor I session 86
 - overhead reduction 105
 - summary 83
 - NMBR/NONMBR sub-option 87
 - SG/NOSG sub-option 87
- DEVICE|NODEVICE (I/O device activity)
 - Postprocessor
 - report generation 226
- DEVJ (device delays job) 164
- DEVJ (resource-oriented device delays) 164
- DEVV (device activity row report)
 - Monitor II display session 192
 - syntax summary 188
- DEVV/NODEVV (device activity row report)
 - Monitor II background session 192
 - syntax summary 188
- DEVV|NODEVV (device activity row report)
 - Postprocessor
 - report generation 227
- DI (data index) 164
- DI (data index) command
 - Monitor III reporter session
 - report limitations 142
 - screen display 141, 143
- dialog
 - filter 360
- dialog button area 359
- DINTV (duration interval) control statement 215
 - Postprocessor 213

- dispatching priority
 - Monitor I session 91
 - RMF 22
- DISPLAY command
 - session command 55
- display session 167, 189
- displaying
 - options in effect for a session 60
- Distributed Data Server
 - customization 28
 - PassTicket support 19
 - RACF access 19
 - selective deactivation of Monitor III reports 30
 - START command 53
 - sysplex-wide management 53
- Distributed Data Server (DDS)
 - setting up for z/OS 26
- Distributed Data Server options 28
- Domino server activity
 - report generation 227
 - SMF records 202
- DOMINO/NODOMINO (Domino server activity)
 - report generation 227
- DSND (data set delays) 164
- DSNJ (data set delays - Job report) 165
- DSNV (data set delays - Volume report) 165
- duplicate connections
 - prevention of 382
- duration interval 215
 - DINTV control statement 213
 - Postprocessor
 - maximum length 217
 - relation to report interval 216
- duration report 204
 - Postprocessor 213
 - reporting period specification 230
- duration reports 201, 215

E

- ENCLAVE (enclave activity) 165
- ENQ (enqueue delays) 165
- ENQ/NOENQ (enqueue activity)
 - Monitor I session
 - description 88
 - summary 83
- ENQ|NOENQ (enqueue activity)
 - Postprocessor
 - report generation 227
- ENQJ (enqueue delays job) 165
- ENQR (resource-oriented enqueue delays) 165
- enqueue activity
 - Monitor I session (ENQ/NOENQ)
 - description 88
 - summary 83
 - Monitor II background session (SENQ|NOSENQ) 196
 - syntax summary 188
 - Monitor II display session (SENQ) 196
 - syntax summary 188
 - Postprocessor
 - report generation 227
- enqueue activity (*continued*)
 - SMF record type 77
 - overview/exception control statements 271
- enqueue contention
 - Monitor II background session (SENQ|NOSENQ)
 - Postprocessor report generation 228
- Enqueue, storage, and SRM reports 169
- Enterprise Disk Systems report
 - report generation 227
- environment file
 - gpm4cim.env 278
- ERB3RP3I - Monitor III Reporter Phase 3
 - installer 387
- ERB9R2SW 298
- ERBMFRES
 - removing 17
- ERBR3SYS 393
- ERBR3WFX 386
- ERBRMF00 (parmlib member)
 - Monitor I session
 - contents 36
- ERBRMF00 (Parmlib member)
 - Monitor I session 83
- ERBRMF01 (parmlib member)
 - Monitor II session
 - contents 37
- ERBRMF01 (Parmlib member)
 - Monitor II background session
 - contents 180
 - use 189
- ERBRMF02 (parmlib member)
 - Monitor I session
 - contents 37
- ERBRMF03 (parmlib member)
 - Monitor II session
 - contents 38
- ERBRMF04 39
- ERBRMF04 (parmlib member)
 - Monitor III session
 - contents 39
- ERBRMF04 (Parmlib member)
 - Monitor III data gatherer session 124
 - content 114
 - session option specification 109
- ERBS2V 135
- ERBS2V CLIST 135
- ERBV2S CLIST 134
- ERBVSEDF CLIST 25
- ESS activity
 - Monitor I session (ESS/NOESS)
 - description 89
- ESS statistics
 - OVW (overview condition) control statement 267
 - report generation 227
- ESS/NOESS (Enterprise Storage Server)
 - Monitor I session
 - summary 83
- ESS/NOESS (ESS activity)
 - Monitor I session
 - description 89
- ESS|NOESS (ESS statistics)
 - report generation 227

- ETOD (exception time of day) control statement
 - Postprocessor 213, 218
- Excel macros 339
- EXCEPT (exceptional value) control statement
 - Postprocessor 213, 218
 - cache activity (type 74-5 SMF record) 261
 - channel path activity (type 73 SMF record) 256
 - CPU activity (type 70-1 SMF record) 238
 - cryptographic hardware activity (type 70-2 SMF record) 244
 - device activity (type 74-1 SMF record) 258
 - enqueue activity (type 77 SMF record) 271
 - examples 220
 - ficon director activity (type 74-7 SMF record) 267
 - I/O queuing activity (type 78-3 SMF record) 272
 - page data set activity (type 75 SMF record) 271
 - paging activity (type 71 SMF record) 248
 - virtual storage activity (type 78-2 SMF record) 272
- EXCEPT control statement
 - syntax 219
- exception report 204
 - Postprocessor 213
 - EXCEPT control statement 218
 - exceptions list 238, 244
 - report period 218
 - specification 219, 220
- EXITS/NOEXITS (user exit)
 - Monitor I session
 - description 89
 - summary 84
- EXITS|NOEXITS (user exit)
 - Postprocessor 213, 220
- EXRPTS (exception report) control statement
 - Postprocessor 213, 220

F

- F (frame display command) 180
- F (frame) command
 - Monitor II display session 180
- FACILITY class
 - RACF resource profiles 20
- FCD/NOFCD (FICON director activity)
 - Monitor I session
 - description 90
 - summary 83
- FCD|NOFCD (FICON Director Activity)
 - report generation 227
- ficon director activity
 - SMF record type 74-7
 - overview/exception control statements 267

- FICON director activity
 - Monitor I session (FCD/NOFCD) description 90
- FICON Director Activity report generation 227
- file with overview control statements
 - OVW file 316, 318, 320
- filter dialog 360
- FIND command
 - Monitor II ISPF session 181
- FIND/RFIND (find) command
 - Monitor III reporter session description 151
- Flash Express
 - statement of direction 405
- frame display command (F) 180
- FREF (forward reference) command
 - Monitor III reporter session description 147
 - parameter exceptions 149
 - RANGE value 150
 - syntax 148
 - without parameters 150

G

- GDDM (Graphical Data Display Manager)
 - Monitor III reporter session graphic mode 152
- global performance data control authority 17
- glossary 407
- GMT-Offset
 - Sysplex dialog field 347
- GO command
 - Monitor II ISPF session 184
- GO mode
 - introduction 140
 - Monitor III reporter session command description 152
- GPM4CIM
 - configuration file 278
 - diagnostic aids 285
 - how to start 281
 - post-installation steps 280
 - READ access to BPX.SMF 284
 - RMF XP started task 276
 - security setup 280
 - setup after SMP/E 280
- gpm4cim configuration file 278
- GPM4CIM options 28
- GPM4CIM server 276
- GPM4CIM started task 276
- gpm4cim.env
 - environment file 278
- GPMSEVERE (Distributed Data Server) 53
- GPMSEVERE options 28
- GPMSEVERE00 (parmlib member) 28
- GRAPHIC (graphic mode) command
 - Monitor III reporter session description 152
- graphic report
 - Monitor III reporter session activating 152
 - printing 154

- graphic report (*continued*)
 - Monitor III reporter session (*continued*) switching to table report 156
- GROUP (group response time) 165
- grouping Monitor II and Monitor III users 41
- GUI session
 - WFEX report 383

H

- H (hardcopy mode) command
 - Monitor II display session 182
- HARDCOPY (print screen/report) command
 - Monitor III reporter session description 153
- hardcopy mode 182
 - for ISPF and TSO/E sessions 182
- HFS (UNIX file system statistics)
 - Monitor II display session 193
 - syntax summary 188
 - Postprocessor report generation 227
- HFS|NOHFS (HFS statistics)
 - Postprocessor report generation 227
- HFSNAME (UNIX file system statistics)
 - Monitor III data gatherer session description 113
- hierarchical file system statistics 169
- High Performance FICON for System z 257
- homepage, RMF 130
- how to add remote resources 313
- HSM (hierarchical storage manager delays) 165
- HSMJ (hierarchical storage manager delays job) 165
- HTTP server activity
 - HTTP/NOHTTP report option 227
 - SMF records 202

I

- I/O device activity 169, 188
 - Monitor I session (DEVICE|NODEVICE) 86
 - classes sub-option 87
 - NMBR/NONNMBR sub-option 87
 - SG/NOSG sub-option 87
 - summary 83
 - Monitor I session (DEVICE|NODEVICE)
 - Postprocessor report generation 226
 - Monitor I session (I/O DEVICE|NODEVICE) overhead reduction 105
 - Monitor II background session row report (DEVV|NODEVV) 192
 - Monitor II background session row report (DEVV|NODEVV)
 - Postprocessor report generation 227

- I/O device activity (*continued*)
 - Monitor II background session table report (DEV|NODEV) 191
 - Monitor II background session table report (DEV|NODEV)
 - Postprocessor report generation 226
 - Monitor II display session row report (DEVV) 192
 - Monitor II display session table report (DEV) 191
 - SMF record type 74-1
 - overview/exception control statements 258
- I/O queuing activity 92, 169
 - IOQ|NOIOQ
 - Postprocessor report generation 227
 - Monitor I session
 - description 91
 - Monitor I session (IOQ|NOIOQ)
 - classes sub-option 92
 - NMBR/NONNMBR sub-option 91
 - summary 83
 - Monitor II background session (IOQUEUE|NOIOQUEUE)
 - Postprocessor report generation 227
 - Monitor II session (IOQUEUE|NOIOQUEUE) 188, 194
 - SMF record type 78-3
 - overview/exception control statements 272
- IBM System z Application Assist Processor (zAAP)
 - definition 409
- IBM System z Integrated Information Processor (zIIP)
 - definition 409
- IBM z/OS Management Facility
 - overview 369
- ICU (Interactive Chart Utility)
 - Monitor III reporter session
 - directory panel 154
 - ICU command description 154
 - printing graphic reports 154
- id (RMF identifier) 55
- IEFPARM DD statement
 - definition 23
- IEFPROC EXEC statement
 - definition 22
- IFASMFDP program
 - creation of non-VSAM data set 202
- IFASMFDP program
 - Postprocessor
 - creation of non-VSAM data set 202
- ILOCK (IRLM long lock detection)
 - Monitor II display session 193
- in reporting hierarchy 71
- INITIAL (option set default values) options
 - Monitor III reporter session 161
- installation directory 323
 - default 323

- INTERVAL (interval length)
 - Monitor I session
 - description 90
 - overhead reduction 105
 - summary 84
 - undesirable results 104, 105
 - Monitor II background session
 - conflicting options 51
- interval measurement
 - SINTV (interval length)
 - Monitor II background session 176, 185
- interval report 204
 - Monitor I session (REPORT/NOREPORT)
 - description 95
 - summary 84
 - undesirable results 105
 - Monitor II background session (REPORT/NOREPORT) 183
 - syntax summary 176
 - Postprocessor
 - EXRPTS control statement 220
 - reporting period specification 230
 - REPORT/NOREPORT option
 - Monitor II background session 51, 198
- interval reports 201
- interval synchronization
 - Monitor I session (SYNC/NOSYNC)
 - description 96
 - summary 84
 - Monitor III data gatherer session option 118
- IOQ/NOIOQ (I/O queuing activity)
 - Monitor I session
 - classes sub-option 92
 - description 91
 - NMBR/NONMBR sub-option 91
 - summary 83
- IOQ|NOIOQ (I/O queuing activity) 227
- IOQUEUE (I/O queuing activity) 165
- IOQUEUE/NOIOQUEUE (I/O queuing activity)
 - Monitor II session 188, 194
- IOQUEUE|NOIOQUEUE (I/O queuing activity)
 - Postprocessor report generation 227
- IOSUB (I/O subsystem)
 - Monitor III data gatherer session
 - description 113
- IRLM long lock detection
 - Monitor II display session (ILOCK) 193
 - syntax summary 188
- IRLM long lock reporting 48
- ISPF (Interactive System Product Facility)
 - Monitor II display session 168
 - Monitor III reporter session 133
- ISPF session
 - hardcopy mode 182
- ISPF session panels
 - scrolling through 180

J

- JCL
 - starting Postprocessor with 205
- JCL (job control language)
 - Monitor III data gatherer session 53
 - RMF cataloged procedure 22
- Jclgen.bat 323
- JES delays
 - Monitor III reporter session command 165
- JES not active 45, 52
- JESJ (JES delays for specified job) 165
- JOB command
 - Monitor III reporter session 165
- job or resource view in resource hierarchy 71

K

- keyboard
 - navigation 399
 - PF keys 399
 - shortcut keys 399

L

- language
 - Monitor III reporter session (LO option) 146
- LANGUAGE (language option panel)
 - command
 - Monitor III reporter session description 160
- library list
 - Monitor II (LLI) description 194
- Library List and OPT Settings Selection Menu 170
- Linux systems
 - monitoring with RMF XP 8
- LLI (library list)
 - Monitor II display session 194
- LLI (program library list)
 - Monitor II display session syntax summary 188
- LOCK (lock reporting)
 - Monitor III data gatherer session description 113
- LOCKSP (Spin locks) 165
- LOCKSU (Suspend locks) 165
- long-term data gathering
 - with Monitor I 4
 - with Monitor III 4
- long-term overview reporting 6
- long-term performance analysis
 - RMF XP
 - long-term performance analysis 12
 - with RMF XP 12
- long-term reporting with the Postprocessor 201
- Lotus Domino server activity
 - SMF records 202
- Lotus Domino Server report 227

- LSCTCNT (logically swapped users for terminal wait)
 - Trace Activity report variable 98
- LSCTMTE (think time for logical swap candidate)
 - Trace Activity report variable 98

M

- M (menu) command
 - Monitor II display session 177
- macros
 - for Report Working Sets 327
- major name
 - Monitor I session
 - ENQ/NOENQ option 88
- MASTER
 - Monitor III data gatherer session 114
- Master Gatherer system 114
- master system 53
- MCVFCNT (pages to be stolen by force steal routine)
 - Trace Activity report variable 98
- MCVMGAGE (expanded storage migration age)
 - Trace Activity report variable 98
- MCVSBLETF (percentage of eligible storage that is actually fixed)
 - Trace Activity report variable 98
- MCVSIPR (common page-in rate)
 - Trace Activity report variable 98
- MCVSTCRI (highest system UIC)
 - Trace Activity report variable 98
- MCVTWSS (common target working set size)
 - Trace Activity report variable 98
- MEMBER (Parmlib member)
 - Monitor I session
 - description 93
 - summary 84
 - Monitor II background session 180
 - syntax summary 176
 - Monitor II background session option processing order 62
 - Monitor III data gatherer session description 114
- menu display (M) command for Monitor II sessions 177
- menu items
 - Monitor II display session 177
- menus
 - Monitor III Primary Menu 136
- Message Browser 368
- messages
 - Monitor III start 135
- metric types 357
- metrics list
 - for DataView 357
- MFEnnnnn (Monitor I session) ddname
 - definition 24
- MFEnnnnn ddname
 - Postprocessor 210
- MFMESSAGE ddname
 - definition 24
- MFINPUT data set
 - Postprocessor
 - DD statement 204

MFPMMSGDS (message output) ddname
 Postprocessor 210
 MFRnnnnn (Monitor I session) ddname
 definition 24
 MFRnnnnn ddname
 Postprocessor 210
 MFXnnnnn (exception interval report
 data) ddname
 Postprocessor 210
 migrating from previous releases 15
 migration
 of Working Sets 323
 minor name
 Monitor I session
 ENQ/NOENQ option 88
 MINTIME (interval length)
 Monitor III data gatherer session
 description 115
 MM (menu default options) command
 Monitor II display session 177
 MNTJ (operator mount job delays) 165
 mode
 Monitor III reporter session
 default 152
 GO display command 152
 GRAPHIC command 152
 STOP mode 152
 MODIFY command 124
 Monitor III data gatherer session 124
 session command
 example 55
 modifying RMF session
 options 55
 syntax 55
 system command
 non-interactive session 52
 passing DISPLAY session
 command 55
 passing MODIFY session
 commands 55
 passing STOP session
 command 56
 Monitor I 60
 data gatherer options 85
 long-term data gathering with 4
 Monitor I session
 cache activity
 CACHE/NOCACHE option
 summary 83
 central processor unit activity
 CPU/NOCPU option
 summary 83
 channel path activity
 CHAN/NOCHAN option
 summary 83
 controlling 24
 CPU/NOCPU option 86
 cryptographic hardware activity
 CRYPTO/NOCRYPTO 86
 CRYPTO/NOCRYPTO option
 summary 83
 cycle length
 CYCLE 86
 CYCLE option summary 84
 overhead reduction 105
 data collection
 skipped interval 91

Monitor I session (*continued*)
 data collection (*continued*)
 statistical sampling 104
 data recording
 RECORD/NORECORD option
 description 95
 RECORD/NORECORD option
 summary 84
 device activity
 DEVICE/NODEVICE option
 description 86
 DEVICE/NODEVICE option
 summary 83
 overhead reduction 105
 dispatching priority 91
 enqueue activity
 ENQ/NOENQ option
 description 88
 ENQ/NOENQ option
 summary 83
 ESS activity
 ESS/NOESS option
 description 89
 ESS/NOESS option summary 83
 FICON director activity
 FCD/NOFCD option
 description 90
 FCD/NOFCD option
 summary 83
 I/O device activity
 DEVICE/NODEVICE option
 description 86
 DEVICE/NODEVICE option
 summary 83
 I/O queuing activity
 classes sub-option 92
 device number suboption 91
 IOQ/NOIOQ option
 description 91
 IOQ/NOIOQ option summary 83
 interval length
 INTERVAL option description 90
 INTERVAL option summary 84
 overhead reduction 105
 interval report 213
 REPORT/NOREPORT option
 description 95
 REPORT/NOREPORT option
 summary 84
 interval synchronization
 SYNC/NOSYNC option
 description 96
 SYNC/NOSYNC option
 summary 84
 NOZZ option 45
 options 83
 command source (input
 source) 93
 conflicting options 104
 OPTIONS/NOOPTIONS option
 description 93
 OPTIONS/NOOPTIONS option
 summary 84
 specifying 51
 undesirable results 24, 104

Monitor I session (*continued*)
 page data set activity
 PAGESP/NO PAGESP option
 description 95
 PAGESP/NO PAGESP option
 summary 83
 Parmlib member
 MEMBER option description 93
 MEMBER option summary 84
 printed report
 SYSOUT option description 97
 SYSOUT option summary 84
 private area data summary and detail
 reports 103
 report interval
 interval length option
 summary 84
 session duration
 STOP/NOSTOP option
 description 96
 STOP/NOSTOP option
 summary 84
 session identifier 52
 session option
 CACHE/NOCACHE (cache
 activity) 83
 CHAN/NOCHAN (channel path
 activity) 83
 CPU/NOCPU (central processor
 unit activity) 83
 CRYPTO/NOCRYPTO
 (cryptographic hardware
 activity) 83
 CYCLE (cycle length) 84, 104, 105
 DEVICE/NODEVICE (I/O device
 activity) 83, 86, 105
 displaying 56
 ENQ/NOENQ (enqueue
 activity) 83, 88
 ESS/NOESS (Enterprise Storage
 Server) 83
 ESS/NOESS (ESS activity) 89
 EXITS/NOEXITS (user exit
 routine) 84, 89
 FCD/NOFCD (FICON director
 activity) 83, 90
 INTERVAL (interval length) 84,
 90, 104, 105
 IOQ/NOIOQ (I/O queuing
 activity) 83, 91
 MEMBER (Parmlib member) 84,
 93
 modifying 55
 OPTIONS/NOOPTIONS (option
 list) 84, 93
 PAGESP/NO PAGESP (page data
 set activity) 83, 95
 PAGING/NO PAGING (system
 paging activity) 83, 95
 RECORD/NORECORD (record to
 SMF data set) 84, 95
 REPORT/NOREPORT (interval
 report) 84, 95
 STOP/NOSTOP (session
 duration) 84, 96
 SYNC/NOSYNC (interval
 synchronization) 30, 84, 96

- Monitor I session (*continued*)
 - session option (*continued*)
 - SYSOUT (sysout class) 84, 97
 - TRACE/NOTRACE (trace activity) 83, 97
 - VSTOR/NOVSTOR (virtual storage activity) 83, 103
 - WKLD/NOWKLD (workload activity) 83, 104
 - SMF record
 - synchronization with Monitor III 30
 - writing data to 84, 95
 - starting 45
 - concurrent with RMF 45
 - stopping 56, 57
 - sysout class
 - SYSOUT option description 97
 - SYSOUT option summary 84
 - system activity measured 10
 - system paging activity
 - PAGING/NOPAGING option description 95
 - PAGING/NOPAGING option summary 83
 - trace activity
 - TRACE/NOTRACE 97
 - TRACE/NOTRACE option summary 83
 - user exit routine
 - EXITS/NOEXITS option description 89
 - EXITS/NOEXITS option summary 84
 - virtual storage activity
 - VSTOR/NOVSTOR option description 103
 - VSTOR/NOVSTOR option summary 83
 - VMGUEST/NOVMGUEST 104
 - workload activity
 - WKLD/NOWKLD option description 104
 - WKLD/NOWKLD option summary 83
- Monitor I session option
 - CACHE/NOCACHE (cache activity) 85
 - CHAN/NOCHAN (channel path activity) 85
 - CPU/NOCPU (CPU activity) 86
 - CRYPTO/NOCRYPTO (cryptographic hardware activity) 86
 - CYCLE (cycle length) 86
 - default 83
 - summary 83
- Monitor II 169
 - background session 63
 - displaying available commands 179
 - finding character string in report, 181
 - ISPF interface to 129
 - leaving options unchanged 179
 - Library List and OPT Settings Selection Menu 170
 - OPT information 170
- Monitor II (*continued*)
 - recalling previously displayed data 181
 - repeat FIND command 182
 - report commands 188
 - report modes 174
 - reports
 - structure of 172
 - resetting default options 179
 - snapshot monitoring with 4
 - snapshot reporting with 6
 - sorting reports 181
- Monitor II and Monitor III users
 - grouping 41
- Monitor II background session
 - controlling 189
 - interval report 213
 - maximum number 51
 - modifying options 55
 - options
 - conflicting 51, 198
 - specifying 189
 - Parmlib member
 - ERBRMF01 189
 - member name generating 180
 - Postprocessor
 - SESSION control statement 230
 - user reports 229
 - session identifier 52
 - starting 52, 61, 189
 - stopping 57
 - system activity measured 10
 - user reports 187
- Monitor II background session options 187
- Monitor II data gatherer session
 - session option 107
- Monitor II display session 167, 189
 - commands
 - delta mode (D) 182
 - frame (F) 180
 - hardcopy mode (H) 182
 - menu default options (MM) 177
 - menu of available reports (M) 177
 - print (P) 186
 - program function (PF) key assignment 185
 - recall previous report (R) 181
 - stop (Z) 187
 - timed update (T) 184
 - display menu 171
 - ending 187
 - ISPF session commands
 - automatic report refresh (GO) 184
 - cancel options (CANCEL) 179
 - display commands (COMMANDS) 179
 - find string (FIND) 181
 - print report page (PRINT) 186
 - repeat find (RFIND) 182
 - reset options (RESET) 179
 - set report options (RO) 177
 - sort report (SORT) 181
 - specify system (SYSTEM) 186
 - stop (RETURN) 177
 - program function (PF) key
 - override assigned keys 185
- Monitor II display session (*continued*)
 - system activity measured 10
 - TSO/E session
 - controlling 187
- Monitor II display session options 187
- Monitor II session
 - starting 52
- Monitor II snapshot reporting 167
- Monitor II User Reports panel 170
- Monitor III
 - BATCH reporter 385
 - before starting 133
 - enabling RMF 133
 - ERBR3WFX 386
 - help facility 163
 - ISPF interface to 129
 - long-term data gathering with 4
 - messages during start 135
 - Primary Menu 136
 - Reporter session
 - data set allocation 133
 - short-term data collection with 3
 - short-term interactive performance analysis with 5
 - STOP mode 140
 - sysplex considerations 135
- Monitor III batch session 382
- Monitor III data gatherer session
 - cataloged procedure 53
 - console output 56
 - controlling recording 120
 - delay information 11
 - description 109
 - local storage buffer size 119
 - modifying options 55
 - option list 115
 - Parmlib member
 - description 109, 114
 - member name generation 114
 - option 114
 - RMFGAT 24
 - sample collection 120
 - time interval 115
 - user defined data set 112
 - session identifier 53
 - session option
 - CACHE (cache subsystem) 111
 - CFDETAIL (coupling facility) 111
 - CYCLE (cycle length) 112
 - DATASET (record samples) 112, 120
 - defaults 110
 - description 111
 - HFSNAME (UNIX file system statistics) 113
 - IOSUB (I/O subsystem) 113
 - LOCK (lock reporting) 113
 - MASTER 114
 - MEMBER (Parmlib member) 114
 - MINTIME (interval length) 115
 - OPD (OMVS process data) 115
 - OPTIONS (option list) 115
 - RESOURCE (JES resource) 117
 - SGSPACE 117
 - specification 109
 - STOP (session duration) 117

Monitor III data gatherer session
(continued)

- session option (continued)
 - SYNC (interval synchronization) 118
 - SYSOUT (sysout class) 118
 - VSAMRLS 118
 - WSTOR (buffer size) 119
 - ZFS 120
 - ZIIPUSE 120
 - starting 51, 53
 - stopping 56, 57
 - system activity measured 10
 - VSAM data set support 25
 - data set selection 121
 - defining data sets 25
 - modifying 124
 - page release control 122
 - starting 121, 122
 - stopping 124
 - ZFS/NOZFS 120
 - ZIIPUSE/NOZIIPUSE 120
- Monitor III gatherer session
- data set support 24
- Monitor III report data
- exception handling for 383
- Monitor III Reporter Phase 3 installer,
ERB3RP3I 387
- Monitor III reporter session
- BREF (backward reference) command
 - description 147
 - parameter exceptions 149
 - RANGE value 150
 - summary 146
 - syntax 148
 - without parameters 150
 - CANCEL (cancel options) command
 - description 151
 - resetting entries 155
 - summary 146
 - changing options 146
 - COLOR (color graphics panel)
 - command 158
 - color graphic options
 - description 146
 - field description 159
 - selection panels 158, 159
 - commands 166
 - session commands 146
 - COMMANDS (help menu for
commands) command
 - description 151
 - summary 146
 - controlling 133
 - CURRENT (current range data)
 - command
 - description 151
 - summary 146
 - cursor-sensitive control
 - description 163
 - data set support
 - alternate system support 134
 - data set preallocation 134
 - DI screen 141, 143
 - delay information 11
 - DI screen 141, 143

Monitor III reporter session (continued)

- FIND/RFIND (find) command
 - description 151
 - summary 146
- FREF (forward reference) command
 - description 147
 - parameter exceptions 149
 - RANGE value 150
 - summary 146
 - syntax 148
 - without parameters 150
- GO (GO mode) command
 - command summary 146
 - description 152
- GRAPHIC (graphic mode) command
 - description 152
 - summary 146
- graphic report
 - switching to table report 156
- HARDCOPY (print screen/report)
 - command
 - description 153
 - summary 146
- ICU (Interactive Chart Utility)
 - command
 - description 154
 - summary 146
- LANGUAGE (language option panel)
 - command 160
- language (LO) option 146
 - description 160
- option set selection menu 146
 - default values (INITIAL) 161
 - description 161
- OPTSET (option set)
 - command 161
 - option 146
- primary menu 133
- program function (PF) key
 - building commands 155
 - changing settings 155
 - PFK command summary 146
- report options
 - description 160
 - RO command 146
- RESET (reset options) command
 - description 155
 - summary 146
- RETRIEVE (display last command)
 - command
 - description 156
 - summary 146
- ROPTIONS (report options panel)
 - command 160
- session commands 146
- session options
 - automatic customization 162
 - creating a new option set 162
 - deleting an option set 162
 - making an option set current 162
 - selecting an option set 162
- SESSION (session options panel)
 - command 156
- SO command 146
- specification 156
- summary 156

Monitor III reporter session (continued)

- STOP mode
 - description 152
 - system activity measured 10
 - tabular report
 - cursor-sensitive control 163
 - displaying 156
 - switching to graphic report 156
 - TABULAR command
 - description 156
 - TABULAR command
 - summary 146
 - TOGGLE (toggle reports) command
 - description 156
 - summary 146
 - user written report 156
- Monitor III reports
- selective deactivation from DDS 30
- MPFLSTxx
- examples 375, 378, 382
 - exception-initiated monitoring without
MPFLSTxx functions 384
- MSGJ (operator message job delays) 165

N

- name 362
- naming parmlib option members 33
- navigating from the Primary Menu 136
- navigation
 - keyboard 399
- NMBR/NONMBR (device number)
 - sub-option
 - Monitor I session device activity
option 87, 91
- non-interactive session
 - description 51
 - options
 - modifying 55
 - specifying 51
 - session control
 - session option processing 59
 - session identifiers
 - displaying 55
 - session options
 - displaying 55
 - starting 51, 53
 - stopping 56
- non-TSO/E session
 - starting
 - RMF cataloged procedure 22
- Notices 403
- NOZZ (no Monitor I session)
 - Monitor I session 45

O

- observing bottlenecks and exceptions 75
- observing response times 73
- OMDGAMRE (messages on the action
message retention facility queue)
 - Trace Activity report variable 98
- OMDGCMDI (commands issued per
second)
 - Trace Activity report variable 98

- OMDGOEB (operator reply entries on system reply queue)
 - Trace Activity report variable 98
- OMDGWQEB (WTO queue elements on system output queue)
 - Trace Activity report variable 99
- OMDGWTLI (write-to-logs issued per second)
 - Trace Activity report variable 99
- OMDGWTOI (total lines of messages, write-to-operators issued per second)
 - Trace Activity report variable 99
- OMVS kernel activity
 - Postprocessor OMVS|NOOMVS report option 228
- online reporting environment 129
- OPD (OMVS process data) 165
 - Monitor III data gatherer session description 115
- Open RMF Overview Spreadsheets
 - usage example 339
- OPT Settings report 188
 - OPT parameter settings 195
- option list
 - Monitor I session (OPTIONS/NOOPTIONS)
 - description 93
 - summary 84
 - Monitor II background session 178
 - syntax summary 176
 - Postprocessor 212
- options
 - default session 84
 - Monitor I data gatherer 85
 - Monitor I session 83
 - conflicting options 104
 - undesirable results 24, 104
 - Monitor II background session 177, 187
 - Monitor II display session 177
 - Monitor III data gatherer session
 - description 111
 - session options
 - MODIFY command 55
- OPTIONS (option selection menu)
 - command
 - Monitor III reporter session
 - description 146
- OPTIONS/NOOPTIONS (option list)
 - Monitor I session
 - description 93
 - summary 84
 - Monitor II background session 178
 - syntax summary 176
 - Monitor III data gatherer session
 - description 115
- OPTSET (option set)
 - Monitor III reporter session
 - command 161
 - Monitor III reporter session
 - option 146
- output data sets
 - Postprocessor 208
- overview and exception conditions 238
- OVERVIEW control statement
 - Postprocessor 214, 221

- overview control statements 295
- Overview Records 295
- overview report
 - Postprocessor 213
 - conditions list 238, 244
 - report generation 221
 - type selection
 - specifying (OVERVIEW control statement) 221
 - specifying (OVW control statement) 222
- Overview Reports 295
- Overview Working Set 302, 311
- OVW (overview condition) control statement 222
 - Postprocessor
 - cache activity (type 74-5 SMF record) 261
 - channel path activity (type 73 SMF record) 256
 - coupling facility activity (type 74-4 SMF record) 259
 - CPU activity (type 70-1 SMF record) 238
 - cryptographic hardware activity (type 70-2 SMF record) 244
 - device activity (type 74-1 SMF record) 258
 - enqueue activity (type 77 SMF record) 271
 - ESS statistics (type 74-8 SMF record) 267
 - ficon director activity (type 74-7 SMF record) 267
 - I/O queuing activity (type 78-3 SMF record) 272
 - page data set activity (type 75 SMF record) 271
 - paging activity (type 71 SMF record) 248
 - PCIE Function statistics (type 74-9 SMF record) 270
 - virtual storage activity (type 78-2 SMF record) 272
 - workload activity (type 72-3 SMF record) 251
- OVW control statement
 - Postprocessor 214

P

- P (print) command
 - Monitor II display session 186
- PA1 (attention) key
 - Monitor III reporter session
 - STOP mode 152
- page data set activity
 - Monitor I session (PAGESP/NOPAGESP)
 - description 95
 - summary 83
 - Monitor I session (PAGESP|NOPAGESP)
 - Postprocessor report
 - generation 228
 - Monitor II background session (PGSP/NOPGSP) 195
- page data set activity (*continued*)
 - Monitor II background session (PGSP|NOPGSP) 228
 - Monitor II display session (PGSP) 195
 - Monitor II session (PGSP/NOPGSP) 188
 - SMF record type 75
 - overview/exception control statements 271
- PAGESP/NOPAGESP (page data set activity)
 - Monitor I session
 - description 95
 - summary 83
- PAGESP|NOPAGESP (page data set activity)
 - Postprocessor
 - report generation 228
- paging activity
 - Monitor II background session (SPAG/NOSPAG)
 - syntax summary 188
 - Monitor II display session (SPAG)
 - syntax summary 188
- PAGING/NOPAGING (system paging activity)
 - Monitor I session
 - description 95
 - summary 83
- PAGING|NOPAGING (system paging activity)
 - Postprocessor
 - report generation 228
- panel
 - Monitor II User Reports 170
- parmlib member
 - ERBRMF00 (Monitor I session) 36
 - ERBRMF01 (Monitor II session) 37
 - ERBRMF02 (Monitor I session) 37
 - ERBRMF03 (Monitor II session) 38
 - ERBRMF04 (Monitor III session) 39
- Parmlib member
 - member name generating
 - Monitor I session 93
 - Monitor II background session 180
 - Monitor I session (MEMBER)
 - description 93
 - summary 84
 - Monitor II background session (MEMBER) 180
 - syntax summary 176
- parmlib naming 34
- parmlib option members
 - naming 33
- PassTicket support 19
- paste DataView to a PerfDesk 350
- PCIE Activity Report (PCIE|NOPCIE) 228
- PCIE Function statistics
 - OVW (overview condition) control statement 270
- PCIE|NOPCIE (PCIE Activity Report)
 - Postprocessor 228
- PerfDesk
 - change name of 350

- PerfDesk (*continued*)
 - closing 349
 - creating 348
 - paste DataView to 350
 - working with 348
- PerfDesk folder
 - creating 348
- performance data
 - accessing across the sysplex 8
 - sysplex
 - accessing data across 8
- performance goals
 - watching 73
- performance management 65
- Performance Monitoring
 - metrics list 357
- performance problem
 - definition of 77
- persistence capability
 - for Linux and AIX performance monitoring 277
- PFK (program function key)
 - Monitor III reporter session 154
- PGSP (page data set activity)
 - Monitor II display session 195
- PGSP/NOPGSP (page data set activity)
 - Monitor II background session 195
 - Monitor II session 188
- PGSP|NOPGSP (page data set activity) 228
- plot area, Series 357
- Plot/Save Series dialog 357
- Postprocessor
 - cache activity (type 74-5 SMF record) 261
 - cataloged procedure 204
 - channel path activity (type 73 SMF record) 256
 - control statements 204, 212, 213
 - examples 234
 - coupling facility activity (type 74-4 SMF record) 259
 - CPU activity (type 70-1 SMF record) 238
 - cryptographic hardware activity (type 70-2 SMF record) 244
 - data set
 - preallocation 210
 - description 204
 - device activity (type 74-1 SMF record) 258
 - duration interval 215
 - duration report 235
 - duration reports 215
 - enqueue activity (type 77 SMF record) 271
 - error conditions 212
 - exception report 218
 - ficon director activity (type 74-7 SMF record) 267
 - I/O queuing activity (type 78-3 SMF record) 272
 - IFASMFDP program
 - creation of non-VSAM data set 202

- Postprocessor (*continued*)
 - IFASMFDP program
 - creation of non-VSAM data set 202
 - input
 - log streams 207
 - SMF data sets 207
 - ISPF interface to 129
 - JCL
 - ddname summary 210
 - long-term overview reporting 6
 - long-term reporting with 201
 - options list 212
 - Options panel 209
 - output data sets 208
 - Overview Records 301
 - page data set activity (type 75 SMF record) 271
 - paging activity (type 71 SMF record) 248
 - report generation 223
 - reporting period 212
 - specification 230
 - return codes 213
 - setup panel 206
 - SMF record 202
 - DD statement for spanned records 204
 - multiple data sets 202
 - multiple systems 202
 - non-VSAM data set 202
 - starting
 - data set preallocation 210
 - sample JCL 204
 - starting from ISPF 206
 - starting with JCL 205
 - virtual storage activity (type 78-2 SMF record) 272
 - workload activity (type 72-3 SMF record) 251
- Postprocessor control statements 213
- Postprocessor HTTP API of the DDS 20
- Postprocessor Input pane 207
- Postprocessor reporting 201
- Postprocessor reports
 - XML output viewing 289
- Postprocessor XML reports
 - sorting 290
- Postprocessor XML Toolkit
 - installation 289
 - installation and use 289
 - using 290
- PPEXTnnn (exception report) ddname
 - Postprocessor 210
- PPORPnnn (overview report) ddname
 - Postprocessor 210
- PPOVWREC (overview records) ddname
 - Postprocessor 211
- PPRPTS (combined interval report) ddname
 - Postprocessor 210
- PPSUMnnn (summary report) ddname
 - Postprocessor 211
- PPXRPTS (combined exception report) ddname
 - Postprocessor 211

- PPXSRPTS (sysplex reports) ddname
 - Postprocessor 211
- PR/SM
 - processor activity 86
- prevention of duplicate connections 382
- Primary Menu
 - navigating from 136
- PRINT command
 - Monitor II ISPF session 186
- printed report
 - Monitor II background session
 - SYSOOT (sysout class)
 - option 176, 184
 - SYSTEM (system ID) option 186
 - Monitor II display session
 - P (print) command 186
 - Monitor III reporter session
 - graphic reports 154
 - screens and tabular reports 153
- priority for RMF
 - specifying 23
- private area
 - detail report
 - Monitor I VSTOR/NOVSTOR
 - option 103
 - summary report
 - Monitor I VSTOR/NOVSTOR
 - option 103
- PROC (processor delays) 165
- procedure
 - RMF procedure 22, 53
- processing session options
 - Monitor I session 63
 - Monitor II background session 63
 - when you modify a session
 - Monitor II background session 62
- PROCJ (processor job delays) 165
- PROCU (processor usage) 165
- program default
 - Monitor I session options 83
 - Monitor II session options 107
- program function (PF) key
 - Monitor II display session
 - override assigned keys 185
 - Monitor III reporter session 154
 - building commands 155
 - changing settings 155
- program library lists
 - Monitor II display session (LLI)
 - syntax summary 188
- program properties table (PPT) 16
- Programming Interface Information 405

Q

- QSCJ (quiescent jobs) 165
- quick start 130

R

- R (recall) command
 - Monitor II display session 181
- RACF access to the Distributed Data Server 19
- RACF FACILITY
 - BPX.WLMSEVER 18

- RACF resource profiles 20
- RAXESCT (common storage pages on expanded storage)
 - Trace Activity report variable 99
- RAXFMCT (frames allocated to common)
 - Trace Activity report variable 99
- RCEAEC (total expanded storage E frames currently on ESTE queue)
 - Trace Activity report variable 99
- RCEAECLO (available expanded storage low threshold)
 - Trace Activity report variable 99
- RCEAECOK (available expanded storage satisfactory threshold)
 - Trace Activity report variable 99
- RCEAFC (total frames currently on all available frame queues)
 - Trace Activity report variable 99
- RCEAFCLO (available central storage low threshold)
 - Trace Activity report variable 99
- RCEAFCOK (available central storage low threshold)
 - Trace Activity report variable 99
- RCEBELFX (total number of fixed pages below 16 megabytes)
 - Trace Activity report variable 99
- RCECOMPI (common area pages paged-in)
 - Trace Activity report variable 99
- RCECOMPO (common area pages paged-out)
 - Trace Activity report variable 99
- RCEDFRS (number of times a deferred frame allocation has been satisfied)
 - Trace Activity report variable 99
- RCEESINU (in-use expanded storage frames)
 - Trace Activity report variable 99
- RCEESREA (non-VIO pages read from expanded storage)
 - Trace Activity report variable 99
- RCEESWRT (pages written to expanded storage frames)
 - Trace Activity report variable 99
- RCEHSPEM (total hiperspace pages migrated from auxiliary storage) expanded to
 - Trace Activity report variable 100
- RCEHSPER (total hiperspace pages read from expanded to central storage)
 - Trace Activity report variable 100
- RCEHSPEW (total hiperspace pages written from central to expanded storage)
 - Trace Activity report variable 100
- RCEHSPPI (total hiperspace pages paged in from auxiliary storage)
 - Trace Activity report variable 100
- RCEHSPPO (total hiperspace pages paged out to auxiliary storage)
 - Trace Activity report variable 100
- RCELPAPI (PLPA and PLPA directory pages paged-in)
 - Trace Activity report variable 100
- RCEMVBEL (pages moved from below 16M in central storage)
 - Trace Activity report variable 100
- RCENWSF (total number secondary and non-working pages migrated)
 - Trace Activity report variable 100
- RCEPAGMV (number of times a frame was moved)
 - Trace Activity report variable 100
- RCEPOOL (frames currently available to the system)
 - Trace Activity report variable 100
- RCESPFER (frames available by swap-out without requiring I/O)
 - Trace Activity report variable 100
- RCESWPPI (total number of pages requiring I/O to swap-in)
 - Trace Activity report variable 100
- RCESWPPO (pages requiring I/O to swap-out)
 - Trace Activity report variable 100
- RCETOTFX (total number of pages currently fixed)
 - Trace Activity report variable 100
- RCETOTPI (total pages paged-in excluding swap-in and VIO page-in)
 - Trace Activity report variable 100
- RCETOTPO (total number pages paged out)
 - Trace Activity report variable 100
- RCEVIOME (VIO pages written to expanded storage)
 - Trace Activity report variable 100
- RCEVIOMG (VIO pages migrated from expanded storage to paging data sets)
 - Trace Activity report variable 100
- RCEVIOPI (total number VIO pages paged-in)
 - Trace Activity report variable 100
- RCEVIOPO (total number VIO pages paged-out)
 - Trace Activity report variable 101
- RCEVIORE (VIO reads from extended storage)
 - Trace Activity report variable 101
- RCEWSDNE (primary working set pages migrated to auxiliary storage)
 - Trace Activity report variable 101
- RCVAFQA (average available frame count)
 - Trace Activity report variable 101
- RCVAVQC (AVQ low count)
 - Trace Activity report variable 101
- RCVCPUA (CPU usage average * 16)
 - Trace Activity report variable 101
- RCVFXIOP (percentage central storage fixed or allocated)
 - Trace Activity report variable 101
- RCVMFXA (average number fixed frames)
 - Trace Activity report variable 101
- RCVPAGRT (total paging rate)
 - Trace Activity report variable 101
- RCVPTR (paging rate)
 - Trace Activity report variable 101
- RCVSWPTM (time used by ASM to transfer pages to/from data set)
 - Trace Activity report variable 101
- RCVUICA (UIC average)
 - Trace Activity report variable 101
- recall display command 181
- RECORD/NORECORD (record to SMF data set)
 - Monitor I session
 - description 95
 - summary 84
 - Monitor II background session 186
 - conflicting options 51, 198
 - syntax summary 176
- REGION parameter
 - parameter adjustment 22
- remote resources
 - how to add 313
- report (RO) option
 - Monitor III reporter session 146
- report commands
 - Monitor II 188
- report from the current session
 - creating 182
- report interval
 - Monitor I session
 - length 90
 - synchronization 30, 84, 96
 - Postprocessor
 - relation to duration interval 216
- Report Listings
 - viewing 307
- report mode
 - delta 172, 174
 - total 172, 174
- report types 304
- Report Working Set 302, 311
- Report Working Sets
 - macros for 327
- REPORT/NOREPORT (interval report)
 - Monitor I session
 - description 95
 - summary 84
 - undesirable results 105
 - Monitor II background session 183
 - conflicting options 51, 198
 - syntax summary 176
- reporters
 - starting 129
- reporting
 - Postprocessor 201
- reporting across time zones, Sysplex 203
- reporting capabilities 127
- reporting hierarchy 70
- REPORTS control statement 223
 - Postprocessor 214
- reserve activity
 - Monitor II background session (SENQR/NOSENQR) 197
 - syntax summary 188
 - Monitor II display session (SENQR) 197
 - syntax summary 188
- RESET (reset options) command
 - Monitor III reporter session
 - description 155

- RESET command
 - Monitor II ISPF session 179
- Resource
 - creating a DataView from 348
 - expanding 347
- RESOURCE (JES resource)
 - Monitor III data gatherer session description 117
- Resource and Metric information 360
- resource directory 299
- resource name
 - format
 - Monitor I session 88
- Resources
 - how to work with 304
- response times
 - observing 73
- RETRIEVE (display last command)
 - command
 - Monitor III reporter session description 156
- return code
 - Postprocessor 213
- RETURN command
 - Monitor II ISPF session 177
- RFIND command
 - Monitor II ISPF session 182
- RLSDS (VSAM RLS activity by data set) 165
- RLSLRU (VSAM RLS activity LRU statistics) 165
- RLSSC (VSAM RLS activity by storage class) 165
- RMCAAWSC (APPC/MVS transaction scheduler wait swap count)
 - Trace Activity report variable 101
- RMCADWSC (detected wait physical swap count)
 - Trace Activity report variable 101
- RMCAEXSC (exchange on recommendation value swap count)
 - Trace Activity report variable 101
- RMCAFHLD (failed swaps because of outstanding HOLD SYSEVENT)
 - Trace Activity report variable 101
- RMCAICSC (improve central storage use)
 - Trace Activity report variable 101
- RMCAIPSC (improve system paging rate)
 - Trace Activity report variable 101
- RMCALWSC (long wait physical swap count)
 - Trace Activity report variable 101
- RMCAMRSC (make room to swap in a user swapped out too long)
 - Trace Activity report variable 101
- RMCANQSC (CPU enqueue exchange swap count)
 - Trace Activity report variable 101
- RMCAOISC (OMVS input wait)
 - Trace Activity report variable 101
- RMCAOOSC (OMVS output wait)
 - Trace Activity report variable 102
- RMCARSSC (central storage shortage swap count)
 - Trace Activity report variable 102
- RMCATISC (terminal input swap count)
 - Trace Activity report variable 102
- RMCATOSC (terminal output swap count)
 - Trace Activity report variable 102
- RMCATSSC (count of translation swaps)
 - Trace Activity report variable 102
- RMCAUSSC (unilateral swap out count)
 - Trace Activity report variable 102
- RMCAXSXC (auxiliary storage shortage swap count)
 - Trace Activity report variable 102
- RMCTTRPC (pages used for transaction elements)
 - Trace Activity report variable 102
- RMF
 - dispatching priority 22
 - overview of capabilities 1
 - REGION parameter 22
 - what you can gather 10
 - what you can report 10
- RMF (Resource Measurement Facility)
 - storing gatherer options 33
- RMF (Resource Measurement Facility)
 - cataloged procedure
 - JCL description 22
 - RMFGAT 53
 - session option processing 59
 - starting non-TSO/E session 22
 - Client/Server control task 385
 - CLISTs 39
 - controlling session options 64
 - dispatching priority 91
 - gathers data using three monitors 3
 - homepage 130
 - identifier 55
 - message help 130
 - non-interactive session
 - option specification 51
 - session control 59
 - starting 51, 53
 - stopping 56
 - online help 130
 - overhead reduction 105
 - Parmlib member 51
 - session identifier 55, 56
 - session option processing
 - conflicts 59
 - controlling 64
 - invalid option value 60
 - source order 59
 - sessions
 - control 51
 - description 51
 - starting 51
 - stopping 56
 - skipped data collection 91
 - specifying priority for 23
 - starting 45
 - session option processing 59
 - without concurrent Monitor I session 45
 - status 55
 - stopping 47
 - non-interactive session 56
 - tutorial 130
- RMF administration 13
- RMF Client/Server 373
- RMF Master Gatherer system 114
- RMF Performance Monitoring
 - help information 342
 - installation and setup 342
 - introduction 341
 - overview 341
 - prerequisites 342
 - prerequisites for the client 342
 - prerequisites for the z/OS sysplex 342
 - putting it all together 342
 - what is monitored 363
- RMF Performance Monitoring (RMF PM) 130
- RMF PM (RMF Performance Monitoring) 130
- RMF Postprocessor XML Toolkit
 - installation and use 289
- RMF procedure 22, 53
 - example 22
- RMF spreadsheet macros
 - behavior of 328
- RMF Sysplex Data Server
 - accessing data 9
 - activating 9
 - data flow 9
 - overview 8
 - programming interface services 9
 - services for Monitor II data 10
 - services for Monitor III data 10
 - services for SMF data 9
- RMF Sysplex Data Server's SMF
 - buffer 9, 20
- RMF XP
 - cross platform monitoring 275
 - GPM4CIM started task 276
 - monitoring AIX, Linux, or Windows systems 8
 - persistence capability 277
 - SMF record type 104 277
- RMF XP options 28
- RMFCS (RMF Client/Server)
 - components 384
 - control session 381
 - customization 378
 - installation 378
 - JCL procedures 384
 - message-initiated monitoring 375
 - overview 373
 - prerequisites 378
 - procedures 387
 - scenarios 375
 - START command 54, 381
 - usage considerations 382
- RMFCSC 385
- RMFGAT (Monitor III cataloged procedure) 53
- RMFM3B 385
- RMFPnnnn (Monitor II interval report)
 - ddname
 - Postprocessor 211
- RMFSCxx (Monitor I or Monitor II background session) ddname
 - definition 24
- RMFSTART CLIST 40
- RMFxxnnn (Monitor II background session) ddname
 - definition 24

- RO command
 - Monitor II ISPF session 177
- ROPTIONS (report options panel)
 - command
 - Monitor III reporter session 160
- RTOD (report time of day) control
 - statement
 - Postprocessor 214, 230

S

- sample interval 355
- scroll forward 180
- scrolling through ISPF session
 - panels 180
- SDELAY|NOSDELAY (serialization
 - delay)
 - Postprocessor 228
- SDEVICE|NOSDEVICE (shared device
 - activity)
 - Postprocessor
 - report generation 233
- SDS (sysplex data server activity)
 - Monitor II display session
 - syntax summary 188
- SDS (Sysplex Data Server activity)
 - Monitor II display session 195
- security
 - for z/OS UNIX 18
- sending comments to IBM xix
- SENQ (enqueue activity)
 - Monitor II display session 196
 - syntax summary 188
- SENQ/NOSENQ (enqueue activity)
 - Monitor II background session 196
 - syntax summary 188
- SENQ|NOSENQ (system enqueue
 - activity)
 - Monitor II background session
 - Postprocessor report
 - generation 228
- SENQR (reserve activity)
 - Monitor II display session 197
 - syntax summary 188
- SENQR/NOSENQR (reserve activity)
 - Monitor II background session 197
 - syntax summary 188
- SENQR|NOSENQR (system enqueue
 - reserve activity)
 - Postprocessor report generation 228
- SERBPWSV 298
- serialization delay reporting
 - (SDELAY|NOSDELAY) 228
- Series
 - add to a DataView 356
 - definition dialog 356
 - plot area 357
 - remove from DataView 353
 - settings for 353
 - time-stamps 356
 - value-list chart 353
- Seriesfind lowest value in
 - Series 353
- service level agreement 68
- session
 - create report from 182
 - displaying options in effect for 60

- SESSION (Monitor II background session)
 - control statement
 - Postprocessor 214, 230
- SESSION (session options panel)
 - command
 - Monitor III reporter session 156
- session (SO) option
 - Monitor III reporter session 146
- session command
 - DISPLAY 55
 - MODIFY command
 - example 55
 - modifying RMF session
 - options 55
 - syntax 55
 - START command
 - Monitor III data gatherer
 - session 122
 - non-interactive session 52, 53
- starting sessions 51
- STOP command 56
 - Monitor II background
 - session 187
 - Monitor III data gatherer
 - session 117
 - non-interactive session 56
- stopping sessions 51
- session control
 - Monitor II background session 189
 - non-interactive session 51, 53
 - session option processing 59
- overview 51
- starting 51
- session duration
 - Monitor I session (STOP/NOSTOP)
 - description 96
 - summary 84
 - undesirable results 105
 - Monitor II background session
 - (STOP/NOSTOP) 187
 - conflicting options 51, 198
 - syntax summary 176
- session identifier
 - guidelines for assigning 52
 - Monitor I 24
 - STOP session command 56
 - use in DISPLAY session command 56
 - use in MODIFY session command 55
- session mode 182
- session option
 - format 52
 - MODIFY command 55
 - Monitor I session
 - displaying 56
 - specifying 52
- session options
 - default 84
- sets, Postprocessor output data 208
- setting up RMF CLISTs 39
- setup panel
 - Postprocessor 206
- SFCB 275
- SG/NOSG (storage group) sub-option
 - Monitor I session device activity
 - option 87
- SGSPACE/NOSGSPACE
 - data gatherer 117

- SGSPACE/NOSGSPACE (*continued*)
 - Monitor III data gatherer session 117
- shared device activity
 - Postprocessor
 - report generation 233
- short-term data collection with Monitor
 - III 3
- short-term interactive performance
 - analysis with Monitor III 5
- shortcut keys 399
- single-value chart 353
- SINTV (interval length)
 - Monitor II background session 185
 - conflicting options 198
 - syntax summary 176
- SLA 68
- Small Footprint CIM Broker (SFCB) 275
- SMCABFLS (records lost because of a
 - shortage of buffers)
 - Trace Activity report variable 102
- SMCABFWT (number of buffers written)
 - Trace Activity report variable 102
- SMCADSCT (records lost because of full
 - data set)
 - Trace Activity report variable 102
- SMCANMFL (current number of full
 - buffers)
 - Trace Activity report variable 102
- SMCARCWT (number of records written)
 - Trace Activity report variable 102
- SMF buffer 9, 20, 22, 47
- SMF dump data 300
- SMF record
 - Monitor I and Monitor III
 - synchronization 30
 - Monitor I session
 - RECORD/NORECORD option
 - description 95
 - RECORD/NORECORD option
 - summary 84
 - Monitor II background session 186
 - syntax summary 176
 - Postprocessor 202
 - spanned records 204
 - type 70-1 (CPU activity)
 - overview/exception control
 - statements 238
 - type 70-2 (Crypto Hardware activity)
 - overview/exception control
 - statements 244
 - type 71 (paging activity)
 - overview/exception control
 - statements 248
 - type 72-3 (workload activity)
 - overview control statements 251
 - type 73 (channel path activity)
 - overview/exception control
 - statements 256
 - type 74-1 (device activity)
 - overview/exception control
 - statements 258
 - type 74-4 (Coupling Facility activity)
 - overview control statements 259
 - type 74-5 (cache activity)
 - overview/exception control
 - statements 261

- SMF record (*continued*)
 - type 74-7 (ficon director activity)
 - overview/exception control statements 267
 - type 74-8 (Ess statistics) 267
 - type 74-9 (PCIE Function statistics) 270
 - type 75 (page data set activity)
 - overview/exception control statements 271
 - type 77 (enqueue activity)
 - overview/exception control statements 271
 - type 78-2 (virtual storage activity)
 - overview/exception control statements 272
 - type 78-3 (I/O queuing activity)
 - overview/exception control statements 272
- SMF record type 104 277
- authorize GPM4CIM 284
- SMF record type 104 collection
 - how to request 283
- SMF record writing
 - activating 30
 - suppressing 31
- SMF synchronization
 - defining 31
- SMFBUF option 22, 47, 56
- Snapshot monitoring with Monitor II 4
- snapshot reporting
 - with Monitor II 6
- SORT command
 - Monitor II ISPF session 181
- sorting XML reports 290
- SPACED (disk space) 165
- SPACEG (storage group space) 165
- SPAG (system paging activity)
 - Monitor II display session 198
 - syntax summary 188
- SPAG/NOSPAG (system paging activity)
 - Monitor II background session 198
 - syntax summary 188
- SPAG|NOSPAG (system paging activity)
 - Monitor II background session
 - Postprocessor report generation 229
- specifying priority for RMF 23
- spreadsheet
 - usage example 336
- spreadsheet macros
 - behavior of RMF 328
- Spreadsheet Reporter
 - concepts 293
 - getting more out of Postprocessor reports 294
 - how to use 299
 - installation 298
 - installation steps 298
 - introduction 7
 - macros 326
 - main dialog 302
 - overview 293
 - prerequisites 298
 - resource directory 299
 - resources 300
- Spreadsheet Reporter (*continued*)
 - specify Postprocessor report types 319
 - specifying general processing options 317
 - specifying reporting periods and duration intervals 322
 - usage example 336, 339
- spreadsheets
 - Spreadsheet Reporter 293
- Spreadsheets
 - creating 311
- SRCS (central storage/processor/SRM) activity
 - Monitor II display session 198
 - syntax summary 188
- SRCS/NOSRCS (central storage/processor/SRM) activity
 - Monitor II background session 198
 - syntax summary 188
- SRCS|NOSRCS (central storage/processor/SRM) activity
 - Postprocessor
 - report generation 229
- START command
 - Monitor III data gatherer session 124
 - session command 189
 - Monitor III data gatherer session 122
 - non-interactive session 52, 53
 - session option processing 59
 - system command 45
 - session option processing 59
- START|STOP suboption
 - DATASET option
 - Monitor III data gatherer session 121
 - Monitor III data gatherer session 121
- starting
 - Monitor II background session 61
- starting Monitor I session 60
- statement of direction
 - coupling facility list structures 405
- statistical sampling
 - Monitor I session 104
- status
 - displaying 55
- STOD (summary report time of day)
 - control statement
 - Postprocessor 214, 231
- STOP command
 - session command 56
 - Monitor III data gatherer session 117
 - system command 47, 56
- STOP mode
 - Monitor III 140
- STOP/NOSTOP (session duration)
 - data gatherer 117
 - Monitor I session
 - description 96
 - summary 84
 - undesirable results 105
 - Monitor II background session 187
 - conflicting options 51, 198
 - syntax summary 176
- STOP/NOSTOP (session duration) (*continued*)
 - Monitor III data gatherer session
 - description 117
- STOR (storage delays) 165
- storage group (SG/NOSG) sub-option
 - Monitor I session device activity option 87
- STORC (common storage) 165
- STORCR (common storage remaining) 165
- STORF (storage frames) 165
- storing data 5
- STORJ (storage job delays) 166
- STORM (storage usage by memory objects) 166
- STORR (resource-oriented storage delays) 166
- STORS (storage delays summary) 166
- subtype 60-64, Windows on System x performance data 283
- SUMMARY (summary report) control statement
 - Postprocessor 214, 231
- Summary of changes xxiv
- summary report 204
 - Postprocessor 213
 - report generation 231
 - reporting period specification 231
- SWITCH|NOSWITCH suboption
 - DATASET option
 - Monitor III data gatherer session 121
 - Monitor III data gatherer session 121
- symbolic substitution 34
- SYNC/NOSYNC (interval synchronization)
 - Monitor I session
 - description 96
 - summary 84
 - synchronizing with Monitor III 30
 - Monitor III data gatherer session
 - description 118
- syntax diagrams
 - how to read xv
- SYS1.PROCLIB system procedure library
 - RMF cataloged procedure 22
- SYSENQ (sysplex enqueue delays) 166
- SYSID (system identifier) control statement
 - Postprocessor 214, 232
- SYSINFO (system information) 166
- SYSOUT (sysout class)
 - Monitor I session
 - description 97
 - summary 84
 - Monitor II background session
 - option 184
 - syntax summary 176
 - Monitor III data gatherer session
 - description 118
- SYSOUT (sysout class) control statement
 - Postprocessor 214, 232
- SYSOUT data set
 - message and report dynamic allocation 23

- sysplex
 - expanding 347
 - open and close 347
 - reporting across time zones 203
 - time-zone 346
- Sysplex
 - considerations for 25
- Sysplex Data Server
 - SMF buffer 9, 20
- sysplex data server activity
 - Monitor II display session (SDS)
 - syntax summary 188
- Sysplex Data Server activity
 - Monitor II display session (SDS) 195
- sysplex view 71
- SYSRPTS (sysplex reports) control statement
 - Postprocessor 214, 232
- SYSRTD (response time distribution) 166
- SYSUM (sysplex summary) 166
- SYSTEM (system ID)
 - Monitor II background session
 - option 186
 - Monitor II display session
 - command 186
- system activities measured
 - Monitor I session 10
 - Monitor II session 10
 - Monitor III session 10
 - Postprocessor 10
- system commands
 - CANCEL command 47
 - MODIFY command
 - non-interactive session 52
 - passing DISPLAY session
 - command 55
 - passing MODIFY session
 - commands 55
 - passing STOP session
 - command 56
 - START command 45
 - STOP command 47, 56
- system enqueue contention 228
- system enqueue reserve activity
 - Postprocessor report generation 228
- system environment
 - customizing 15
- system identifier
 - Postprocessor
 - SYSID control statement 232
- System Overview Report 330
- system paging activity 188
 - Monitor I session
 - (PAGING/NOPAGING)
 - description 95
 - summary 83
 - Monitor I session
 - (PAGING|NOPAGING)
 - Postprocessor report
 - generation 228
 - Monitor II background session
 - (SPAG/NOSPAG) 198
 - Monitor II background session
 - (SPAG|NOSPAG)
 - Postprocessor report
 - generation 229

- system paging activity (*continued*)
 - Monitor II display session
 - (SPAG) 198
 - SMF record type 71
 - overview/exception control
 - statements 248
- system view 71
- Systems
 - how to maintain 314
- SYSWKM (work manager delays) 166

T

- T (timed update) command
 - Monitor II display session 184
- T (tutorial) command
 - RMF interactive reporting
 - description 130
- TABULAR (tabular report) command
 - Monitor III reporter session
 - description 156
- tabular report
 - Monitor III reporter session
 - cursor-sensitive control 163
 - displaying 156
 - switching to graphic report 156
- this document
 - who should use xiii
- throughput definition 74
- time zones
 - sysplex reporting across 203
- time-stamps
 - Series 356
- time-zone 346
- timed update (T command)
 - Monitor II display session 184
- TOGGLE (toggle reports) command
 - Monitor III reporter session
 - description 156
- total report mode 174
- TRACE/NOTRACE (trace activity) 97
 - Monitor I session
 - summary 83
- TRACE|NOTRACE (trace activity)
 - Postprocessor
 - report generation 229
- trademarks 405
- transaction rate 74
- TSO/E session
 - ending 187
 - hardcopy mode 182
- tutorial
 - RMF interactive reporting
 - description 130

U

- UNIX System Services
 - access to UNIX System Services 18
 - OMVS kernel activity 228
 - OMVS process data 115
 - UNIX file system statistics 193
- USER (user exit) option
 - Monitor II display session 198
 - syntax summary 188
- USER command 139

- user exit routine
 - Monitor I session (EXITS/NOEXITS)
 - description 89
 - summary 84
 - Monitor II background session
 - (USER/NOUSER) 198
 - syntax summary 188
 - Monitor II display session
 - (USER) 198
 - syntax summary 188
 - Postprocessor
 - EXITS control statement 220
- user interface
 - ISPF 399
 - TSO/E 399
- user reports
 - Monitor II session
 - specifying for Postprocessor 229
- user-written reports
 - Monitor III reporter session
 - displaying 156
- USER/NOUSER (user exit)
 - Monitor II background session 198
 - syntax summary 188

V

- value-list chart 353
- viewing local Report Listings 307
- virtual storage activity
 - Monitor I session
 - (VSTOR/NOVSTOR)
 - description 103
 - summary 83
 - Postprocessor report generation 229
 - SMF record type 78-2
 - overview/exception control
 - statements 272
- VMGUEST/NOVMGUEST 84
 - Monitor I session 104
- VSAM (Virtual Storage Access Method)
 - data set 25
 - Monitor III data gatherer session
 - data set definition 25
 - data set selection 121
 - modifying 124
 - page release control 122
 - starting 121, 122
 - stopping 124
 - Monitor III reporter session
 - alternate system support 134
 - data set preallocation 134
 - DI screen 141, 143

- VSAMRLS
- Monitor III data gatherer session
 - description 118
- VSTOR/NOVSTOR (virtual storage activity)
- Monitor I session
 - description 103
 - summary 83
- VSTOR|NOVSTOR (virtual storage activity)
- Postprocessor report generation 229

W

- WFEX (workflow/exceptions) 166
- WFEX report
 - starting GUI session 383
- what's new 130
- who should use this document xiii
- WHOLD (page release storage value)
 - suboption
 - DATASET option
 - Monitor III data gatherer session 121
 - Monitor III data gatherer session 122
 - relationship to WSTOR suboption 122
- Windows on System x performance data:
 - subtype 60-64 283
- Windows systems
 - monitoring with RMF XP 8
- WKLD/NOWKLD (workload activity)
 - Monitor I session
 - description 104
 - summary 83
- WLM OPT parameters 195
- WLMGL (workload activity)
 - Postprocessor
 - SMF record type 72-3 251
- WLMGL|NOWLMGL (sysplex workload activity)
 - Postprocessor
 - report generation 233
- work scope 362
- work scope type 362
- Working Set
 - Overview Working Set 302, 311
 - Report Working Set 302, 311
- Working Set migration 323
- workload activity
 - Postprocessor
 - SMF record type 72-3 251
 - SYSRPTS control statement 233
 - SMF record type 72-3
 - overview control statements 251
- workstation
 - reporting capabilities on 287
- WSTOR (buffer size)
 - Monitor III data gatherer session
 - description 119
- WHOLD suboption
 - relationship in controlling page space 122

X

- XCF (cross-system coupling facility) 166
- XCF transport class
 - defining 16
- XCF|NOXCF (cross-system coupling facility)
 - Postprocessor
 - EXRPTS control statement 230
- XCFJ (cross system coupling facility job delays) 166
- XML output
 - ddnames 212
- XML Postprocessor reports
 - viewing in a browser 289

- XML Postprocessor reports (*continued*)
 - viewing with the HTTP API of the DDS 289
 - viewing with the Spreadsheet Reporter 289
- XML reports
 - sorting 290
- XPOVWRPT ddname 212
- XPRPTS ddname 212
- XPXSRPTS ddname 212

Z

- Z (stop) command
 - Monitor II display session 187
- z/OS Management Facility
 - overview 369
- z/OSMF
 - overview 369
 - Resource Monitoring plug-in 369
- ZFS/NOZFS 120
- ZFSACT (zFS activity) 166
- ZFSSUM (zFS file system summary) 166
- zHPF 257
- ZIIPUSE/NOZIIPUSE 120
- ZZ (Monitor I session identifier) 52



Product Number: 5650-ZOS

Printed in USA

SC34-2664-01

