OS/390

Resource Measurement Facility User's Guide

SC28-1949-04

OS/390

Resource Measurement Facility User's Guide

SC28-1949-04

Note

Before using this information and the product it supports, be sure to read the general information under “Notices” on page xvii.

Fifth Edition, March 1999

This is a major revision of SC28-1949-03.

This edition applies to Version 2 Release 7 of OS/390 (5647-A01) and to all subsequent releases and modifications until otherwise indicated in new editions or technical newsletters.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address below.

IBM welcomes your comments. A form for readers' comments may be provided at the back of this publication, or you may address your comments to the following address:

IBM DEUTSCHLAND ENTWICKLUNG GMBH
RMF DEVELOPMENT
SCHOENAICHER STRASSE 220
71032 BOEBLINGEN
GERMANY

If you prefer to send comments electronically, use one of the following methods:

FAX (RMF Development): Your International Access Code +49+7031+16+4240
IBM Mail Exchange: DEIBMRMF at IBMMAIL
Internet: deibmrmf@ibmmail.com

Internet

Visit our homepage at http://www.ibm.com/s390/rmf/

If you would like a reply, be sure to include your name, address, telephone number, or FAX number.

Make sure to include the following in your comment or note:

- Title and order number of this book
- Page number or topic related to your comment

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Part 3. Operation

Part 4. Performance Management

Part 5. Data Gathering Reference

Part 6. Reporting Reference

Part 7. Analysis on the Workstation

Figures

Notices

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates.

Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, programs, or services except those expressly designated by IBM, are the user's responsibility.

Any pointers in this publication to non-IBM Web sites are provided for convenience only, and do not in any manner serve as an endorsement of these Web sites. IBM accepts no responsibility for the content or use of non-IBM Web sites specifically mentioned in this publication or accessed through an IBM Web site that is mentioned in this publication.

IBM may have patents or pending patent applications covering subject matter 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 the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785, USA.

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

The following terms, **DENOTED BY AN ASTERISK (*)** when used for the first time in this publication, are trademarks of the IBM Corporation in the United States or other countries:

- CICS
- CUA
- DFSORT
- ESCON
- GDDM
- Hiperspace
- IBM
- IMS
- Language Environment
- Multiprise
- MVS/ESA
- OS/2
- OS/390
- Processor Resource/Systems Manager

- PR/SM
- RACF
- Resource Measurement Facility
- RMF
- S/390 Parallel Enterprise Server
- VTAM

UNIX is a registred trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

The following terms, **DENOTED BY TWO ASTERISKS (**)** when used for the first time in this publication, are trademarks of other companies:

- i486
- Lotus
- Microsoft
- Quattro
- Windows
- 1-2-3

About This Book

The Resource Measurement Facility* (RMF) is the strategic IBM product for performance management in an OS/390 host environment.

This book describes RMF, what it can do, and how to use RMF sessions. For information about analyzing the various reports that RMF produces, see *RMF Report Analysis*, and *RMF Performance Management Guide* for information on how to go about the task of performance management.

Who Should Use This Book

This book 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 an OS/390 system, this book assumes that the reader has extensive knowledge of the OS/390 system.

How This Book Is Organized

This book contains the following parts:

Part 1, Introduction
This chapter describes the different components of RMF, and explains how to use them for data gathering, data reporting, and performance management.

Part 2, Administration
This chapter 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
Here, you get information about operator tasks for starting, modifying, and stopping the different types of monitoring sessions.

Part 4, Performance Management
The chapters of 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
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.

Part 6, Reporting Reference
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

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 Spreadsheet Reporter (RMFPP) 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.

  The Spreadsheet Converter (RMF2SC) is part of the Spreadsheet Reporter for handling data from the RMF report displays or from report data sets (Monitor II, Monitor III, and Postprocessor) using techniques familiar to every spreadsheet user. You can find a description in the *RMF Programmer's Guide*.
- Performance Monitoring of OS/390 (PM of OS/390) provides an interface between the OS/2 workstation and the OS/390 sysplex that gives you the flexibility to create unique scenarios that monitor the performance of your system.
- 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 MVS systems in your network that have an APPC or TCP/IP connection configured to your PWS.

The OS/390 RMF Library

This table shows the shortened titles, the full titles, and the order numbers of the books in the RMF library for OS/390. This book uses the shortened titles when referring to other books.

Figure 0-1. RMF Library

Short Title Used in This Book	Title	Order Number
Books available as Hardcopy and Softcopy		
RMF User's Guide	*OS/390 RMF User's Guide*	SC28-1949
RMF Report Analysis	*OS/390 RMF Report Analysis*	SC28-1950
RMF Performance Management Guide	*OS/390 RMF Performance Management Guide*	SC28-1951
RMF Programmer's Guide	*OS/390 RMF Programmer's Guide*	SC28-1952
RMF Reference Summary	*OS/390 RMF Reference Summary*	SX22-0044
RMF Messages and Codes	*OS/390 RMF Messages and Codes*	GC28-1948
Softcopy documentation as part of the *OS/390 Collection* (SK2T–6700)		
RMF Diagnosis Guide	*OS/390 RMF Diagnosis Guide*	SC33-6592
RMF NewsFLASH	*OS/390 RMF NewsFLASH*	SC28-1986

Related Information

For additional information on OS/390, see the *OS/390 Information Roadmap*, GC28-1727.

Syntax Notation

This book uses “railway-line” syntax diagrams show you how to specify commands and options. The diagrams identify clearly the keywords and variables to be used, the required and optional parameters, and the defaults. They are not used for commands that consist of a single keyword with no operands.

How to Read the Diagrams

To read a syntax diagram, follow the path of the line, starting from left to right and moving from top to bottom.

- The ▶▶—— >>—— symbol indicates the beginning of a syntax diagram.
- The ——▶ ——> symbol, at the end of a line, indicates that the syntax diagram continues on the next line.
- The ▶—— >—— symbol, at the beginning of a line, indicates that a syntax diagram continues from the previous line.
- The ——▶◀ ——>< symbol indicates the end of a syntax diagram.

Syntax items (for example, a keyword or variable) may be:

- Directly on the line (required)
- Above the line (default)
- Below the line (optional)

Symbols

You **must** code these symbols exactly as they appear in the syntax diagram

#	Number sign
:	Colon
,	Comma
=	Equals sign
-	Hyphen
( )	Parenthesis
.	Period

Variables

Highlighted lowercase letters denote variable information that you must substitute with specific information. For example:

```
►►──┬─────────────────────┬──────────────────────────────►◄
    └─,──USER──=────user_id──┘
```

Here you must code USER= as shown and supply an ID for user_id. You may, of course, enter USER in lowercase, but you must not change it otherwise.

Repetition

An arrow returning to the left means that the item can be repeated.

```
      ┌──────────┐
►►────▼─repeat─┴─────────────────────────────────────────►◄
```

A character within the arrow means you must separate repeated items with that character.

```
      ┌─,────────┐
►►────▼─repeat─┴─────────────────────────────────────────►◄
```

A footnote (1) by the arrow references a limit that tells how many times the item can be repeated.

```
      ┌──────────┐
►►────▼─repeat─┴─(1)─────────────────────────────────────►◄
```

Note:

1 Specify *repeat* up to 5 times.

Defaults

Defaults are above the line. The system uses the default unless you override it. You can override the default by coding an option from the stack below the line. For example:

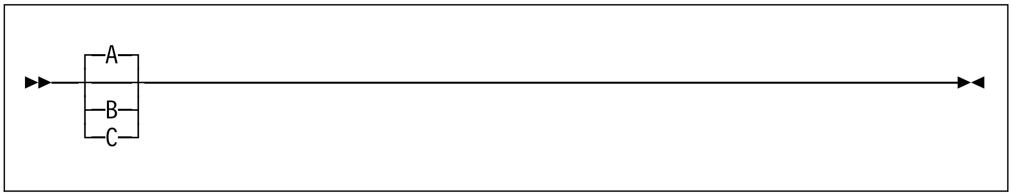

In this example, `A` is the default. You can override `A` by choosing `B` or `C`.

Required Choices

When two or more items are in a stack and one of them is on the line, you **must** specify one item. For example:

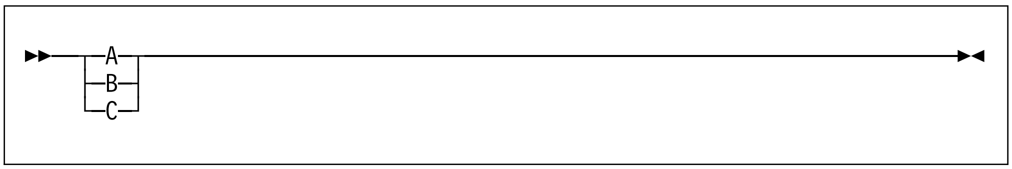

Here you must enter either `A` or `B` or `C`.

Optional Choice

When an item is below the line, the item is optional. Only one item **may** be chosen. For example:

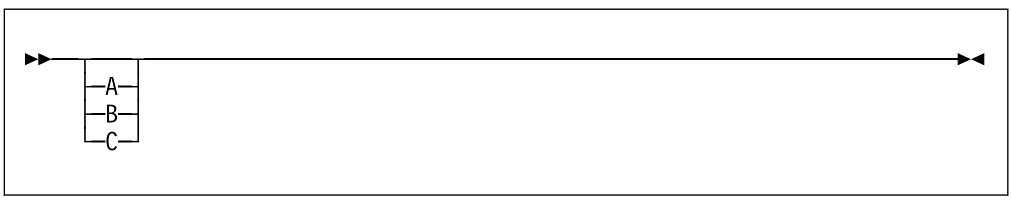

Here you may enter either `A` or `B` or `C`, or you may omit the field.

Summary of Changes

What's New in OS/390 Version 2 Release 7

Summary of Changes
for SC28-1949-04
OS/390 Version 2 Release 7

This book contains the information previously presented in *RMF User's Guide*, SC28-1949-03, which supports the OS/390 Resource Measurement Facility.

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

The following information describes the enhancements that are being distributed with OS/390 Version 2 Release 7.

Online Reporting for Cache Subsystems

New Monitor III reports provide high-level health indicators for the cache subsystems as well as detailed information about the I/O activities.

CACHSUM	Cache Summary report
CACHDET	Cache Detail report

This extends the Monitor III capabilities of monitoring the entire I/O subsystem in the sysplex in an optimal way.

A new option CACHE/NOCACHE is available for Monitor III data gathering to support this function.

Enhanced Enclave Support

The enclave support in Monitor III has been enhanced to assist you in managing business units of work in their sysplex. Many new types of applications (for example, DDF, SAP R/3, DSOM) create enclave transactions executing in several address spaces but they need to be managed as own single business units of work. Therefore, a new report showing resource consumption and delays by enclave will improve significantly performance management for these new applications:

ENCLAVE	Enclave Activity report

Support for OS/390 UNIX System Services

Two new reports (Monitor II and Postprocessor) provide performance information on the Hierarchical File System (HFS).

HFS	Monitor II HFS Statistics report
HFS	Postprocessor HFS Statistics report

This will enable customers to tune their system with regards to the HFS resources. Based on the storage information and utilization statistics in the new reports, customers can size the HFS buffers optimally.

Data gathering for HFS file statistics in the Postprocessor report will be performed in the Monitor III gatherer session. Therefore, a new option HFSNAME is available for Monitor III data gathering to define the appropriate files.

Combined Postprocessor Reports

By supporting new ddnames, the Postprocessor offers the capability to combine the interval or exceptions reports for all intervals belonging to a measurement session into a single data set or output file:

PPRPTS Combined interval report
PPXRPTS Combined exception report

Documentation

Due to the fact that the Spreadsheet Converter (RMF2SC) is part of the Spreadsheet Reporter (RMFPP), you find most information about it in Chapter 18, “RMF Spreadsheet Reporter (RMFPP)” on page 18-1. Only specific functions (conversion of some Monitor II and Monitor III reports) which are outside the scope of RMFPP can be performed with RMF2SC only. If you are interested in these functions, you should refer to the *RMF Programmer's Guide*. The complete description of RMF2SC has been moved to this publication.

History of Changes

What's New in OS/390 Version 2 Release 6

Summary of Changes
for SC28-1949-03
OS/390 Version 2 Release 6

This book contains the information previously presented in *RMF User's Guide*, SC28-1949-02, which supports the OS/390 Resource Measurement Facility.

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

The following information describes the enhancements that are being distributed with OS/390 Version 2 Release 6.

Year 2000 Support for RMF

RMF is an element of OS/390. Beginning with OS/390 Version 1 Release 2, OS/390 is certified as a Year 2000–ready operating system by the Information Technology Association of America (ITAA). Follow-on releases are also Year 2000 ready.

Previous products, such as OS/390 Version 1 Release 1, and all releases of MVS/ESA SP V5, are Year 2000 ready with maintenance applied. Previous products that are not Year 2000 ready will not be marketed after December 31, 1997.

For RMF, the following changes were made to ensure Year 2000 readiness.

RMF expands the year to four digits where necessary and provides a sliding window when specifying a time frame for reporting. This is relevant for the Postprocessor DATE option and the DATE parameter in the BREF and FREF commands of Monitor III.

To select the time frame of interest, the Postprocessor provides a DATE and 'xTOD' option. The DATE option specifies the begin and end date of the reporting time frame, and the syntax allows you to specify a two- or four-digit year.

Sliding Window

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.

- Postprocessor - DATE(yyddd,yyddd)
- Monitor III - BREF DATE=mm/dd/yy

In both cases, a two-digit value will be set properly that it falls into the sliding window.

Online Reporting for the Coupling Facility

Three new Monitor III reports enable you to analyze bottlenecks or problems in the coupling facility area that might result in a performance degradation in the parallel sysplex. To avoid critical situations for production data bases and transaction processing systems, you can see the immediate state of the coupling facility and the structures, and you have all relevant data at a glance in a granularity you can choose. In addition, you can see the results of tuning actions in this area, for example after having removed a coupling structure from an overloaded coupling facility to another one.

Duration Report for the Coupling Facility

In this release, the Postprocessor provides a duration report for the coupling facility.

Sysplex Overview and Spreadsheet Reporting

Overview data is available for the three Postprocessor sysplex reports based on the same technology as used for overview and exception data for single-system reports:

- Coupling Facility Activity report
- Workload Activity report (goal mode)
- Shared Device Activity report

This is a major step towards an integrated Spreadsheet Reporter solution that is not restricted to spreadsheet data from Postprocessor single-system reports.

You can choose from a variety of overview criteria to tailor the appearance of an overview report or the contents of overview records. Especially when using sysplex overview criteria, the Postprocessor provides a sysplex and a single-system view to shared resources, for example, service classes, devices or coupling facility structures.

The new Collector function of the Spreadsheet Reporter allows you to create and to submit Postprocessor jobs directly from the workstation without a logon to the host system.

Repackaging of RMF Load Libraries

RMF creates two new target libraries SERBLINK and SERBLPA. This requires some additional steps in the customization process.

PM of OS/390 Communication with the Distributed Data Server

The Distributed Data Server provides the ability to serve multiple clients in a single-server address space. This capability will by used by PM of OS/390.

Documentation

As part of the name change of OS/390 OpenEdition to OS/390 UNIX System Services, occurences of OpenEdition have been changed to OS/390 UNIX System Services, or its abbreviated name OS/390 UNIX.

Most references to a 4381 processor have been eliminated. If you are still using this processor type, please refer to a previous edition of this publication, no changes have been implemented in this area.

What's New in OS/390 Version 2 Release 4

Summary of Changes
for SC28-1949-02
OS/390 Version 2 Release 4

This book contains the information previously presented in *RMF User's Guide*, SC28-1949-01, which supports the OS/390 Resource Measurement Facility.

This book includes terminology, maintenance, and editorial changes. Technical changes or additions to the text are indicated by a vertical line to the left of the change, here is a brief summary of the changes and additions made to this book.

Data Set Level Reporting

As an extension of the current spectrum of resource oriented reports, RMF offers reporting of data set usage. This new capability is one of the key requirements of many RMF users to help them identify device problems on a data set level. Three new Monitor III reports provide information on how individual data sets on a specific device are being utilized:

- DSND - Data Set Delays report
- DSNJ - Data Set Delays - Job report
- DSNV - Data Set Delays - Volume report

By using this data, the customer can easily identify data sets that should be moved to another device to avoid contentions or bottlenecks.

IMS* Long Lock Detection

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 new Monitor II ILOCK report enables you to identify locking situations that are caused by serialization effects when sharing data among several IMS instances in a sysplex.

Spreadsheet Reporter

The Spreadsheet Reporter (RMFPP) is an enhancement of the Spreadsheet Converter (RMF2SC) to provide a more usable support for converting Postprocessor listings and Overview records into spreadsheets. In addition, it provides sample macros to help you in presenting and analyzing performance data at a glance.

You find many report samples in the *RMF Performance Management Guide*.

Postprocessor Cache Reporting Enhancements

The RMF Postprocessor Cache support has been enhanced:

- The Postprocessor Cache Subsystem report contains a Summary report providing a comprehensive view about all control unit and DASD data belonging to the cache subsystem.
- The Exception and Overview reports enable customer-defined reporting and provide data for Spreadsheet Converter processing.

Storage Utilization Fields

RMF collects additional storage data to allow planning for storage consumption. They are available in SMF records type 71 and in the Postprocessor Exception and Overview reports.

PM of OS/390 Enhancements

Analysis Support

Performance Monitoring of OS/390 has been enhanced by a set of analysis functions providing 'drop in' performance analysis simplifying the task of performance analysts and enhancing their effectiveness through:

- 'Point and shoot' navigation
- PerfDesks for performance analysis actions

TCP/IP Support

TCP/IP support is available as an alternative host connection to APPC.

Documentation

In previous editions of this publication, there was sometimes a differentiation between 4381 processors and 3090 or ES/9000 processors. As you know, new CMOS processors are available as

- IBM 9672
- IBM S/390 Parallel Enterprise Server* - Generation 3 and Generation 4

- IBM S/390 Multiprise* 2000

Therefore, the reference to processors has been changed in that way that they refer to all of these processor types, and only differences are shown explicitly that belong to the 4381 processor.

Part 1. Introduction

This chapter gives you an overview about all the capabilities of RMF.

- Data Gathering with Monitor I, Monitor II, and Monitor III
- Reporting with Monitor III, Monitor II, and the Postprocessor
- Monitoring on the workstation with PM of OS390
- Creating spreadsheets with the Spreadsheet Reporter

And you will be informed about the Sysplex Data Server for accessing data across the sysplex.

Chapter 1. RMF - Your Performance Management Tool

This is OS/390 RMF

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
- Performance Monitoring of OS/390
- Client/Server Enabling
- Spreadsheet Converter / Reporter
- Sysplex Data Server

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.

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.

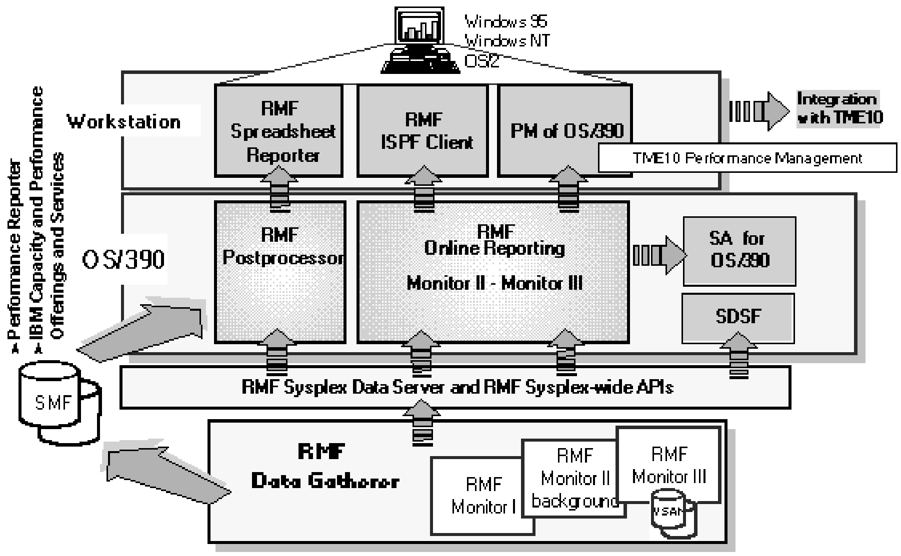

You can run data gathering on each OS/390 system and use the RMF Sysplex Data Server to have all data available on the one system on which you run your performance management tasks.

Figure 1-1. 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 three modes for reporting on the performance of your system:

- Monitor II TSO/E 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 3270 local session (non-SNA): You define a BTAM connected terminal for online and interactive reporting. This makes your performance reporting run independently of TSO/E.
- 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 a 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 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 very fast way to access performance data without having to go through dumping, sorting, and merging of all kinds of SMF records.

The Postprocessor offers different types of report:

Interval reports: they draw a picture of 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 — practically no time limitation.

Summary, plot, exception and overview reports: these are other capabilities of the Postprocessor that 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 optimal base for further spreadsheet processing on the workstation.

Viewing Reports on Spreadsheets

The **Spreadsheet Converter (RMF2SC)** makes many Monitor I and Monitor II filed reports, and Monitor II and Monitor III display reports, available on spreadsheets. You can use the spreadsheet product of your choice to manipulate the data.

The **Spreadsheet Reporter (RMFPP)** is a function that exploits the capabilities of the Spreadsheet Converter, it simplifies the task of converting data to spreadsheet format and offers sample macros to perform data analysis and to create graphical reports.

Monitoring on the Workstation

Performance Monitoring of OS/390 (PM of OS/390) gives you the capability to construct monitoring scenarios and use them whenever necessary. This is done on an OS/2* workstation, and the access to the current performance data of your OS/390 systems is made via the standard PM Common Functions interface, without the need to have a TSO/E session running.

PM of OS/390 is the first step of the integration with TME10 performance management.

Client/Server Enabling (RMFCS) uses the client/server concept to support performance management for OS/390 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 OS/390 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.

Accessing Data across the Sysplex

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.

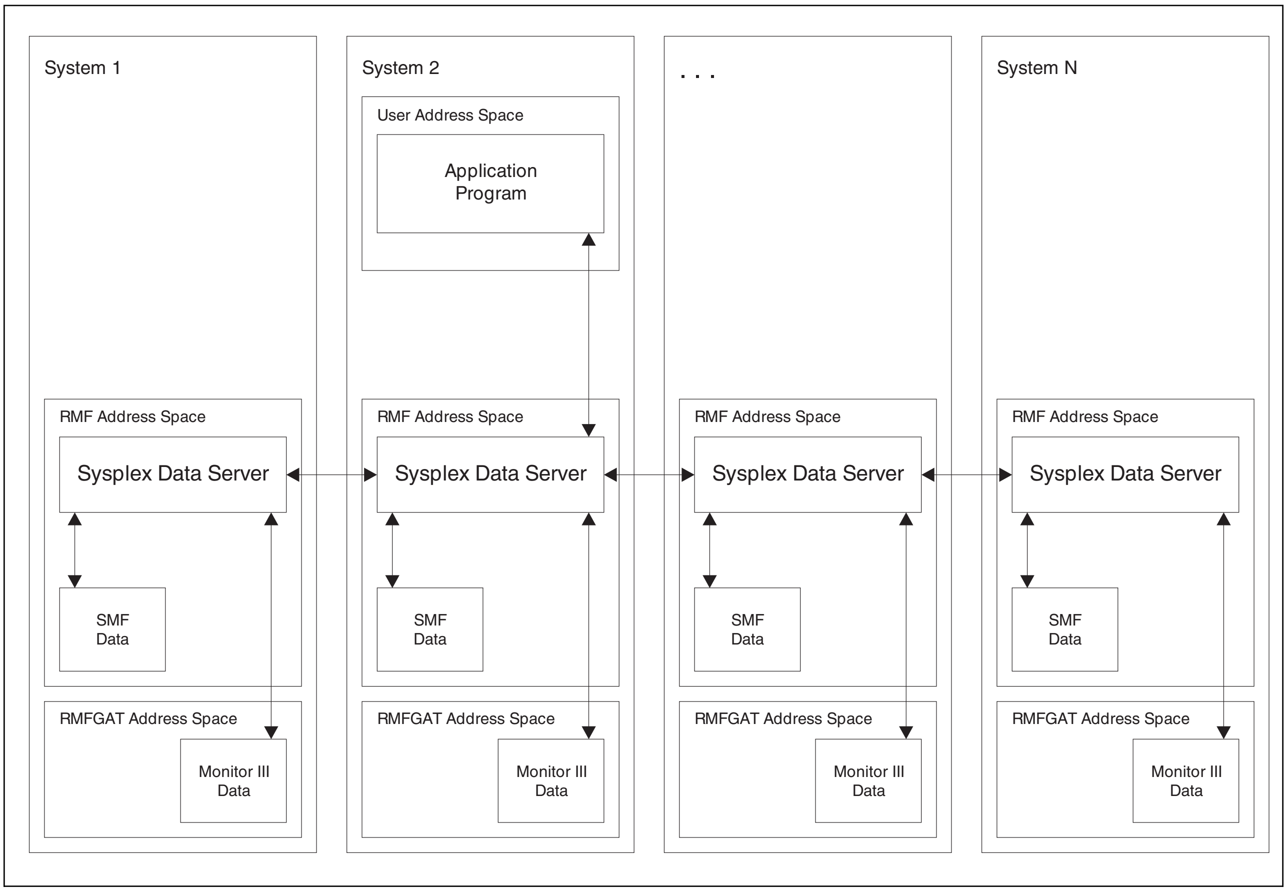

Figure 1-2. RMF Sysplex Data Server Data Flow. The application program may call the RMF Sysplex Data Services from any point in the sysplex.

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 General-Use 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 “Ensure Access to SMF Data for the Sysplex Data Services” on page 2-4.

Sysplex Data Services for SMF Data

RMF or other MVS products store SMF data in a wrap-around buffer. You can choose to create the RMF Buffer for SMF data when you start RMF. The size of the buffer and the types of SMF records stored in it can be specified as a RMF startup parameter. The RMF Sysplex Data Server services return SMF data when the data 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 and the data service for SMF data, which collects only records created by active monitor sessions.

---

[1] Authorization of application programs is provided by the OS/390 Security Server (RACF* element), 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.

1

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

Figure 1-3. Activities that RMF Monitors 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	*			*
	*	79.9	74.1	Device	*	*	*	*
	*	79.10		Domain	*	*	*	*
	*			Enclave	*			
	*	79.7	77	Enqueue	*	*	*	*
		79.15		IRLM long locks		*		
	*	79.13/14	78.1/3	I/O queuing	*	*	*	*
74.3/6	*			OS/390 UNIX	*	*		*
		79.11	75	Page/Swap data set		*	*	*
		79.4	71	Paging		*	*	*
	*	79.3	70	Processor	*	*	*	*
		79.6		Reserve		*		*
72.2/4	*	79.3		Storage	*	*		*
			76	System counters			*	*
		79.8		Transactions		*		*
	*		78.2	Virtual storage	*		*	*
	* *		72.1 72.3	Workload Performance groups Service 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
- Domains and performance groups
- Service classes and workload groups
- Enclaves

For any domain and performance group or service 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

Part 2. Administration

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

The following topics are dealt with in this part:

What Administration Involves

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 10-5 and Part 6, "Reporting Reference" on page 13-19.

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

Chapter 2. Setting Up RMF

What to Do Before You Start

This chapter deals with:

- 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
- The installation of workstation functions

Overview on Administration Tasks

After installing RMF, you have to set up certain functions for RMF:

Administration Tasks

- Activating RMF
- Setting Up RMF Control Session including Monitor I and Monitor II
- Setting Up the Monitor III Gatherer Session RMFGAT
- Setting Up the Distributed Data Server
- Synchronizing SMF Recording Intervals
- Storing Gatherer Options
- Considering Reporting Aspects
- Installing Workstation Functions

Activating RMF

Activating RMF consists of the following steps:

- Customizing the System Environment
- Specifying Access Definitions

Customizing the System Environment

Define RMF Library Authorization

Since OS/390 Release 6, RMF has two new target libraries SERBLINK and SERBLPA:

- In previous releases, all RMF load modules resided in SYS1.LINKLIB and SYS1.LPALIB. They have been moved to new libraries SYS1.SERBLINK and SYS1.SERBLPA (except the two extended Router SVC routines IGX00007 and IGX00022).
- The load modules IGX00007 and IGX00022 will remain in SYS1.LPALIB. These two load modules will be changed to reside "above the line" (RMODE=ANY).

Therefore, the installation jobs have been changed accordingly.

Details can be found in the current version of the *OS/390 2.7.0 Program Directory*.

If you are activating RMF for the first time, you have to define these new libraries as 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 *OS/390 MVS Initialization and Tuning Reference*. For information about the syntax of the SETPROG command, see *OS/390 MVS System Commands*.

Ensure Linkage to Language Environment*

Two components of RMF, the Postprocessor and PM of OS/390, use the services of the Language Environment (formerly known as LE/370). To do so, they need access to the data set SYS1.SCEERUN. There are two ways of providing this access:

- The recommended way is to include the data set SYS1.SCEERUN in the LINKLST of the system on which RMF is running. No further action need then be taken when starting the separate components.
- If, for any reason, you do not wish to include SYS1.SCEERUN in the LINKLST, you must specify it as the STEPLIB of the job step that starts the component.

  Sample JCL members are provided in SYS1.SAMPLIB and SYS1.PROCLIB to help you do this. They are, for the Postprocessor, member ERBSAMPP, and for PM of OS/390, members GPMTPSAM and GPMTCSRV.

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 *OS/390 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.

In the past, there was a recommendation given to define a specific XCF transport class, this recommendation is no more valid. It is recommended you 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 *WSC Flash 9723*.

Check the Program Properties Table (PPT)

OS/390 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 you to specify a PPT entry for ERBMFMFC and ERB3GMFC in a SCHEDxx Parmlib member, which must include the RMF defaults and user overrides. If you migrate directly from MVS/ESA to OS/390, make sure that the PPT specification for OS/390 RMF match your specifications for RMF on the previous release of MVS/ESA.

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.

Specifying Access Definitions

Ensure Access to SMF Data for the Sysplex Data Services

Users of applications that call sysplex data services to access SMF data must have RACF authorization to do so. One such application is the RMF Postprocessor, another one is the data gatherer of the Monitor II ILOCK command.

Use the following RACF commands:

1. To activate the resource class:

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

2. To define the resource name:

   ```
   RDEFINE FACILITY resname UACC(NONE)
   ```

   where `resname` is the resource name, for example, ERBSDS.SMFDATA, or a generic resource name, for example, ERBSDS.*.

3. To grant the `userid` of the application program READ access:

   ```
   PERMIT resname CLASS(FACILITY) ID(userid) ACC(READ)
   ```

4. Activate changes:

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

Ensure Access to OS/390 UNIX System Services or Resources

If a cataloged procedure starts a program that uses UNIX System services or resources, the procedure should be defined to the security program. Therefore, Monitor III must be defined so that it can obtain the correct data. The cataloged procedure RMFGAT must be defined or defaulted to have a user ID of RMFGAT.

The same has to be performed for the cataloged procedure GMPSERVE to start the Distributed Data Server.

Enter the following RACF commands to give RMFGAT and GPMSERVE an OMVS user ID (UID) and to designate the root directory as its home directory:

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

For details, please refer to *OS/390 UNIX System Services Planning*.

Setting Up 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
- Storing Gatherer Defaults
- Preallocating Monitor I and Monitor II Reporter Data Sets

Customizing the RMF Control Session

IBM provides the cataloged 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 control session RMF 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 2-8).

This example shows the RMF procedure as supplied:

RMF Control Session

```
//IEFPROC   EXEC  PGM=ERBMFMFC,REGION=32M,
//          PARM=''
```

PARM
Can be used to specify gatherer options, and the SMF Data Buffer options that the RMF Sysplex Data Server will use. The format of this option is described in “Controlling the SMF Buffer” on page 3-4. The defaults mean that specifying **SMFBUF** is equivalent to:

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.

RMF reads its ERBRMFxx members from the Parmlib concatenation as defined in the MVS LOADnn 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:

Example

```
//IEFPROC   EXEC  PGM=ERBMFMFC,REGION=32M,
//          PARM=''
//IEFPARM   DD    DSN=parmlibname,DISP=SHR
```

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

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

SMF Data Buffer Example

```
//IEFPROC   EXEC  PGM=ERBMFMFC,REGION=32M,
//          PARM='SMFBUF'
```

Specifying Priority for RMF

We recommend that the started tasks RMF and RMFGAT have the second- highest priority in the system, next to the system address spaces. How you ensure this, depends on whether the system is running in compatibility mode or in goal mode. If the priority is too low, it can happen that RMF is not dispatched when its interval time expires, and data collection for jobs running with higher priority might be incomplete. This could result in incorrect measurement reports.

- **Compatibility mode**

  Use the system parameters in IEAICSxx and IEAIPSxx to put RMF and RMFGAT in a performance group of their own, and give this performance group a high dispatching priority.

- **Goal mode**

  Use the WLM application to put RMF and RMFGAT in service class SYSSTC, then, its dispatching priority will always be above any installation-defined service class.

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 2-13.

Preallocating Monitor I and Monitor II Reporter Data Sets

RMF dynamically allocates all Monitor I and Monitor II message and report data sets to SYSOUT. 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.

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

Figure 2-1 (Page 1 of 2). 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. You cannot change it in the RMF options.
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). 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	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.

Figure 2-1 (Page 2 of 2). ddnames for Monitor I and Monitor II Data Sets

ddname	Session	Contains	Allocations	Notes
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.
RMFDMxxx	Monitor II local 3270	Hardcopy output	One data set allocated for each session.	**xxx** is the session identifier for Monitor II Local 3270 sessions.
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
- Ensuring Common Storage Tracking

IBM provides the cataloged procedure you need to start the Monitor III gatherer session, it is stored in SYS1.PROCLIB(RMFGAT):

Monitor III Gatherer Session

```
//IEFPROC  EXEC PGM=ERB3GMFC,REGION=40M,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.

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 3-3). 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 to a large storage buffer, or optionally, to user-defined VSAM data sets. Without the VSAM data sets, the data will be over-written 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. At a minimum, 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. Based on practical experience, we recommend you to define six VSAM data sets, each with 50 cylinders of 3380 disk space or its equivalent. On small and medium systems, this will allow for about two days of data. 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 15-14) 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 2-14 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 Syntax

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

Where:

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

volume
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
The primary extent of the VSAM data set (the default is 150 tracks)

Example

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

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

"Controlling Data Set Recording" on page 13-12 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

The preparation of a Distributed Data Server (DDS) host session as server address space for users of PM of OS/390 requires the customization of a GPMSRVxx Parmlib member that is needed for the GPMSERVE procedure to start the Distributed Data Server.

IBM provides a default Parmlib member GPMSRV00. You may tailor this according to your needs.

GPMSRV00

```
/*********************************************************************/
/*                                                                   */
/* NAME:        GPMSRV00                                             */
/*                                                                   */
/* DESCRIPTION: PARMLIB MEMBER FOR THE RMF DISTRIBUTED DATA SERVER   */
/*              HOST ADDRESS SPACE (GPMSERVE)                        */
/*                                                                   */
/*********************************************************************/
TIMEOUT(0)              /* NO SESSION TIMEOUT                        */
MAXSESSIONS_UNIX(5)     /* MAX LOCAL CLIENTS, E.G. PM_OF_OS390 APPC */
MAXSESSIONS_INET(5)     /* MAX REMOTE TCP/IP CLIENTS                 */
SESSION_PORT(5001)      /* TCP/IP PORT FOR INCOMING SESSIONS         */
```

TIMEOUT(seconds)

The number of seconds after that GPMSERVE will timeout inactive sessions. A value of 0 indicates no timeout.

Default: TIMEOUT(0)

MAXSESSIONS_UNIX(numsess)
The maximum number of local sessions (via UNIX domain sockets) that the server is able to handle concurrently. If the maximum number of sessions is reached, the server will reject further session requests.

MAXSESSIONS_UNIX is an integer between 0 and 100.

The parameter includes the number of parallel PM of OS390 sessions initiated via APPC, because the APPC transaction program on the OS/390 host communicates with the DDS address space via UNIX domain sockets.

Default: MAXSESSIONS_UNIX(5)

MAXSESSIONS_INET(numsess)
The maximum number of TCP/IP sessions that the server is able to handle concurrently. If the maximum number of sessions is reached, the server will reject requests to establish a new session.

MAXSESSIONS_INET is an integer between 0 and 100.

The parameter includes the number of parallel PM of OS390 sessions initiated via TCP/IP.

Default: MAXSESSIONS_INET(5)

SESSION_PORT(port)
The TCP/IP port number on which the OS/390 server will listen for incoming connection requests for a session-oriented communication.

SESSION_PORT is an integer between 0 and 65535.

For TCP/IP sessions coming from PM of OS390, you must choose a port between 5000 and 9999.

Default: SESSION_PORT(5001)

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

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 (in SMFPRMxx) 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.

Monitor I / II Option RECORD (in ERBRMFxx) specifies SMF record collection.

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

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

See *OS/390 MVS Initialization and Tuning Reference* for details.

Defining SMF Synchronization

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

SMF Options

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:

Monitor I Options

SYNC(SMF) Synchronization with SMF - this is the default and means that records will be written as specified via 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 III Synchronization

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)

Notes:

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 for every 30 minutes (taking three intervals from SYSA and two intervals from SYSB).

Use the following values to synchronize SMF record writing:

Recommendation

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

For information about SMF record format and printing, see the *RMF Programmer's Guide*. See the *OS/390 MVS System Management Facilities (SMF)* book for descriptions and formulas of the fields for each SMF record RMF produces.

Storing Gatherer Options

The following tasks help you in tailoring the options for data gathering according to your requirements:

- Naming Parmlib Option Members
- Generalizing Parmlib Members
- Defining Parameters for Monitor I
- Defining Parameters for Monitor II
- Defining Parameters for Monitor III

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

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

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

Naming Parmlib Option Members

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

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

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

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

- ERBRMF02 for Monitor I
- ERBRMF03 for Monitor II

You can use the default and alternative members as they are, or you can alter them to suit your needs. You can also create new Parmlib members from scratch, following the naming convention of `ERBRMFxx`. For the options and their syntax, see Chapter 11, Long-term Data Gathering with Monitor I, Details of Session and Report Commands and Chapter 13, Short-term Data Gathering with Monitor III, 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.

Parmlib Concatenation

With the support of Parmlib concatenation in OS/390, it is recommended to define one or more *customer* Parmlibs that can be specified in the LOADnn Parmlib member. Then you can distinguish between system-supplied members (for example through the SMP/E installation process) which will be stored by default in SYS1.PARMLIB, and customer-modified members in an additional Parmlib data set.

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

Generalizing Parmlib Members

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

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

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

RMF supports the use of symbolic names in:

- All RMF Parmlib members

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

Example

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

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

Example

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

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

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

Example

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 OS/390.

- 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 OS/390 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, please refer to *OS/390 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

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

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                        */
  DEVICE(DASD)             /* DIRECT ACCESS DEVICES MEASURED        */
  DEVICE(NOTAPE)           /* NO TAPE DEVICES MEASURED              */
  DEVICE(NOCHRDR)          /* NO CHARACTER READER DEVICES MEASURED */
  DEVICE(NOUNITR)          /* NO UNIT RECORD DEVICES MEASURED       */
  DEVICE(NOCOMM)           /* NO COMMUNICATION DEVICES MEASURED     */
  DEVICE(NOGRAPH)          /* NO GRAPHICS DEVICES MEASURED          */
  DEVICE(NONMBR)           /* NO SELECTION BY DEVICE NUMBERS        */
  DEVICE(NOSG)             /* NO SELECTION BY STORAGE GROUPS        */
NOENQ                      /* NO ENQUEUES MEASURED                  */
  IOQ(DASD)                /* DASD I/O QUEUEING MEASURED            */
  IOQ(NOTAPE)              /* NO TAPE I/O QUEUEING MEASURED         */
  IOQ(NOCHRDR)             /* NO CHARACTER READER I/O QUEUEING      */
  IOQ(NOUNITR)             /* NO UNIT RECORD DEVICE I/O QUEUEING    */
  IOQ(NOCOMM)              /* NO COMMUNICATION I/O QUEUEING         */
  IOQ(NOGRAPH)             /* NO GRAPHICS DEVICE I/O QUEUEING       */
  IOQ(NONMBR)              /* NO SELECTIVITY BY LCU NUMBERS         */
  PAGESP                   /* PAGE/SWAP DATASET STATISTICS          */
  PAGING                   /* PAGING DATA                           */
NOTRACE                    /* NO TRACE REPORT                       */
  VSTOR(S)                 /* VIRTUAL STORAGE SUMMARY DATA          */
  WKLD(PERIOD)             /* WORKLOAD MANAGER DATA                 */
/********************************************************************/
/* PART 2: TIMING                                                   */
/********************************************************************/
  CYCLE(1000)              /* SAMPLE EVERY SECOND (1000 MSEC)       */
NOSTOP                     /* ACTIVE UNTIL OPERATOR ISSUES STOP     */
  SYNC(SMF)                /* USE INTVAL/SYNCVAL FROM SMFPRMXX      */
/********************************************************************/
/* PART 3: REPORTING / RECORDING OF DATA                            */
/********************************************************************/
NOOPTIONS                  /* OPTIONS NOT DISPLAYED, NO REPLY       */
  RECORD                   /* WRITE SMF RECORDS EVERY INTERVAL      */
NOREPORT                   /* NO WRITTEN REPORTS TO SYSOUT          */
  SYSOUT(A)                /* REPORTS TO CLASS A, IF REPORT         */
/********************************************************************/
/* PART 4: USER EXITS                                               */
/********************************************************************/
NOEXITS                    /* DO NOT TAKE USER EXITS                */
```

Note: If you miss gathering options for the coupling facility, for UNIX System Services or XCF, keep in mind that this 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.

The options are:

ERBRMF02

```
/*******************************************************************/
/* PART 1: MEASUREMENTS                                            */
/*******************************************************************/
  CACHE                    /* CACHE STATISTICS                     */
  CHAN                     /* CHANNEL STATISTICS                   */
  CPU                      /* CPU STATISTICS                       */
  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                    */
  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/SWAP DATASET STATISTICS         */
  PAGING                   /* PAGING DATA                          */
  TRACE(RCVUICA,END)       /* TRACE 'UIC AVERAGE'                  */
  TRACE(RCVCPUA,END)       /* TRACE 'CPU USAGE*16'                 */
  TRACE(RCVPTR,END)        /* TRACE 'PAGING RATE'                  */
  TRACE(DMDTCMPL,END)      /* TRACE 'CURRENT MPL'                  */
  TRACE(DMDTRUA,END)       /* TRACE 'AVG NUMBER READY USERS'       */
  VSTOR(D)                 /* VIRTUAL STORAGE DETAIL DATA          */
  WKLD(PERIOD)             /* WORKLOAD MANAGER DATA BY PERIOD      */
  WKLD(SYSTEM)             /* WORKLOAD MANAGER DATA FOR SYSTEM     */
  WKLD(RANGE)              /* WORKLOAD MANAGER DATA BY RANGES      */
  WKLD(DOMAIN)             /* WORKLOAD MANAGER DATA BY DOMAINS     */
  WKLD(TIME)               /* WORKLOAD MANAGER DATA BY TIME SLICE  */
/*******************************************************************/
/* PART 2: TIMING                                                  */
/*******************************************************************/
  CYCLE(250)               /* SAMPLE EVERY 250 MILLISECONDS        */
  STOP(8H)                 /* STOP AFTER 8 HOURS                   */
  SYNC(SMF)                /* USE INTVAL/SYNCVAL FROM SMFPRMXX     */
/*******************************************************************/
/* PART 3: REPORTING / RECORDING OF DATA                           */
/*******************************************************************/
  OPTIONS                  /* OPERATOR MAY EXAMINE/CHANGE OPTIONS  */
  RECORD                   /* WRITE SMF RECORDS EVERY INTERVAL     */
  REPORT(REALTIME)         /* WRITE REPORTS EACH INTERVAL          */
  SYSOUT(A)                /* REPORTS TO CLASS A, IF REPORT        */
/*******************************************************************/
/* PART 4: USER EXITS                                              */
/*******************************************************************/
NOEXITS                      /* DO NOT TAKE USER EXITS             */
```

Note: This option set specifies that real-time reports will be created. If you run your system in goal mode, you will get no Workload Activity report based on the WKLD options, because this report you only get as sysplex report from the Postprocessor. There is no need to change the WKLD options for goal mode, because the suboptions (as PERIOD or DOMAIN) are for reporting purpose only and are not relevant for data gathering.

Defining Parameters for Monitor II

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 ?                       */
NODDMN                     /* DOMAIN DATA ?                        */
NODEV                      /* DEVICE DATA ?                        */
NODEVV                     /* DEVICE DATA FOR SPECIFIC DEVICE ?    */
NOIOQUEUE                  /* I/O QUEUING DATA ?                   */
NOPGSP                     /* PAGE/SWAP DATASET MEASUREMENTS ?     */
NOSENQ                     /* SYSTEM ENQUEUE CONTENTION ?          */
NOSENQR                    /* SYSTEM ENQUEUE RESERVE DATA ?        */
NOSPAG                     /* SYSTEM PAGING ACTIVITY ?             */
NOSRCS                     /* SYSTEM REAL STORAGE/CPU/SRM DATA ?   */
NOTRX                      /* TRANSACTION ACTIVITY 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 options it contains 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 ?                       */
  DDMN                       /* DOMAIN DATA ?                        */
  DEV                        /* DEVICE DATA ?                        */
NODEVV                       /* DEVICE DATA FOR SPECIFIC DEVICE ?    */
  IOQUEUE                    /* I/O QUEUING DATA ?                   */
  PGSP                       /* PAGE/SWAP DATASET MEASUREMENTS ?     */
  SENQ                       /* SYSTEM ENQUEUE CONTENTION ?          */
  SENQR                      /* SYSTEM ENQUEUE RESERVE DATA ?        */
  SPAG                       /* SYSTEM PAGING ACTIVITY ?             */
  SRCS                       /* SYSTEM REAL STORAGE/CPU/SRM DATA ?   */
  TRX                        /* TRANSACTION ACTIVITY 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

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)     */
  IOSUB                      /* I/O SUBSYSTEM GATHERING ACTIVE       */
NOCFDETAIL                   /* NO COUPLING FACILITY DETAILS         */
  CACHE                      /* ACTIVATE CACHE GATHERING             */
```

Considering Reporting Aspects

- Setting Up the RMF CLISTs
- Grouping Monitor II and Monitor III Users

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.

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.

Grouping Monitor II and Monitor III Users

Since 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 performance group 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 III users are put in a separate performance group or service class.

Installing Workstation Functions

- Installing the RMF Spreadsheet Reporter - RMFPP
- Installing Performance Monitoring of OS/390 - PM of OS/390
- Installing RMF Client/Server Enabling - RMFCS

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 will be automatically installed in the RMF distribution libraries on your host system through SMP/E. Then, it is the responsibility of each user to perform the installation on the workstation according to the provided description.

Installing the RMF Spreadsheet Reporter - RMFPP

Download the RMF Spreadsheet Reporter to your workstation.

For instructions on downloading, as well as activation and usage of RMFPP functions, please refer to chapter Chapter 18, “RMF Spreadsheet Reporter (RMFPP)” on page 18-1.

Installing Performance Monitoring of OS/390 - PM of OS/390

The installation of PM of OS/390 comprises the following steps. Depending on your preference for an APPC connection or a TCP/IP connection, one of the first two steps is optional, but both steps are recommended.

1. Preparation of an APPC connection
   a. Host - Definition of an APPC TP profile
   b. Workstation - Definition of the OS/2 APPC setup
2. Preparation of an TCP/IP connection
   a. Host - Setup of the server start procedure
   b. Workstation - Server bring-up by the client
3. Installation of Performance Monitoring Common Functions

   The Performance Monitoring Common Functions (PM Common Functions) is a prerequisite for PM of OS/390. Install this product before installing PM of OS/390. If the Performance Monitoring Common Functions is already installed on your workstation, you can omit this step.
4. Installation of PM of OS/390
5. Tailoring PM of OS/390 to your needs

These steps are described in detail in “Installing PM of OS/390” on page 19-13.

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 20-6.

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 10-5 and Part 6, "Reporting Reference" on page 13-19.

Chapter 3. Starting and Stopping RMF

The RMF Control Session

This chapter 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

- How to start the Distributed Data Server

Starting RMF

Enablement

RMF is an optional element of OS/390. 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 form of the START command for RMF is:

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

parm

The SMF buffer options (see “Controlling the SMF Buffer” on page 3-4) or the session options for Monitor I (see Chapter 11, “Long-term Data Gathering with Monitor I” on page 11-1).

If you specify options, each must have the form:

`option [(value)]`

Multiple options must be separated by commas and enclosed in parentheses.

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 field. See Chapter 5, How RMF Processes Session Options.

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

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

  ```
  START RMF,,,(WKLD(SYSTEM),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 follows:

RMF in a Sysplex

```
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 2-8).

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 3-2.

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 4-2) 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, and any local 3270 display session. 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 4-7. The form 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, but they can also be written to an in-storage, wrap-around SMF buffer (see “Accessing Data across the Sysplex” on page 1-6). You can control the size of this buffer and the SMF record types that RMF writes to it, using the SMFBUF option.

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

- A size of 32 megabytes
- Collection of SMF record types 70 to 78, with all subtypes of each

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 RMF Control Session including Monitor I and Monitor II” on page 2-5). 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(70:78,79(1:15))
```

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

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=32M,
//       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))
```

Starting the Distributed Data Server

The Distributed Data Server provides the ability to serve multiple clients in a single-server address space. This capability will by used by PM of OS/390.

Example

To start the Distributed Data Server, assuming that you have stored the corresponding parameters in Parmlib member GPMSRV01, issue the command:

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

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

You find details about the Distributed Data Server in “Host - Definition of an APPC TP Profile” on page 19-15 and “Host - Setup of the Server Start Procedure” on page 19-28.

Chapter 4. Starting and Controlling Monitor Sessions

Gatherer and Display Sessions

This chapter 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 II local 3270 display sessions
- Monitor III gatherer session
- RMF Client/Server sessions (RMFCS)

You will find details of the options and commands for all RMF sessions in Part 5, Data Gathering Reference and Part 6, Reporting Reference, where each RMF session has a chapter to itself.

Specifying Session Options

When starting or modifying the sessions described in this chapter, 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 RMF Control Session including Monitor I and Monitor II" on page 2-5 for more details).
- The RMF Parmlib member, or other equivalent data set member containing session options. See "Storing Gatherer Options" on page 2-13 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 5-1.

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 10-5 and Part 6, "Reporting Reference" on page 13-19.

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 2-8. 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 4-4.

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.
- **A 3-digit device number** for a Monitor II local 3270 display session. Use the device number of the terminal to be used
- **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 5-1 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 11, “Long-term Data Gathering with Monitor I” on page 11-1.

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 Session and Report Commands” on page 16-23.

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 II Local 3270 Display Session

The value for `session-id` for a local 3270 display session is **nnn**, the three-digit hexadecimal device number of the display station to be used for the session. Because of BTAM restrictions, four-digit numbers are not possible.

You cannot specify options in the START command. You control a Monitor II local 3270 display session with display commands, which are described in Chapter 16, “Snapshot Reporting with Monitor II” on page 16-1.

Example

To start a Monitor II local 3270 display session for a terminal with a device number of 642, issue the command:

`MODIFY RMF,START 642`

When you issue a START command, RMF readies the Monitor II display functions your session requires, and issues a message to the operator console when RMF is ready to receive your display commands. RMF displays the menu on the screen.

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 2-8). 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 13, “Short-term Data Gathering with Monitor III” on page 13-1.

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)`

Starting an RMF Client/Server Session (RMFCS)

RMF Client/Server Enabling (RMFCS) is a concept that supports performance management for OS/390 systems without an active TSO/TCAS subsystem on the OS/390 host. RMFCS allows you to establish as many sessions as you want with any OS/390 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 OS/2 windows by using the ISPF/SPLIT function, which allows 32 logical screens. Each SPLIT creates a new OS/2 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 OS/390 performance management with a state-of-the-art user front end.

Hitherto, one or more 3270 TSO sessions were used for online monitoring of OS/390 performance data. The new concept of RMFCS uses an OS/2-based workstation as the single point of control for multiple OS/390 systems.

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 COMMNDxx to have the task started automatically 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 20, “RMF Client/Server Enabling (RMFCS)” on page 20-1.

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 5-1.

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)
  ```

As described in “CFDETAIL” on page 13-4, the default option for data gathering for the coupling facility should be `NOCFDETAIL` to avoid significant system overhead. If there is a need to gather detailed data because of a problem with the coupling facility, you can activate this with the following command:

Activate CFDETAIL Option

```
ROUTE *ALL,MODIFY RMF,MODIFY III,CFDETAIL
```

It is very important to route the MODIFY command to all members in the sysplex. Only one member is the master gatherer (as seleceted by the Sysplex Data Server), and this is transparent and cannot be modified externally by a command or option. This concept ensures that detailed data gathering will be performed only on one member in the sysplex.

Displaying RMF Status

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

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

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

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

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

Examples

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

  ```
  MODIFY RMF,DISPLAY ACTIVE
  ```

  or

  ```
  F RMF,D
  ```

  to use the shortest form.

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

  ```
  F RMF,D ZZ
  ```

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

  ```
  F RMF,D ALL
  ```

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

  ```
  F RMF,D III
  ```

- To display the SMFBUF option, issue the command:

  ```
  F RMF,SMFBUF
  ```

Stopping RMF Sessions

You can end sessions in three ways:

- By issuing the system command STOP, which stops all active background sessions. See “Stopping RMF” on page 3-4.
- By specifying a time value in the STOP option for a specific session. See Part 5, “Data Gathering Reference” on page 10-5.
- By issuing a STOP session command to stop a specific session. All other active sessions continue processing. See “Stopping a Specific Session” on page 4-8.

  Note: You can stop a Monitor II local 3270 display session with a display command or a STOP session command.

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`

- To stop a Monitor II local 3270 session on the terminal with an ID of 642, issue the session command:

  `MODIFY RMF,STOP 642`

  Alternatively, you can issue the following display command from the terminal:

  `END`

  Note: You can also use the QUIT, QQ, X, Z, or STOP command to stop the Monitor II session.

Chapter 5. How RMF Processes Session Options

Different Ways to Specify 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 chapter 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 **later** in the list. For example, if you specify MEMBER(02,07), RMF first notes the options from ERBRMF02, then processes those from ERBRMF07, overwriting any conflicting options in the process. This means that, if ERBRMF02 specifies ENQ(DETAIL) and ERBRMF07 specifies ENQ(SUMMARY), RMF establishes ENQ(SUMMARY) for the session.

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

4. The RMF default Parmlib members.

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

5. Program defaults (lowest priority).

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

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

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

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

Displaying a List of Options in Effect for a Session

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

- Issuing the DISPLAY session command from the operator console. For information about issuing a DISPLAY session command, see “Displaying RMF Status” on page 4-6.
- 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,,,(WKLD(GROUP),MEMBER(10),CYCLE(100),DEVICE(COMM))
```

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

```
WKLD(GROUP)
CYCLE(100)
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:

```
WKLD(GROUP)
CYCLE(100)
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=32M,
//             PARM='CYCLE(150),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:

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

RMF ignores CYCLE(150) 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:

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

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

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

RMF ignores WKLD(SYSTEM) and 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(GROUP) -- COMMAND       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(100) -- COMMAND        MEMBER(02) -- EXEC
RECORD -- DEFAULT            MEMBER(10) -- COMMAND
SYSOUT(A) -- DEFAULT         IOQ(NOCHRDR) -- DEFAULT
NOREPORT -- DEFAULT          IOQ(NOCOMM) -- DEFAULT
NOPAGESP -- MEMBER           IOQ(DASD) -- DEFAULT
DEVICE(NOCHRDR) -- 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 Session and Report Commands” on page 16-23. 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(100), MINTIME(200) and STOP(40M).

If you issue the command:

```
MODIFY  RMF,MODIFY III,CYCLE(100),MINTIME(200),STOP(40M)
```

the options will be modified as you want.

If, however, member ERBRMF10 includes the options:

```
CYCLE(100)
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(100), 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 MVS 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 *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

If your particular performance-management situation is not dealt with here, refer to *RMF Performance Management Guide*, which goes into more detail.

Chapter 6. Performance Administration

Prerequisites

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

- What you show know about MVS workload management
- How to migrate from compatibility to goal mode
- What the hierarchical view of performance looks like

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 1-11, 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 to. However, meeting the business needs of the users is the reason the system exists.

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 also important with regard to the enhanced capabilities of performance management in an OS/390 system running in goal mode. 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.

If your system is still running in compatibility mode, some data in the RMF reports (for example, execution velocity) help you in planning the migration to goal mode.

Defining performance goals for your installation is an investment for the future.

MVS Workload Management Solution for Today and Tomorrow

When running in compatibility mode, MVS requires you to translate your data processing goals from high-level objectives about what work needed to be done into the extremely technical terms that the system could understand. This translation requires highly skilled staff, and can be protracted, error-prone, and eventually in conflict with the original business goals. The multi-system, sysplex, parallel processing, and data-sharing environments further add to the complexity.

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

Workload management requires a shift of focus from tuning at a system resource level to defining performance expectations. This requires a basic shift in philosophy towards goal-oriented systems management.

Fewer, simpler, and consistent system externals

Workload management provides a way to define MVS externals and tune MVS 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. The IPS and ICS Parmlib members, as well as certain parts of the OPT Parmlib member need no longer be used to specify resource management.

Externals reflect customer expectations

Workload management provides new MVS 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 MVS, 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*. Some installations might already have this kind of information in an SLA. 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 MVS-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 or she 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.

  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:

$$\mathit{Velocity} = \frac{\mathit{Using\ samples}}{\mathit{Using\ samples} + \mathit{Delay\ samples}} \times 100$$

Please, refer to “Workflow and Velocity” on page 7-5 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.

Using Modes to Migrate to Workload Management

The process of converting from your present installation to workload management requires some changes. All changes are transparent to users, and to your applications. No JCL changes are required.

For ease of migration, workload management provides

Two modes of operation

- The traditional method of performance management with the ICS and the IPS is called *compatibility mode*.
- The goal-oriented performance management with the service definition is called *goal mode*.

You can define your goal-oriented policies in the service definition while running your installation in compatibility mode with the existing ICS and IPS. Then, when you are comfortable with your service definition, you can activate a goal-oriented policy, and switch into goal mode.

You can toggle between the two modes of management dynamically by an operator command, without a system IPL. When switching back to compatibility mode, the previous IPS and ICS parameters will be restored automatically.

Run Sysplex with Systems in Compatibility Mode

In order to use workload management, you must be running a sysplex. How you accomplish this depends on what the test and production environments are in your installation. You can be running a multi-system sysplex, or a monoplex.

You should run OS/390 in compatibility mode, with your existing IPS and ICS, until you complete the steps for migrating to workload management, and are comfortable switching into goal mode.

Defining Performance Goals

In the process of defining performance goals, RMF supports you by providing response times and velocity data while you are running your systems in compatibility mode.

Postprocessor Workload Activity Report

- The report shows the average response time for your CICS and IMS subsystems if you have used the SRVCLASS parameter in your IEAICSxx Parmlib member.
- A new field called EX VEL% shows the execution velocity for each performance group in compatibility mode.

The SRVCLASS parameter lets you associate a service class with a report performance group in compatibility mode. You can define a service policy, with service classes representing CICS and IMS work. You define classification rules for the service classes, install the service definition, and activate the service policy. All work in that service class is reported under the report performance group in the Workload Activity report.

To get the response time information while in compatibility mode with the SRVCLASS parameter, you should do the following:

- Define a service definition with:
  - One or more service classes representing the CICS or IMS work for which you want the response time information.

    Note: You need to have the levels of CICS and IMS that support workload management installed.

    You can assign any kind of response time goal to these service classes, because the system is NOT processing towards the goals.
  - Classification rules assigning the work to these service classes.
  - At least one policy.

- In your IEAICSxx member, list the service classes that you defined under the subsystem (CICS or IMS) and associate them with a report performance group.

  For example, if you defined 3 service classes for your IMS work in your service definition, you would define the following in your IEAICSxx member:

  Example

  ```
  SUBSYS=IMS
         SRVCLASS=IMSHI,RPGN=50
         SRVCLASS=IMSMED,RPGN=60
         SRVCLASS=IMSLOW,RPGN=70
  ```

- Activate the service policy
- Issue a SET ICS= to process with the revised ICS
- View the response time information in the Workload Activity report under the report performance groups

You find more information about planning for workload management and migrating from compatibility to goal mode in *OS/390 MVS Planning: Workload Management*.

Reporting Hierarchy

The mode your system or sysplex is running in, compatibility or goal mode, makes a slight difference in the way you control performance management with RMF. You can see the difference with the help of 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. You may start with this level of report if you are running a single MVS system.

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.

The workload-related sysplex reports are available only in goal mode. All other reports are available in both modes, and let you manage your system's performance in a similar way. The only difference is, that compatibility mode reports show performance data by domain and performance group, and goal mode reports by workload groups and service classes.

Chapter 7. Performance Monitoring

What to Watch For

This chapter describes what to do from day to day to keep your finger on the pulse of the system, so to speak.

If you do this, 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

If you have set performance goals for your workload within a sysplex by means of WLM, 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.

If you want to evaluate the attainment of performance goals for time intervals in the past, use the

Postprocessor - Workload Activity Report

This sysplex report is an extension of the single-system Workload Activity report that is familiar from previous versions of RMF.

It 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.

The “traditional” version is still available as single-system report for systems running in compatibility mode.

As long as your system is running in compatibility mode, you do not specify performance goals within WLM. But you will have defined, explicitly or implicitly, the targets you want to reach to offer the optimum service to your users.

Monitor III - Workflow/Exceptions Report

This can be the report to start with when monitoring the performance of each of your systems. With the *workflow* or *speed* indicator, you get direct information how well a system or a specific application, defined by domains or performance groups, is performing.

This report is also available in goal mode, and shows data for the active service classes.

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.

If you want to get the answer for a sysplex running in goal mode, you will call the

Monitor III - Sysplex Summary Report

The report displays the average response time for each service 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

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

The report shows the average response time for all domains, performance groups or 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.

Another aspect you might be interested in is the response time for specific transactions. This you get with the help of the

Monitor II - Transaction Activity Report

The average response times for all TSO transactions, which you have defined in your IEAIPSxx Parmlib member, and all TSO and batch performance periods, are shown here.

Note: This report is available only in compatibility mode.

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

Depending on the mode in which you are running your systems, this can be a:

- Sysplex report in goal mode: 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.
- Single-system report in compatibility mode: response time values for domains and performance groups.

Monitoring Throughput

There are various indicators that show throughput values.

Transaction Rate

The *classical* definition is

$$Throughput = \frac{Number\ of\ transactions \quad or \quad Number\ of\ jobs}{Time}$$

You get this information in various reports.

Monitor II Reports

Data is available for address spaces, TSO and batch performance periods and TSO transactions:

- Address Space SRM Data report

  Here you get the transaction count and the total session time for each address space.
- Transaction Activity report

  This report shows (in compatibility mode only) the transaction rate for all TSO and batch performance periods.

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 domains, performance groups, or service classes.

- Group Response Time report.

  This report shows detailed values for one selected performance group, domain or service class period.

Postprocessor Reports

You might start with the

- Workload Activity report.

  On either system level (compatibility mode) or sysplex level (goal mode), you get transaction rates for all domains, performance groups, workloads, or service classes.

- Summary report.

  In compatibility mode you see the transaction rate for the total system for each interval of your reporting period.

- Exception report.

  You can define exception criteria for transaction rates of specific domains or performance groups.

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:

$$Value = \frac{Using\ samples}{Using\ samples + Delay\ samples} \times 100$$

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

And you find velocity data in the

- Sysplex Summary report

Postprocessor Reports

The

- Workload Activity report

shows the velocity values. You can use the values reported in compatibility mode to define the goals when you are planning the migration to goal mode.

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 have a look at *exceptions* or *delays* that might be the source of a performance problem. Monitor III is best supports this approach.

To define your performance exceptions directly, use the

Monitor III - Workflow/Exceptions Report

There are many types of exception 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 exception, or thresholds, on the basis of CPU, I/O, workload (compatibility mode only), 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 domain, performance group or service class. 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 PM of OS/390.

PM of OS/390 - PerfDesks

PM of OS/390 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 for systems running in goal mode, 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

In Case of a Problem ...

Monitoring your systems, as described in the previous chapter, should be an ongoing process. Hopefully, analyzing performance problems will be a task that has to be performed only from time to time — but the full capability of RMF will help you here, too.

This chapter

- 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

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

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

You will find a very detailed description on how to analyze and solve performance problems in *RMF Performance Management Guide*, so this chapter gives only a high-level overview of the different reports that are available for performance analysis.

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 delay values in your system.
- Job report

  An address space can be delayed for one of 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 report 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.

The reason for a performance problem can also be the overutilization 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:

Postprocessor Reports

Some of the reports you might use are the:

- CPU Activity report
- Channel 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.

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

Monitor II Reports

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

If you see overutilization of resources because of:

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

you might be interested in evaluating your performance inhibiters.

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.

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:

Postprocessor Workload Activity Report

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

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

Postprocessor Trace Activity Report

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.

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.

Spreadsheet Reporter Macros

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

- Process RMF Summary Report
- Create RMF DASD Report
- Create Workload, CPU, Paging Report
- System Week-Summary Report
- Workload Trend Report
- Device Trend Report

Chapter 9. System Tuning

RMF and Tuning

Probably one of the key aspects in providing service to your users is offering optimum response-time and throughput values; the best service possible with the resources available at your installation.

This requires optimum usage of these resources. RMF provides data to measure and analyze this usage, as discussed in the previous chapter, but more is required than just measurement: *tuning your system*.

Tuning as Iterative Process

In the context of tuning, you might hear about *balanced systems*. What does this mean? A system has many components, such as processors, channels, I/O devices and different types of storage. All these components participate in the task of processing your applications. To get the best results, you should ensure that all of them have comparable power and capability. Of course, you would not configure a large processor with 4 DASD devices, because you would never get the processing power you expect due to bottlenecks in I/O processing. To take the other extreme, an I/O environment of 3000 DASD devices would probably not utilized efficiently by a small processor.

The message is: *the system has to be balanced*. If you run into capacity problems with your processor (100% busy all the time with long queues), in most cases it is not sufficient just to increase the processing power without enlarging the I/O environment. If you do so, you might not be able to run the workload that you would expect to be processed after an upgrade.

How can you find out about the capability and capacity of the different resources? Each one allows a different utilization, and an overutilization creates queues; one of the key causes of performance problems. You might run a processor with 95% utilization, but you never would never be able to get 95% utilization of a disk volume or a channel.

Typical Tuning Approach

If you see tuning as *balancing capacity*, you should consider these steps:

- Measuring performance (for example, response times)
- Measuring utilization of all (or key) resources
- Applying capacity rules
- Reacting to results with some hardware or software reconfigurations
- Measuring performance (for example, response times)

  If you are satisfied with the results: congratulations, otherwise:
- Measure utilization of all (or key) resources
- Apply capacity rules
- ...
- ...

This approach might look easier than it is, as you will agree if you ever have tried it. Especially, what are *capacity rules*? There are no hard and fast rules, like *use this device with up to 34.2% utilization and the performance is good, higher utilization will create problems*. In any specific situation, there are always many components and dependencies that define what is *good* and *bad* performance. By the way, this is not only a question of your resources, but also of your expectations. But we might discuss that another time.

There may be no hard and fast rules, but there are *rules-of-thumb*. You can use them to get a good feel of whether the data you have measured are in a range that indicates performance problems, or whether the numbers tell you that this component of the system is in good shape.

You will find many of these rules-of-thumb in the *RMF Performance Management Guide*. This book covers processors, storage, and I/O environment and provides information on interpreting measured data and using them as performance indicators.

Chapter 10. Capacity Planning

Looking Ahead...

The previous chapters concentrated on allocating existing resources to meet service objectives. Capacity planning is a means of predicting the resources needed to meet future service objectives.

This chapter shows you how to gather data that show up the trends that are happening in your system.

This enables you to react in time to developments that demand a more rigorous management of existing resources, or the procurement of new hardware or software.

Some Aspects to be Considered

Capacity Planning

A process of planning for sufficient computer capacity in a cost-effective manner to meet the service needs of all users.

Capacity planning involves asking the following questions:

- What proportion of your computer resources is being used (processor, storage, I/O devices)?
- Which workloads or applications are consuming the resources?
- What are the expected growth rates?
- When will the demands on current resources impact service levels?

An effective capacity planning process provides:

- A mapping of business objectives (user requirements) to quantifiable *information technology (IT)* resources.
- Management oriented reporting of service, resource usage, and cost. This quantifies what is involved in providing users with good performance.
- Input for making business decisions which involve IT.
- A way to avoid surprises.

There are different methods of capacity planning, for example:

- Rules-of-thumb
- Comparison with other systems
- Parametric methods
  - Transaction profile (10 read calls, 2 update calls, 8 physical I/Os)
  - Cost of function (CICS: 15ms per physical I/O in a 3390)
- Analytic (queuing theory) models: For example the IBM capacity planning tool CP2000
- Simulation, using a computer program that has the essential logic of the system: for example the Snap/Shot modelling system from IBM
- Benchmarks, *Teleprocessing Network Simulator (TPNS)*

RMF Provides Data for Capacity Planning

In addition, capacity planning requires input data from the system; in particular, workload and utilization data. Many tools and programs work with the SMF records that are created by RMF. Of course, you can evaluate the RMF reports directly that cover longer time ranges.

Postprocessor Duration Report

This report provides summary data for intervals of up to 100 hours that can be used as basis for capacity planning.

As mentioned above, details about resource consumption for the different workloads in your system are needed to make projections. Typically, you would use the

Postprocessor Workload Activity Report

You can get values on a very detailed level for performance groups, service classes or workloads for consumption of processor, I/O, and storage resources.

RMF offers you valuable capabilities in evaluating the gathered performance data.

Performance Data Base

The Postprocessor Overview report is comprehensive, and provides additional data that you can use for your own applications.

You might define your own performance data base with the resulting records, which you can tailor to meet your requirements.

You can also exploit the spreadsheet support within RMF.

Spreadsheet Reporter

The Spreadsheet Reporter enables you to transfer Postprocessor data from the host system to your workstation for further processing, and to store them in spreadsheets. There, you can exploit the capabilities of your spreadsheet application to create reports and perform the calculations that you need for analyzing and planning purposes.

In addition, a set of powerful macros is available to help you in displaying and analyzing the key data of your system.

In addition to RMF, a data reduction tool such as *Service Level Reporter (SLR)* or *Performance Reporter for MVS* can be helpful. This kind of tool can help you to review large quantities of data, identify peaks and trends, and correlate data from different sources.

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 11. Long-term Data Gathering with Monitor I

About Monitor I Session Options

This chapter 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 4-2)
- The PARM field of the EXEC statement in the RMF cataloged procedure (see "Setting Up RMF Control Session including Monitor I and Monitor II" on page 2-5)
- The RMF Monitor I Parmlib member ERBRMF00 (see page "ERBRMF00" on page 2-17)

RMF merges the input to a final set of options for the session. See Chapter 5, "How RMF Processes Session Options" on page 5-1 for details.

Summary of Session Options

Figure 11-1 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

Figure 11-1 (Page 1 of 2). Summary of Monitor I Session Options

Option	Description	Details on
Activity Options		
CACHE/NOCACHE	Cache activity	page 11-4
CHAN/NOCHAN	Channel path activity	page 11-4
CPU/NOCPU	Processor activity	page 11-5
DEVICE(type)/NODEVICE	Device activity	page 11-5
{(SUMMARY)} ENQ{(DETAIL[,majorname[,minorname]])}/NOENQ	Enqueue contention activity	page 11-7
IOQ (option,option...)/NOIOQ	I/O queuing activity	page 11-10
PAGESP/NOPAGESP	Page/swap data set activity	page 11-14
PAGING/NOPAGING	System paging activity	page 11-14
TRACE(variable[,opt list])/NOTRACE	Trace variables for the Trace Activity report	page 11-17
{(S)} VSTOR{(D [,jobname1,jobname2,...])}/NOVSTOR	Virtual storage activity	page 11-21
WKLD(list)/NOWKLD	Workload activity	page 11-22
Time-frame Options		
{(1000)} CYCLE{(nnn)}	The length of the cycle at the end of which RMF makes sampling observations	page 11-5
({30M}) INTERVAL({nnn[M]})	The length of the reporting interval in minutes used in combination with the SYNC(RMF,mm), SYNC(RMF,mmM), or NOSYNC options	page 11-8
[M] STOP (value[H])/NOSTOP	Desired duration of the Monitor I session, in minutes (M), or hours (H)	page 11-15
{(SMF)} {(RMF,mm)} SYNC {(RMF,mmM)}/NOSYNC	Interval synchronization with the SMF or the RMF interval synchronization on the minute	page 11-16
Reporting Options		
{OPTIONS}/{NOOPTIONS} {OPTN } {NOOPTN }	Option list for the session to be printed at the operator console	page 11-12
RECORD/NORECORD	Specifies whether measured data is to be written to the SMF data set	page 11-14
{(REALTIME)} REPORT{(DEFER)}/NOREPORT	Specifies production of printed interval reports of measured data	page 11-15
SYSOUT(class)	SYSOUT class to which the formatted printed reports are directed	page 11-17

Figure 11-1 (Page 2 of 2). Summary of Monitor I Session Options		
Option	**Description**	**Details on**
	Environment Options	
EXITS/NOEXITS	User exit routines to be executed during session processing to gather or report additional data	page 11-8
MEMBER(list)	Parmlib member containing Monitor I session options	page 11-11

Default Session Options

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

Figure 11-2. Monitor I Default Session Options	
Default Option	**Description**
CACHE	Measure cache activity
CHAN	Measure channel path activity
CPU	Measure processor activity
DEVICE(DASD)	Measure DASD device activity (not other classes of device)
IOQ(DASD)	Measure I/O queuing activity on logical control units for DASD devices
PAGESP	Measure page/swap data-set activity
PAGING	Measure system paging activity
RECORD	Write the measured data to the SMF data set
VSTOR(S)	Measure virtual storage activity for summary reports
WKLD(PERIOD)	Measure system workload (if reporting is specified for a system running in compatibility mode, it is done for performance group periods)
CYCLE(1000)	Take sample measurements once every second (1000 milliseconds)
INTERVAL(30M)	Combines data every 30 minutes (value is ignored for SYNC(SMF))
STOP(8H)	End the session 8 hours after it was started
SYNC(SMF)	Synchronize the reporting interval with SMF
NOENQ	Do not measure contention activity
NOEXITS	Execute no user exits when gathering and reporting
NOREPORT	Do not produce printed interval reports
NOTRACE	Do not trace any variables (no Trace Activity report)
OPTIONS	Print a list of session options at the operator console at the start of the session, allowing the operator to change options. In the interests of a fast start-up of Monitor I, we recommend that you specify NOOPTIONS unless changes at start-up are really necessary.

SMF Records

Monitor I creates SMF records type 70 — 78, you find an overview in “Activity Monitoring” on page 1-8.

Description of Monitor I Options

CACHE

```
     ┌─CACHE───┐
►►───┤         ├──────────────────────────────────────────────────────►◄
     └─NOCACHE─┘
```

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. Please, refer to the example (page 2-16) that 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

```
     ┌─CHAN───┐
►►───┤        ├───────────────────────────────────────────────────────►◄
     └─NOCHAN─┘
```

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.

In an Enterprise Systems Connection (ESCON*) Multiple Image Facility (EMIF) environment running in LPAR mode, if CHAN is specified, channel utilization for the individual partitions is measured in addition to total channel utilization.

CPU

```
     ┌─CPU───┐
►►───┤       ├──────────────────────────────────────────────────────►◄
     └─NOCPU─┘
```

Specifies processor activity measurement.

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

CYCLE

```
     ┌─CYCLE(1000)──┐
►►───┤              ├───────────────────────────────────────────────►◄
     └─CYCLE(value)─┘
```

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 11-23 for considerations that apply to choosing a cycle length.

DEVICE

```
     ┌─DEVICE(DASD)──────────────┐
►►───┤                           ├──────────────────────────────────►◄
     │          ┌─,────────┐     │
     ├─DEVICE(──▼─┤ Type ├─┴──)──┤
     └─NODEVICE──────────────────┘

Type:
   ┌─NONMBR─────┐ ┌─DASD───┐ ┌─NOCHRDR─┐ ┌─NOCOMM─┐ ┌─NOGRAPH─┐ ┌─NOTAPE─┐
├──┤            ├─┤        ├─┤         ├─┤        ├─┤         ├─┤        ├──►
   └─NMBR(nmbr)─┘ └─NODASD─┘ └─CHRDR───┘ └─COMM───┘ └─GRAPH───┘ └─TAPE───┘
   ┌─NOUNITR─┐
►──┤         ├──────────────────────────────────────────────────────────┤
   └─UNITR───┘
```

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/NOCHRDR Character reader devices

  COMM/NOCOMM Communications equipment

  DASD/NODASD Direct access storage devices

  GRAPH/NOGRAPH Graphics devices

  TAPE/NOTAPE Magnetic tape devices

  UNITR/NOUNITR Unit record devices

-

  Note: 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 *OS/390 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

```
      ┌─NOENQ────────────────────────────┐
►►────┼──────────────────────────────────┼──────────────────────────────────────────►◄
      └─ENQ(─┬─SUMMARY─────────────────┬─)─┘
             └─DETAIL─┬────────────────┤
                      └─,major─┬───────┤
                               └─,minor┘
```

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.

EXITS

```
   ┌─NOEXITS─┐
►►─┴─EXITS───┴──────────────────────────────────────────────────►◄
```

Specifies whether Monitor I user exit routines are to be executed during session processing to gather and report on additional data. See the *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.

INTERVAL

```
   ┌─INTERVAL(30M)───────┐
►►─┴─INTERVAL(nnn─┬───┬─)─┴─────────────────────────────────────►◄
                  └─M─┘
```

Specifies the length of the Monitor I reporting interval, where **n** 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 2-11 for more information.

Notes:

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 RMF Control Session including Monitor I and Monitor II” on page 2-5 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

```
      ┌─IOQ(─┤ Options ├─)────┐
▶▶────┤                        ├──────────────────────────────────────────────────▶◀
      │       ┌─,──────────┐   │
      ├─IOQ(──▼┤ Options ├─┴─)─┤
      └─NOIOQ──────────────────┘
```

Options:

```
    ┌─NONMBR────┐  ┌─DASD───┐  ┌─NOCHRDR─┐  ┌─NOCOMM─┐  ┌─NOGRAPH─┐  ┌─NOTAPE─┐
├───┤           ├──┤        ├──┤         ├──┤        ├──┤         ├──┤        ├──────▶
    └─NMBR(nmbr)┘  └─NODASD─┘  └─CHRDR───┘  └─COMM───┘  └─GRAPH───┘  └─TAPE───┘
    ┌─NOUNITR─┐
▶───┤         ├───────────────────────────────────────────────────────────────────┤
    └─UNITR───┘
```

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 (IOP, or SAP on a

9672 system), 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 an MVS 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

```
     ┌─MEMBER(00)──────┐
►►───┼─────────────────┼──────────────────────────────────────────────►◄
     │         ┌─,──┐  │
     └─MEMBER(─▼─xx─┴(1)─)─┘
```

Note:
[1] Up to five members can be specified.

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

For the Monitor I session, the default is 00, indicating Parmlib member ERBRMF00. The contents of ERBRMF00 are described in Storing Gatherer Options. If you have created your own Parmlib data set, make sure you specify it in the RMF cataloged procedure. See “Setting Up RMF Control Session including Monitor I and Monitor II” on page 2-5.

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 5-1.

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

OPTIONS

```
    ┌─OPTIONS───┐
►►──┼───────────┼──────────────────────────────────────────────────────►◄
    ├─OPTN──────┤
    ├─NOOPTIONS─┤
    └─NOOPTN────┘
```

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 11-3 shows each possible option followed by its input source.

Figure 11-3. Monitor I 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.
REPLY	The option was changed from the operator console in reply to message ERB306I.

```
ERB305I ZZ : PARAMETERS
ERB305I ZZ :  WKLD (PERIOD) -- DEFAULT
ERB305I ZZ :  SYNC(SMF) -- DEFAULT
ERB305I ZZ :  TRACE (RCVPTR,ALL) -- MEMBER
ERB305I ZZ :  TRACE (RCVCPUA,ALL) -- MEMBER
ERB305I ZZ :  TRACE (RCVUICA,ALL) -- MEMBER
ERB305I ZZ :  NOEXITS -- MEMBER
ERB305I ZZ :  ENQ(DETAIL) -- MEMBER
ERB305I ZZ :  OPTIONS -- COMMAND
ERB305I ZZ :  STOP(168H) -- MEMBER
ERB305I ZZ :  CYCLE(250) -- MEMBER
ERB305I ZZ :  RECORD -- MEMBER
ERB305I ZZ :  SYSOUT(A) -- MEMBER
ERB305I ZZ :  REPORT(REALTIME) -- MEMBER
ERB305I ZZ :  PAGESP -- MEMBER
ERB305I ZZ :  DEVICE(NOCHRDR) -- DEFAULT
ERB305I ZZ :  DEVICE(NONMBR) -- DEFAULT
ERB305I ZZ :  DEVICE(GRAPH) -- MEMBER
ERB305I ZZ :  DEVICE(COMM) -- MEMBER
ERB305I ZZ :  DEVICE(DASD) -- MEMBER
ERB305I ZZ :  DEVICE(TAPE) -- MEMBER
ERB305I ZZ :  DEVICE(UNITR) -- MEMBER
ERB305I ZZ :  DEVICE(NOSG) -- DEFAULT
ERB305I ZZ :  CACHE -- DEFAULT
ERB305I ZZ :  CHAN -- MEMBER
ERB305I ZZ :  CPU -- DEFAULT
ERB305I ZZ :  PAGING -- MEMBER
ERB305I ZZ :  IOQ(NOCHRDR) -- DEFAULT
ERB305I ZZ :  IOQ(NONMBR) -- MEMBER
ERB305I ZZ :  IOQ(GRAPH) -- MEMBER
ERB305I ZZ :  IOQ(COMM) -- MEMBER
ERB305I ZZ :  IOQ(DASD) -- MEMBER
ERB305I ZZ :  IOQ(TAPE) -- MEMBER
ERB305I ZZ :  IOQ(UNITR) -- MEMBER
ERB305I ZZ :  VSTOR(S) -- MEMBER
ERB305I ZZ :  MEMBER(07) -- COMMAND
```

Figure 11-4. Console Output with Monitor I OPTIONS in Effect

PAGESP

```
      ┌─PAGESP───┐
►►────┤          ├──────────────────────────────────────────────────►◄
      └─NOPAGESP─┘
```

Specifies whether page/swap data set activity is to be measured.

PAGING

```
      ┌─PAGING───┐
►►────┤          ├──────────────────────────────────────────────────►◄
      └─NOPAGING─┘
```

Specifies whether system paging activity is to be measured.

RECORD

```
      ┌─RECORD───┐
►►────┤          ├──────────────────────────────────────────────────►◄
      └─NORECORD─┘
```

Specifies whether measured data is to be written to the SMF data set. In order for RECORD to take effect, the complementary SMF enabling procedures must first be performed. These procedures are described in *OS/390 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

```
        ┌─NOREPORT──────────────────┐
►►──────┤                           ├──────────────────────────────────────────►◄
        └─REPORT(──┬─REALTIME─┬──)──┘
                   └─DEFER────┘
```

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.

Notes:

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

```
        ┌─STOP(8H)─────────────────┐
►►──────┤                          ├───────────────────────────────────────────►◄
        │                 ┌─M─┐    │
        ├─STOP(range──────┤   ├─)──┤
        │                 └─H─┘    │
        └─NOSTOP───────────────────┘
```

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

```
   ┌─SYNC(SMF)──────────────┐
►►─┼────────────────────────┼──────────────────────────────────────────────►◄
   ├─SYNC(RMF,mm─┬───┬─)────┤
   │             └─M─┘      │
   └─NOSYNC─────────────────┘
```

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.

Notes:

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 2-11 for more information.

SYSOUT

```
      ┌─SYSOUT(A)──────┐
►►────┤                ├──────────────────────────────────────────────────►◄
      └─SYSOUT(class)──┘
```

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

TRACE

```
      ┌─NOTRACE───────────────────────────────┐
►►────┤                                       ├───────────────────────────►◄
      │                   ┌─ALL──────────────┐│
      │                   │ ┌─,──────────┐   ││
      └─TRACE(variable────┴─▼─┴─opt list─┴───┴┴─)─┘
```

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
CCVWTDB	number of wait-time dispatches when the LPAR is the guest

The following group of variables, beginning with DMD, are available in compatibility mode only:

DMDTCIDX	domain contention index * 100
DMDTCMPL	current multiprogramming level
DMDTINCU	users per domain in storage
DMDTOUTU	users per domain out of storage
DMDTRUA	average number of ready users * 16
DMDTTWET	average elapsed time in 1024 microsecond units for first period TSO transaction (does not include delay imposed by SRM to meet response time objective)

DMDTTWSR	time weighted service rate for a domain

Note: The domain trace variables (DMDT*xxxx*) specified are different from the variables reported on the trace activity report. Tracing the domain variables causes the domain numbers to be appended to an abbreviation of the variable on the report as follows, where *xxx* is the domain number:

Specified	Reported
DMDTCIDX	DMCDXxxx
DMDTCMPL	DMMPLxxx
DMDTINCU	DMINUxxx
DMDTOUTU	DMOUTxxx
DMDTRUA	DMRUAxxx
DMDTTWET	DMTWExxx
DMDTTWSR	DMTWSxxx

The remaining variables are available in goal mode and compatibility mode:

LSCTCNT	current number of logically swapped users for terminal wait
LSCTMTE	maximum think time allowed for logical swap candidate
MCVFRCNT	number of pages needed to be stolen by force steal routine
MCVMGAGE	expanded storage migration age
MCVSBLTF	long term percentage of eligible storage that is actually fixed
MCVSIPR	common page-in rate
MCVSTCRI	highest system UIC
MCVTWSS	common target working set size
OMDGAMRE	maximum number of messages on the action message retention facility (AMRF) queue. If a large number of action messages are retained on the AMRF queue for a particular period, it may mean more operators are needed for that period.
OMDGCMDI	number of commands issued per second.
OMDGOREB	maximum number of operator reply entries (OREs) on the system reply queue. To eliminate thrashing, use this number to monitor and adjust the ORE buffer limit set at IPL time. To dynamically adjust this limit, use the CONTROL M command.
OMDGWQEB	maximum number of WTO queue elements (WQEs) on the system output queue. To eliminate thrashing (excessive data movement which confines system to doing little useful work), use this number to monitor and adjust the WTO buffer time limit set at IPL time. To dynamically adjust this limit, use the CONTROL M command.
OMDGWTLI	number of write-to-logs (WTLs) issued per second, indicating the number of records going to SYSLOG within a time period. To control the number of data sets produced during the day, vary the number of records per SYSLOG data set.
OMDGWTOI	total number of lines of messages, write-to-operators (WTOs) issued per second. Use it to determine the peak message rate period and the average message rate.
RAXESCT	number of common storage pages on expanded storage
RAXFMCT	number of frames allocated to common storage
RCEAEC	total number of expanded storage E frames currently on the ESTE queue
RCEAECLO	available expanded storage low threshold

RCEAECOK	available expanded storage satisfactory threshold
RCEAFC	total number of frames currently on all available frame queues
RCEAFCLO	available central storage low threshold
RCEAFCOK	available central storage satisfactory threshold
RCEBELFX	total number of fixed pages below 16 megabytes in central storage, which is the sum of page-fixed LSQA, SQA (excluding reserved SQA) and V=R allocated pages.
RCECOMPI	number of common area pages paged-in
RCECOMPO	number of common area pages paged-out
RCEDFRS	number of times a deferred frame allocation has been satisfied
RCEESINU	number of in-use expanded storage frames
RCEESREA	number of non-VIO pages read from expanded storage
RCEESWRT	number of pages written to expanded storage frames
RCEHSPEM	total number of hiperspace* pages migrated from expanded storage to auxiliary storage
RCEHSPER	total number of hiperspace pages in the system read from expanded storage to central storage
RCEHSPEW	total number of hiperspace pages written from central storage to expanded storage
RCEHSPPI	total number of hiperspace pages paged in from auxiliary storage
RCEHSPPO	total number of hiperspace pages paged out to auxiliary storage
RCELPAPI	number of PLPA and PLPA directory pages paged-in
RCEMVBEL	number of pages moved from below 16 megabytes in central storage
RCENWSF	total number of secondary and non-working set pages migrated to auxiliary storage.
RCEPAGMV	number of times a frame was moved from one frame to another
RCEPOOL	number of frames currently available to the system, including frames backing permanent storage (nucleus frames, hardware storage area frames, FLPA frames or fixed BLDL frames), bad frames and offline frames
RCESPFR	number of frames available by swap-out without requiring I/O
RCESWPPI	total number of pages requiring I/O to swap-in
RCESWPPO	total number of pages requiring I/O to swap-out
RCETOTFX	total number of pages currently fixed, the sum of page-fixed LSQA, SQA (excluding reserved SQA) and V=R allocated pages
RCETOTPI	total number of pages paged-in excluding swap-in and VIO page-in
RCETOTPO	total number of pages paged-out, excluding swap-out, move-out of VIO pages, and page-out of VIO pages
RCEVIOME	number of VIO pages written to expanded storage
RCEVIOMG	number of VIO pages migrated from expanded storage to paging data sets
RCEVIOPI	total number of VIO pages paged-in, excluding swap-in
RCEVIOPO	total number of VIO pages, excluding swap-out, moved out or paged-out
RCEVIORE	number of VIO reads from extended storage
RCEWSDNE	total number of primary working set pages migrated to auxiliary storage
RCVAFQA	average available frame count
RCVAVQC	AVQ low count
RCVCPUA	CPU usage average * 16
RCVFXIOP	percentage of central storage that is fixed or allocated for paging

RCVMFXA	average number of fixed frames for the system
RCVPAGRT	total paging rate
RCVPTR	paging rate
RCVSWPTM	time (in milliseconds) used by ASM to process a request to transfer a group of pages to or from a data set
RCVUICA	UIC average
RMCAAWSC	APPC/MVS transaction scheduler (ASCH) wait swap count
RMCADWSC	detected wait physical swap count
RMCAEXSC	exchange on recommendation value swap count
RMCAFHLD	number of swaps failed because of an outstanding HOLD SYSEVENT
RMCAICSC	improve central storage use
RMCAIPSC	improve system paging rate
RMCALWSC	long wait physical swap count
RMCAMRSC	make room to swap in a user who was swapped out too long.
RMCANQSC	CPU enqueue exchange swap count
RMCAOISC	OMVS input wait
RMCAOOSC	OMVS output wait
RMCARQSC	requested swap count
RMCARSSC	central storage shortage swap count
RMCATISC	terminal input swap count
RMCATOSC	terminal output swap count
RMCATSSC	count of transition swaps
RMCAUSSC	unilateral swap out count
RMCAXSSC	auxiliary storage shortage swap count
RMCTTRPC	number of pages used for transaction elements
SMCABFLS	number of records lost because of a shortage of buffers
SMCABFWT	number of buffers written
SMCADSCT	number of records lost because of a full data set
SMCANMFL	current number of full buffers
SMCARCWT	number of records written

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

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

VSTOR

```
     ┌─VSTOR(S)──────────────────────┐
►►───┼───────────────────────────────┼────────────────────────────────────────────►◄
     ├─VSTOR(D)──────────────────────┤
     │              ┌─,──────────┐   │
     ├─VSTOR(D,───▼─jobname─(1)──┴─)─┤
     └─NOVSTOR───────────────────────┘
```

Note:
1 You can specify up to 25 jobnames

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.

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: Please, note that 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

```
    ┌─WKLD(─┤ List ├─)──────┐
►►──┼───────────────────────┼──────────────────────────────────────────────────────────►◄
    │      ┌─,─────────┐    │
    ├─WKLD(─▼─┤ List ├─┴─)──┤
    └─NOWKLD────────────────┘
List:
   ┌─PERIOD───┐  ┌─NOGROUP─┐  ┌─NORANGE─┐  ┌─NODOMAIN─┐  ┌─NOSYSTEM─┐  ┌─NOTIME─┐
├──┼──────────┼──┼─────────┼──┼─────────┼──┼──────────┼──┼──────────┼──┼────────┼─────────┤
   └─NOPERIOD─┘  └─GROUP───┘  └─RANGE───┘  └─DOMAIN───┘  └─SYSTEM───┘  └─TIME───┘
```

This option has two purposes:

1. To specify whether the system workload is to be measured. `WKLD` is the default, so measuring will be done automatically, unless you specify `NOWKLD`.
2. To specify the type of reporting to be done if you have defined `REPORT(REALTIME)` or `REPORT(DEFER)` (as opposed to `NOREPORT`). This is possible only in compatibility mode, in goal mode the following options are ignored.

When WKLD is specified, **list** must indicate the type of reporting required. Any or all of the following can be specified in **list**:

PERIOD/NOPERIOD	Reporting by performance group period
GROUP/NOGROUP	Reporting by performance group
RANGE/NORANGE	Reporting by performance objective, domain, and performance group in a range of performance group numbers
DOMAIN/NODOMAIN	Reporting by domain number
SYSTEM/NOSYSTEM	Reporting by system
TIME/NOTIME	Reporting by time-slice group

When you omit the option, the defaults are as underscored above. Specify WKLD explicitly only if you want to change these values using entries in **list**.

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

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

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

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

STOP, INTERVAL, and REPORT Options

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

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 addition to reducing RMF overhead, limiting the devices can also decrease the amount of SQA storage required by a Monitor I session as well as decrease the length of the SMF device records. In the case of 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 12. Snapshot Data Gathering with Monitor II

Monitor II in the Background

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 16, “Snapshot Reporting with Monitor II” on page 16-1.

Chapter 13. Short-term Data Gathering with Monitor III

Monitor III Reference Material

This chapter tells you:

- The syntax and effect of the options
- How to control VSAM data set recording

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 ERBRMF04. See Storing Gatherer Options for its contents, and the MEMBER option in “Description of Data Gatherer Options” on page 13-3 for how to specify other Parmlib members.
- The **parm** field of the START session command that starts the session. See “Starting a Specific Monitor” on page 4-2.

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 4-5.
- The response to the OPTIONS option.

Figure 13-1 gives a summary of the Monitor III gatherer session options. The referenced pages describe the options in detail.

Figure 13-1. Monitor III Data Gatherer Session Options

Option	Effect	Details on
CACHE	Defines cache data gathering.	page 13-3
CFDETAIL	Defines level of detail for data gathering for the coupling facility.	page 13-4
CYCLE(nnnn)	Sets the length of the cycle at the end of which RMF samples data.	page 13-4
DATASET(suboption...)	Controls data set recording of sampled data.	page 13-5
HFSNAME(suboption...)	Controls data set recording for OS/390 UNIX file systems.	page 13-6
IOSUB	Controls data set recording of I/O-subsystem and channel-path activity.	page 13-6
MEMBER(list)	Specifies Parmlib members containing session options.	page 13-7
MINTIME(nnn)	Specifies the interval at which data samples are summarized.	page 13-7
OPTIONS	Controls display of the current options at the start of a session.	page 13-8
RESOURCE(...)	Specifies the job entry subsystem (JES) to be used.	page 13-9
STOP(value)	Sets the duration of the data gatherer interval.	page 13-10
SYNC	Synchronizes MINTIME within the hour.	page 13-10
SYSOUT(class)	Specifies the SYSOUT class for gatherer messages.	page 13-11
WSTOR	Sets the size of the RMF local storage buffer.	page 13-11

Default Session Options

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

Figure 13-2. Monitor III Default Session Options

Default Option	Description
CACHE	Defines cache data gathering.
NOCFDETAIL	Defines partial data gathering for the coupling facility.
CYCLE(1000)	Take data samples once a second (1000 milliseconds)
DATASET(STOP,NOSWITCH)	No data-set recording will be done
IOSUB	Record I/O-subsystem and channel-path data
MINTIME(100)	Build a set of samples every 100 seconds
NOOPTIONS	Session options are not displayed at the operator console at the start of the session
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
SYNC(0M)	MINTIME is synchronized on the hour
WSTOR(32)	Sets the RMF local storage buffer to 32 megabytes

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 1-8.

Description of Data Gatherer Options

CACHE

```
      ┌─CACHE───┐
►►────┴─NOCACHE─┴──────────────────────────────────────────────►◄
```

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. To avoid having duplicated data, you should gather cache activity data on one system only. Please, refer to the example (page 2-16) that shows how to set up gathering options.

To suppress the gathering of cache data, specify NOCACHE.

CFDETAIL

```
   ┌─NOCFDETAIL─┐
►►─┤            ├──────────────────────────────────────────────────►◄
   └─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 NOCFDETAIL. To start collection, specify CFDETAIL when starting or modifying the Monitor III session. When you specify CFDETAIL on a MODIFY command, collection starts at the end of the current Mintime.

With CFDETAIL, a large amount of data is being gathered that enables you to get many details about the usage of each structure in the coupling facility. Please, consider that this data gathering is done only on one member of the sysplex. This is called *sysplex master gathering* and has been implemented to reduce performance overhead on non-master members and to reduce the amount of data in SSHs and SMF records. The RMF Sysplex Data Server determines internally which member of the sysplex will be the master. This cannot be controlled externally by the operator or system administrator and is transparent for performance monitoring. Therefore, it is very important that the MODIFY command has to be routed to all members:

```
ROUTE *ALL,MODIFY RMF,MODIFY III,CFDETAIL
```

CYCLE

```
   ┌─CYCLE(1000)─┐
►►─┤             ├─────────────────────────────────────────────────►◄
   └─CYCLE(nnnn)─┘
```

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.

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

Sysplex Reporting

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

DATASET

```
►►──┬─DATASET(STOP,NOSWITCH)──────┬──────────────────────────────────────────►◄
    │      ┌─,───────────────┐   │
    └─DS(──▼─┤ Suboption ├──┴─)─┘

Suboption:
├──┬──┬─┬──────────────────────────┬─┬──┬───────────────────────────────────────┤
   │  │ │      ┌─,───────────────┐ │ │  │
   │  │ ├─ADD──▼─(data-set-name)─┴─┤ │  │
   │  │ │      ┌─,───────────────┐ │ │  │
   │  │ └─DEL──▼─(data-set-name)─┴─┘ │  │
   ├──────START───────────────────────┤
   ├──────SWITCH──────────────────────┤
   │             ┌─(7)─────┐          │
   └────WHOLD────┴─(value)─┴──────────┘
```

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 13-12.

HFSNAME

```
►►──┬────────────────────────────────────┬──────────────────────────────────►◄
    │           ┌─,───────────────┐      │
    └─HFSNAME(──▼──┤ Suboption ├──┴──)───┘

Suboption:
├──┬──────────────────────────────┬──────────────────────────────────────────┤
   │      ┌─,───────────────────┐ │
   ├─ADD──▼─(file-system-name)──┴─┤
   │      ┌─,───────────────────┐ │
   └─DEL──▼─(file-system-name)──┴─┘
```

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

```
    ┌─IOSUB───┐
►►──┤         ├──────────────────────────────────────────────────────────────►◄
    └─NOIOSUB─┘
```

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.

MEMBER

```
      ┌─MEMBER(04)────────┐
►►────┤                   ├──────────────────────────────────────►◄
      │        ┌─,─┐      │
      └─MEMBER(─▼xx─┴──)──┘
```

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

For the Monitor III gatherer session, the default is 04, indicating Parmlib member ERBRMF04. If you have created your own Parmlib, make sure you specify it on the IEFRDER DD statement in the RMF cataloged procedure. See "Setting Up RMF Control Session including Monitor I and Monitor II" on page 2-5.

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

ERBRMF04 contains the Monitor III gatherer option:

```
RESOURCE(*JES2,JES2)
```

MINTIME

```
      ┌─MINTIME(100)─┐
►►────┤              ├──────────────────────────────────────►◄
      └─MINTIME(nnn)─┘
```

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

Please 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 2-11 for more information about using MINTIME values to synchronize Monitor I and III recording intervals.

OPTIONS

```
      ┌─NOOPTIONS─┐
►►────┼───────────┼──────────────────────────────────────────────────────────►◄
      ├─OPTIONS───┤
      ├─OPTNS─────┤
      └─NOOPTNS───┘
```

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 13-4 on page 13-9 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 13-12.)

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:

Figure 13-3. Monitor I 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 :  NOCFDETAIL -- 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 :  OPTIONS -- COMMAND
ERB305I  III :  RESOURCE(*JES2,JES2) -- MEMBER
ERB305I  III :  SYSOUT(A) -- DEFAULT
ERB305I  III :  MEMBER (04) -- COMMAND
```

Figure 13-4. Console Output with OPTIONS in Effect

RESOURCE

```
      ┌─RESOURCE(*JES2,JES2)───────┐
►►────┼────────────────────────────┴──────────────────────────────────►◄
      ├─RESOURCE(*JES2─┬───────┬─)─┤
      │                └─,parm─┘   │
      └─RESOURCE(*JES3─┬───────┬─)─┘
                       └─,parm─┘
```

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

```
    ┌─NOSTOP──────────────┐
►►──┤                     ├────────────────────────────────────────────►◄
    └─STOP(value─┬─M─┬─)──┘
                 └─H─┘
```

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

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

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

SYNC

```
    ┌─SYNC(0M)──────────┐
►►──┤                   ├──────────────────────────────────────────────►◄
    ├─SYNC(mm──┬───┬─)──┤
    │          └─M─┘    │
    └─NOSYNC────────────┘
```

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 2-11 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

```
      ┌─SYSOUT(A)──────┐
►►────┤                ├──────────────────────────────────────────────────►◄
      └─SYSOUT(class)──┘
```

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.

WSTOR

```
      ┌─WSTOR(32)───┐
►►────┤             ├─────────────────────────────────────────────────────►◄
      └─WSTOR(nnn)──┘
```

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.

Notes:

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.

Controlling Data Set Recording

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

```
    ┌─DATASET(STOP,NOSWITCH)──┐
►►──┼─────────────────────────┼──────────────────────────────────────────────────►◄
    │      ┌─,─────────────┐  │
    └─DS(──▼─┤ Suboption ├─┴─)┘

Suboption:
├──┬──┬────────────────────────┬─┬────────────────────────────────────────────────┤
   │  │      ┌─,──────────────┐ │ │
   │  ├─ADD──▼(data-set-name)─┴─┤ │
   │  │      ┌─,──────────────┐ │ │
   │  └─DEL──▼(data-set-name)─┴─┘ │
   ├──START───────────────────────┤
   ├──SWITCH──────────────────────┤
   │        ┌─(7)──────┐          │
   └─WHOLD──┼──────────┼──────────┘
            └─(value)──┘
```

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 page “WSTOR” on page 13-11) 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 5-1.

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 15-3.

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 of another system, 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 4-5.

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.

Part 6. Reporting Reference

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

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

Chapter 14. The Online Reporting Environment

Menu and Help

To concentrate the reporting capabilities of RMF, a common ISPF interface gives you access to Monitor II, Monitor III and the Postprocessor.

To further ease the task of reporting, online help is provided for Monitor II and Monitor III dialogs and the Postprocessor.

This chapter tells you about:

- The Performance Management menu
- Online help
- 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               OS/390 2.7.0 RMF
Selection ===>

Enter selection number or command on selection line.


  1 Postprocessor   Postprocessor reports for Monitor I, II, and III      (PP)
  2 Monitor II      Snapshot reporting with Monitor II                    (M2)
  3 Monitor III     Interactive performance analysis with Monitor III     (M3)

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

  R RMFPP           Performance analysis with the Spreadsheet Reporter
  P PM of OS/390    PM of OS/390 using the OS/2 workstation
  N News            What's new in OS/390 2.7 RMF

                              T TUTORIAL    X EXIT

  RMF Home Page: http://www.ibm.com/s390/rmf/

         5647-A01 (C) Copyright IBM Corp. 1994,1999. All Rights Reserved
                     Licensed Materials - Property of IBM
```

Figure 14-1. 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

By selecting `U` or the abbreviation `US`, you gain access to 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 (RMFPP)" on page 18-1.

PM of OS/390 Using the OS/2 Workstation

Performance Monitoring of OS/390 provides an interface between the OS/2 workstation and the OS/390 sysplex that gives you the flexibility to create unique scenarios that monitor the performance of your system.

You find all details in Chapter 19, "Performance Monitoring of OS/390" on page 19-1.

What's New in OS/390 2.7 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 Home Page in the Internet? That's our address:

RMF Home Page: http://www.ibm.com/s390/rmf/

Here, you find 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 *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 for RMF contains two types of highlighted phrases. In CUA* terms, 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 15. Interactive Performance Analysis with Monitor III

Overview of Monitor III Reporting

This chapter tells you how to find your way about 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 chapter tells you

- What to do with the Primary menu
- How to get to the other screens you need
- What 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

Enablement

RMF is an optional feature of OS/390. 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 following the gap.

Note: If you are allocating data sets from a sysplex, it is of key impartance that you allocate **all data sets** of the sysplex to enable complete reporting.

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 2-14), then you would use this allocation:

```
ALLOC FI(RMFDS00) DA('SYS1.ERB.SYSTEMA.VSAM') SHR
ALLOC FI(RMFDS01) DA('SYS1.ERB.SYSTEMB.VSAM') SHR
ALLOC FI(RMFDS02) DA('SYS1.ERB.SYSTEMC.VSAM') SHR
ALLOC FI(RMFDS03) DA('SYS1.ERB.SYSTEMD.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 15-14.

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 an 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)]
```

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*.

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)
```

For details on how to analyze the data in the VSAM data sets, see “Before You Start Monitor III” on page 15-2.

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

The Primary Menu

The first panel that RMF displays in response to the RMF command is the Primary menu:

```
                    RMF Monitor III Primary Menu              OS/390 2.7.0 RMF
Selection ===>

Enter selection number or command on selection line.


  S SYSPLEX         Sysplex reports, Coupling Facility, Cache, Data Index  (SP)

  1 OVERVIEW        Workflow/Exceptions, System information, and delays    (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

       5647-A01 (C) Copyright IBM Corp. 1986, 1999. All Rights Reserved
                     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 15-1. Monitor III 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 Figure 15-20 on page 15-31. These commands are called *report commands*.
3. 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 lists the individual single-system reports.

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 *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 `O` or `OPTIONS` to reach the Option Selection menu. This guides you in specifying the report options for the session.
2. By using *report option commands* to call up the data-entry panels for the options you want to specify. You will find these commands listed in Figure 15-19 on page 15-30.

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`

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 changing the time, date, or range fields in the header of the report panel:

```
                       RMF 2.7.0  TITLE                         Line 1 of 30
Command ===>                                                  Scroll ===> HALF

Samples: nnn     System: syst  Date: mm/dd/yy  Time: hh.mm.ss  Range: 100   Sec
```

Figure 15-2. 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 performance group or 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.

Changing Options

To change the options of an RMF 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                           (SO)
  2 COLOR          Set Graphics Colors and/or Patterns           (CO)
  3 LANGUAGE       Set Language and Date/Time Options            (LO)
  4 ROPTIONS       Select report options for
                     REPORT ===> ________                        (RO)
  5 OPTSET         Change or Select Option Set                   (OS)
```

Figure 15-3. Option Selection Menu

Select what you want to change:

SESSION Specify options that are valid for all reports displayed this session.

COLOR Define the colors you like to have in the reports.

LANGUAGE Define language-specific display of date and time.

ROPTIONS Set or modify 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 Figure 15-20 on page 15-31.

OPTSET Build a set of options and store it for later use. If you build several different sets of options, you can select the appropriate one for a given session.

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

Changing Session Options

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

Figure 15-4 on page 15-8 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 15-12.

From the Session Options panel, you can:

- Select the panel you want to appear when you start an RMF session
- Select the display mode (STOP or GO)
- 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 (HARDCOPY)” on page 15-25.
- 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

```
                           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 15-4. Session Options Panel

For more information about the parameters on this 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* is installed on your MVS system, 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.

Figure 15-7 describes the fields on the color graphic options panel.

```
                         RMF Color Graphic Options              Screen 1 of 2
Command ===>

 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

   1   2   3   4   5   6   7   Colors
```

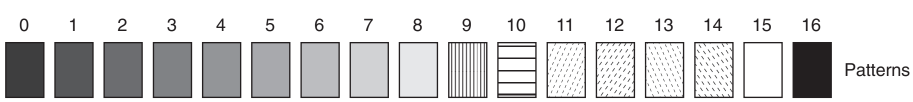

Figure 15-5. Color Graphic Options - Panel 1

```
                         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 ===> ________   ===> _  ===> __
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
```

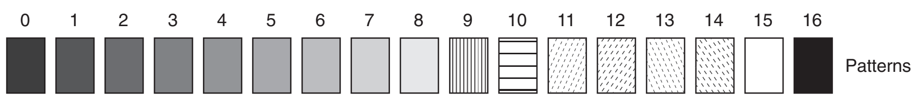

Figure 15-6. Color Graphic Options - Panel 2

Figure 15-7. 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 (CANCEL)” on page 15-23.

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 15-8. Language Options Panel

For more information about the parameters on this panel, 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 compatibility mode, you can specify job groups, and in goal mode 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 15-9 shows a Report Options panel for the DELAY report. For a complete description of these panels for each Monitor III report, see *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 *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 15-9. DELAY Report Options Panel - Goal Mode

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 15-13).

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:      IEAIPSAQ on AQTS
Recommended Option Set:  IEAIPSAQ on AQTS

Action    Name     System     Description
   _      ________            ___________________________________________
   _      ________            ___________________________________________
   _      IEAIPSAQ  AQTS      Generated from IPS member 07/14/99
   _      INITIAL             RMF options from release 2.6
```

Figure 15-10. 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 installation performance specifications (IPS) member, and fills in the system ID. 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 based on IPS definitions 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 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 2.7.0  Data Index - RMFPLEX1         Line 1 of 22
Command ===>                                                    Scroll ===> HALF

Samples: 118     System: MVS2  Date: 04/27/99  Time: 10.12.00  Range: 120   Sec

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

MVS1   04/27/99 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   04/27/99 09.11.00 SYS00002 RMF.MONITOR3.DATASET1.MVS2
                09.14.00
MVS2   04/27/99 10.03.00 SYS00003 RMF.MONITOR3.DATASET2.MVS2
                10.12.00             * * *      Currently active      * * *
MVS2   04/27/99 10.03.00
                10.12.00             * * *      In-storage buffer     * * *

MVS3   04/27/99 09.11.00 SYS00002 RMF.MONITOR3.DATASET1.MVS3
                09.14.00
MVS3   04/27/99 10.03.00 SYS00003 RMF.MONITOR3.DATASET2.MVS3
                10.12.00             * * *      Currently active      * * *
MVS3   04/27/99 10.03.00
                10.12.00             * * *      In-storage buffer     * * *
TEST
                                     * * *      No response        * * *
```

Figure 15-11. 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.

```
                        RMF 2.7.0  Data Index - RMFPLEX1           Line 1 of 4
Command ===>                                                    Scroll ===> HALF

Samples: 118     System: RMFE  Date: 04/27/99  Time: 10.12.00  Range: 120   Sec

       ------Begin------          -------End-------
System --Date-- --Time--          --Date-- --Time--

MVS1   04/27/99 10.03.20          04/27/99 10.12.00
MVS2   04/27/99 09.11.00          04/27/99 10.12.00
MVS3   04/27/99 09.11.00          04/27/99 10.12.00
TEST                                    * * *   No response        * * *
```

Figure 15-12. 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(IEASYSxxx) parameter SYSNAME.

The four-character SMF system ID, defined in the SYS1.PARMLIB(SMFPRMxx) parameter SID(xxxx) 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 that 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

Figure 15-13. 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      Include DDNAMES / DSNAMES information (YES NO)
   Sort Order       ==> ASCEND   Sort data set names (ASCEND DESCEND)
```

Figure 15-14. 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 15-11 on page 15-15, and a panel giving a condensed list of systems belonging to the sysplex, as shown in Figure 15-12 on page 15-16.

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 Commands

This section describes, in alphabetical order, the Monitor III session commands. Figure 15-15 shows a summary of the commands and their parameters.

Figure 15-15 (Page 1 of 2). Summary of Monitor III Session Commands

Command	Parameters	Result	Abbreviation	Details on
BREF	DATE = TIME = RANGE = SYSTEM =	Changes date, time and range (backward) and system (sysplex environment)		page 15-20
CANCEL		Restores options to state at panel entrance (except report job options panel)		page 15-23
COMMANDS		Displays RMF help menu for commands	COM, CMD	page 15-23
CURRENT		Retrieves current data for display	CU	page 15-24
FIND		Searches for character string on report panels	F, FI	page 15-24
FREF	DATE = TIME = RANGE = SYSTEM =	Changes date, time and range (forward) and system (sysplex environment)		page 15-20
GO		Switches to GO mode processing		page 15-24
GRAPHIC	ON OFF	Switch to graphic mode Switch to tabular mode	GR	page 15-25
HARDCOPY	ON OFF REPORT SCREEN	Print all displayed reports Print no reports Print tabular copy of reports Print a copy of the screen image	HC	page 15-25
ICU		Sends report data to ICU and starts an ICU session		page 15-26
PFK		Displays list of PF keys		page 15-27
RESET		Resets options (excluding JOBNAME options) to default values shipped with RMF, or, for the WFEX report in compatibility mode, sets meaningful options automatically		page 15-28

Figure 15-15 (Page 2 of 2). Summary of Monitor III Session Commands

Command	Parameters	Result	Abbreviation	Details on
RETRIEVE		Displays last command entered on the command line		page 15-28
RFIND		Repeats the FIND command		page 15-24
TABULAR	ON OFF	Switch to tabular mode Switch to graphic mode	TAB	page 15-28
TOGGLE		Switches between tabular and graphic display	TOG	15-28

Backward and Forward Referencing (BREF/FREF)

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 15-16, rather than on BREF or FREF commands:

```
                        RMF 2.7.0  TITLE                         Line 1 of 30
Command ===>                                                   Scroll ===> HALF

Samples: nnn     System: syst  Date: mm/dd/yy  Time: hh.mm.ss  Range: 100   Sec
```

Figure 15-16. Header of Monitor III Single-System Reports

The syntax of the BREF and FREF commands is:

```
►►──┬─BREF─┬──┤ D ├──┤ T ├──┤ R ├──┤ S ├──────────────────────────────►◄
    └─FREF─┘

D:
├──┬─DATE─┬──=──mm──┬─.─┬──dd──┬─.─┬──yy──────────────────────────────┤
   └─D────┘         ├─/─┤      ├─/─┤
                    └─-─┘      └─-─┘

T:
├──┬─TIME─┬──=──hh──┬─:─┬──mm──┬─:─┬──ss──────────────────────────────┤
   └─T────┘         ├─-─┤      ├─-─┤
                    └─.─┘      └─.─┘

R:
├──┬─RANGE─┬──=──┬──┬────────┬──┬────────────────────────────────────┤
   └─R─────┘     │  └─nnnM───┘  │
                 │        ┌─S─┐ │
                 └─nnnn───┴───┴─┘

S:
├──┬─SYSTEM─┬──=──systemname──────────────────────────────────────────┤
   └─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.

Note: 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.

Notes:

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. The following diagram shows how these BREF and FREF commands extend the range backward and forward, as well as the intervals included in the display reports as a result of the commands.

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 (CANCEL)

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 CANCEL on the command line. 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 Information for RMF Commands (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.

Displaying Current Range Data (CURRENT)

To display a report with data from the current time for the length of the current range value, issue CURRENT on the command line.

“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 (FIND/RFIND)

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 (GO)

To switch from STOP mode to GO mode, enter GO on the command line of any report. 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, ensure that it has a PA1 key defined. 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 (GRAPHIC)

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:

```
          ┌─ON──┐
►►──GRAPHIC──┼─────┼──────────────────────────────────────────►◄
          └─OFF─┘
```

RMF activates GRAPHIC mode, and if your terminal supports graphics, and your installation has the Graphical Data Display Manager* (GDDM*) program product, 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 (HARDCOPY)

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

```
             ┌─ON─────┐
►►──HARDCOPY──┼────────┼──────────────────────────────────────►◄
             ├─OFF────┤
             ├─SCREEN─┤
             └─REPORT─┘
```

The parameters, which are optional, ave 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 *OS/390 ISPF User's Guide* (SC28-1239).

Printing Graphic Reports (ICU)

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

```
Command ====> 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 15-17 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 APR 1999 10:24 AM
      001 DELAY    GDF      ADMGDF    0  09 APR 1999 10:24 AM RMF
      002 WFEX     GDF      ADMGDF    0  08 APR 1999  1:43 PM RMF
      003 WFEX1    GDF      ADMGDF    0  03 APR 1999  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 15-17. 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.RMFOS270.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 (PFK)

Issue the PFK command to display the program function keys.

You can also use the ISPF KEYS command.

Figure 15-18 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.

Figure 15-18. 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 (RESET)

To ensure the RMF default settings for option panels are in effect, enter RESET on the command line or the respective panel. 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 from the current installation performance specifications (IPS) member.

Retrieving the Last Command (RETRIEVE)

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

Activating Tabular Mode (TABULAR)

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

```
               ┌─ON──┐
►►──TABULAR────┼─────┼──────────────────────────────────────────────►◄
               └─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 (TOGGLE)

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 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 *RMF Programmer's Guide* for more information on user-written reports.

Using the PA Keys

PA1

Some terminals use PA1 as their ATTN key. When the session is in GO mode, pressing PA1 changes the mode to STOP and displays the current screen.

Note: If you are running a reporter session in GO mode, the *only* way to get out of GO mode is to use the ATTN key, or the key that is defined as ATTN, such as PA1.

PA2

Pressing PA2 refreshes the currently displayed screen.

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 14-4.

Monitor III Option and Report Commands

This section lists the commands that you use to reach the Monitor III panels on which you can specify session and report options. You can enter each of the commands from the command line of any Monitor III panel.

There are two types of command:

- Report option commands, which allow you to influence the appearance of the display
- Report commands, which allow you to influence the content of the display

Monitor III Report Option Commands — How to Report

Figure 15-19 shows, in alphabetical order, the commands with which you call up panels to change the appearance of the session. The corresponding panels are described in "Changing Options" on page 15-7

Figure 15-19. Report Option Commands

Command	Displays	Abbreviation
COLOR	Color graphic options panel	CO
EXIT	Leave Monitor III dialog	X
JOBS	Job delay report menu	JS
LANGUAGE	Language option panel	LO
OPTIONS	Option selection panel	OP, OPT
OPTSET	Option set selection panel	OS
OVERVIEW	Workflow/exceptions, system information, and delays	OV
RESOURCE	Processor, device, enqueue and storage menu	RS
ROPTIONS	Report options panel for displayed report	RO
SESSION	Session options panel	SO
SUBS	Subsystem selection (HSM, JES, and XCF)	SUB
SYSPLEX	Sysplex reports, coupling facility, cache, and data index	SP
USER	User-written reports panel	SO

Monitor III Report Commands — What to Report

Figure 15-20 on page 15-31 shows, in alphabetical order, the Monitor III report commands you use to display reports. See *RMF Report Analysis* for a detailed description of each report.

Enter the commands on any command line.

The **Parameters** column indicates what parameters, if any, you can specify on the respective commands:

class

ALL (A)
ASCH (AS)
BATCH (B)
OMVS (O)

STC (S)
TSO (T)

Notes:

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.

sel In compatibility mode:

`DMNnnn` for domain
`PGnnn` for performance group
`ALL`

`n` is the number of the domain or performance group.

In goal mode:

A workload group name
A service class name
A report class name

Note: This parameter is optional.

group In compatibility mode:

`DMNnnn` for domain
`PGnnn` for performance group
`ALL`

`n` is the number of the domain or performance group.

In goal mode:

A service class name

period A service class period

cfname A coupling facility name

] **ssid** A cache subsystem identifier

] **sstype** The name of a subsystem that schedules enclaves

Example

- Display the delay for all batch address spaces in performance group 1:

  `DELAY B,PG1`

- Display GROUP report of service class TSOHIGH for service class period 1:

  `GROUP TSOHIGH,1`

Figure 15-20 (Page 1 of 3). 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

Figure 15-20 (Page 2 of 3). Report Commands

Command	Parameters	Displays	Abbreviation
CFSYS	cfname	Coupling facility system report	CS
CHANNEL		Channel path activity report	CHAN, CH
DELAY	class, sel	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	class, sel	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
DSND	dsname	Data set delays report for all or specified data sets	DSN
DSNJ	jobname	Data set delays - Job report for specified jobname	DSJ
DSNV	volser	Data set delays - Volume report for specified volume	DSV
DSINDEX		Data index information	DS, DI
ENCLAVE	sstype	Enclave activity report	ENCL
ENQ	class, sel	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 name	Enqueue delays for all or specified resources	ER
GROUP	group, period	Group response time breakdown	GP, GRP, GD, RT, GRT
HSM	class, sel	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	class, sel	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
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
PROC	class, sel	Processor delays report for all jobs or specified job groups	PD
PROCJ	jobname	Processor delays variation of job report for specified job	PJ

Figure 15-20 (Page 3 of 3). Report Commands			
Command	**Parameters**	**Displays**	**Abbreviation**
QSCJ	jobname	Operator delays variation for quiesce command of job report for specified jobname	QJ
STOR	class, sel	Storage delays report for all jobs or specified job group	SD
STORC	class,sel	Common storage report	SC
STORCR		Common storage remaining at end of job report	SCR
STORF	class, sel	Detailed information on frame counts for all jobs or specified job group	SF
STORJ	jobname	Storage delays variation of job report for specified job	SJ
STORR		Storage space and page/swap activity report for all system volumes	SR
STORS	sel	Summarized storage information by domain and performance group, or workload group and service class	SS
SYSENQ		Sysplex enqueue delays report	ES
SYSINFO	sel	System information, total and by user groups	SY, SYS, SI
SYSRTD	service class, period	Response time distribution report	RTD
SYSSUM	sel	Sysplex summary	SUM
SYSWKM	service class, period	Work manager delays report for CICS and IMS subsystems	WKM
WFEX		Workflow/exceptions screen	WE, WF
XCF	class, sel	Cross-system coupling facility delays report	XD
XCFJ	jobname	XCF delays variation of the job report for specified jobname.	XJ

Figure 15-21 contains commands for the examples of user-written reports that were delivered with RMF.

Figure 15-21. User-Written Report Commands			
Command	**Parameters**	**Displays**	**Abbreviation**
DEVN		Device activity report	DA
DEVT		Device trend report	DT
DSD		Detailed storage delays report	
MSI		Migration SYSINFO report including execution velocity	
RG		Resource group data report	
SYSTREND		System trend report	ST

Chapter 16. Snapshot Reporting with Monitor II

About the Monitor II Dialog

This chapter guides you in using the Monitor II sessions, and describes:

- The ISPF session
- The TSO/E and local 3270 sessions
- The background session
- Session commands and report commands

Monitor II Sessions

There are 4 types of Monitor II report sessions:

- An ISPF display session

  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.

- A TSO/E display session

  Start this session with the TSO/E command:

  `RMFMON`

- A local 3270 display session

  The operator starts this session, for example for address 624, with the command:

  `MODIFY RMF,START 642`

- A background session

  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`

You can obtain a **printout** of a Monitor II session report:

- During a display session
- During or at the end of a background session

You can get the same reports in all sessions. There is just a small difference in the syntax used to request them:

- Display Session

  The reports are called by *commands* that conform to TSO/E syntax rules:

  Example: `ASD T,A,3`

- Background Session

  The reports are called via *options* that conform to option syntax rules:

  Example: `ASD(T,A,3)`

The Monitor II ISPF Session

Enablement

RMF is an optional feature of OS/390. 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.” As you can see in Figure 16-1 , 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                OS/390 2.7.0 RMF
Selection ===>

Enter selection number or command on selection line.



  1 Address spaces      Address space reports
  2 I/O Subsystem       I/O Queuing, Device, Channel, and HFS reports
  3 Resource            Enqueue, Storage, SRM, and other resource reports

  L Library Lists       Program library information
  U User                User-written reports (add your own ...)

                              T TUTORIAL    X EXIT



          5647-A01 (C) Copyright IBM Corp. 1994,1999. All Rights Reserved
                  Licensed Materials - Property of IBM
```

]

Figure 16-1. Monitor II Primary Menu

The selection **U** displays the User Reports menu. This option is only meaningful if you have written some reports of your own and included them on the User Reports menu. For information on how to do this, see the *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 14-4. When selected, each category of Monitor II report presents a more detailed selection panel showing the individual reports. The categories are:

- Address space reports
- I/O queuing, device and channel reports

- Enqueue, storage and SRM reports

The selection LLI does not lead to the generation of activity reports. It allows you to specify the program libraries that are to be reported on.

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 16-2. 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 16-3. 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/Swap data set activity
  4 SPAG              System paging activity
  5 SRCS              Central storage / processor / SRM
  6 DDMN              Domain activity (compatibility mode only)
  7 TRX               Transaction activity (compatibility mode only)

  Other Resources

  8 SDS               RMF Sysplex Data Server activity
  9 ILOCK             IRLM Long Lock detection
```

Figure 16-4. 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 Information

This panel appears in response to the choice **L** in the Primary menu. The available types of library list are presented for selection.

```
                  RMF Monitor II Library List Selection Menu
Selection ===>

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)
```

Figure 16-5. Monitor II Library List 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.

Entering the command, rather than making a selection, allows you alter the scope of the resulting library 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 16-35.

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 *RMF Programmer's Guide*:

```
                          RMF Monitor II User Reports
Selection ===>

Enter selection number or command on selection line.

  1 USER              User report 1
```

Figure 16-6. 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 Monitor II TSO/E and Local 3270 Sessions

Enablement

RMF is an optional feature of OS/390. 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, this is the panel you see first.

You get the same panel in a local 3270 display session.

```
                         RMF DISPLAY MENU
NAME      PFK#  DESCRIPTION
                                                          IPS = IEAIPSXI
ARD        1    ADDRESS SPACE RESOURCE DATA
ASD        2    ADDRESS SPACE STATE DATA
ASRM       3    ADDRESS SPACE SRM DATA
CHANNEL    4    CHANNEL PATH DISPLAY
DDMN       5    SYSTEM DOMAIN DISPLAY
DEV        6    SYSTEM DEVICE DATA
PGSP       7    SYSTEM PAGING SPACE DATA
SENQ       8    SYSTEM ENQUEUE CONTENTION
SENQR      9    SYSTEM ENQUEUE RESERVE
SPAG       10   PAGING DATA
SRCS       11   CENTRAL STORAGE / CPU / SRM DATA
TRX        12   TRANSACTION ACTIVITY DATA
ARDJ            RESOURCE DATA FOR SPECIFIC JOBNAME
ASDJ            STATE DATA FOR SPECIFIC JOBNAME
ASRMJ           SRM DATA FOR SPECIFIC JOBNAME
DEVV            SYSTEM DEVICE DATA FOR A SPECIFIC VOL/NUMBER
IOQUEUE         I/O QUEUING ACTIVITY DATA
SDS             RMF SYSPLEX DATA SERVER
ILOCK           IRLM LONG LOCK DETECTION
USER            USER PICTURE
```

Figure 16-7. Monitor II Display Menu - Compatibility Mode

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 Monitor II 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, please refer to “Starting a Specific Monitor” on page 4-2.

Structure of Monitor II Reports

This chapter provides some general information about Monitor II reports.

- For display session reports:

  You can have different report headers where you can enter commands:

  – ISPF session

  – TSO/E and local 3270 session

- For all sessions:

  You can have different formats:

  – Table reports

  – Row reports

  You can have different report modes:

  – Total mode

  – Delta mode

Furthermore, the different ways to display and modify report options depending on the session type are explained in “Display and Set Options” on page 16-13.

Display Session Report Fields

If You are Using an ISPF Session

Each report consists of

- A header line identifying the report.
- A line for commands and scroll amount field.
- A status line for MIG, CPU, UIC, PFR and System. For a description of these fields, please refer to Figure 16-10 on page 16-9.

  This line also contains the current setting of the DELTA mode.

- A variable number of data of data lines.

```
                         RMF - Name of Report                     Line 1 of nn
COMMAND ===>                                                  Scroll ===> HALF

                      MIG= nnnK CPU=nnn/nnn UIC=nnn PFR=   n   System=sid Total
```

Figure 16-8. Header of a Monitor II ISPF Session Report

Each ISPF report has the standard layout as you know it from Monitor III reports and other ISPF applications.

If You are Using the TSO/E or Local 3270 Session

Each report consists of:

- A title line
- Two lines of heading information
- A variable number of data of data lines

```
                    ┌    ►Input Area                                          ┌     ►Report Name Area
                    │                      ┌    ►Status/Message Area          │      ┌    ►Mode Area
                    ┴                      ┴                                  ┴      ┴
┌                                                                                      ┐
 F              MIG=786K       CPU= 41/ 50 UIC=255 PFR=   0 SYS1 CHANNEL    TH
  19:13:47   CHANNEL PATH    UTILIZATION(%)        CHANNEL PATH    UTILIZATION(%)
             ID  TYPE SHR   PARTITION  TOTAL       ID  TYPE SHR   PARTITION  TOTAL     ├►Header Area
├                                                                                      ┤
             00   BY           0.00     0.00       01   BL           0.63     0.63
             03   BL           0.00     0.00       05   BL           0.57     0.57     ├►Output Area
             07   BL           0.00     0.00       08   BL           3.73     3.73
└                                                                                      ┘
```

Figure 16-9. Header of a Monitor II TSO Session Report

Figure 16-10 (Page 1 of 2). 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	Normally, this area contains: **MIG** Migration age (in seconds). **CPU** Current average processor utilization. This information depends on the activity of Monitor I. If Monitor I CPU gathering is **active**: • In **LPAR mode,** the header line shows two views separated by a slash (/): – The MVS view of the CPU utilization – The LPAR view of the CPU utilization • In **Basic mode** (non-LPAR mode), the header line shows the MVS view. If Monitor I CPU gathering is **not active**: • In **LPAR mode,** the header line shows: – The SRM view (in compatibility mode, it includes queue length and can have a value of up to 128) of the CPU utilization. – '***' due to missing CPU measurement data for the LPAR view • In **Basic mode** (non-LPAR mode), the header line shows the SRM view. **UIC** The highest unreferenced interval count. **PFR** 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.

Figure 16-10 (Page 2 of 2). Monitor II Display-Session Areas	
Area	**Function**
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 Figure 16-10 on page 16-9.

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 no 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. They can be displayed in:

- **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 I 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

Display-Session Commands

Figure 16-11 (Page 1 of 2). Monitor II Display-Session Commands

Task	ISPF Command	Parm	TSO/3270 Command	Parm	Result
"Display the Menu" on page 16-13	RETURN (PF4)		M		Returns to the Primary menu.
"Display and Set Options" on page 16-13	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 16-15	RESET				**ISPF** On Report Option panels, resets all optional values to the defaults specified in the menu ERBFMENU
"Leave Options Unchanged" on page 16-15	CANCEL				**ISPF** On Report Option panels, ends the option dialog without making any changes
"Display Commands" on page 16-15	COMMANDS				**ISPF** Shows you all the available commands
"Scroll Through Report Frames" on page 16-15	PF8/PF7		F		**ISPF** Forward/backward scrolling **TSO/E** Forward scrolling
"Recall the Previous Data" on page 16-16	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 16-16	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 16-17	FIND	string			**ISPF** Searches for a text string in a report
"Repeat Find" on page 16-17	RFIND (PF5)				**ISPF** Repeats a previously entered FIND command
"Set Delta Mode" on page 16-17	D	ON OFF	D	ON OFF	Sets the DELTA mode Sets the TOTAL mode
"Create a Hardcopy Report" on page 16-18	H	ON OFF	H	ON OFF	Prints all displayed reports Prints no reports

Figure 16-11 (Page 2 of 2). Monitor II Display-Session Commands

Task	ISPF Command	Parm	TSO/3270 Command	Parm	Result
"Refresh a Report Automatically" on page 16-19	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.
"Print a Report Page" on page 16-21	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 16-21	SYS[TEM]	smf_id	SYS[TEM]	smf_id	Identifies system to be monitored
"Assign PF Keys" on page 16-20	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 16-22	=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

Figure 16-12 (Page 1 of 2). Monitor II Background Session Options

Task	Background Session Option	Result
"Refresh a Report Automatically" on page 16-19	`{     (30S)}` `SINTV {(value[S])}`	Specifies number of seconds in each measurement interval.
"Stop the Session" on page 16-22	`          {M}` `STOP(value {H})/NOSTOP`	Desired duration of the Monitor II session, in minutes (M), or hours (H).
"Set Delta Mode" on page 16-17	`DELTA/NODELTA`	Specifies whether RMF should report total values or values that reflect changes since the previous measurement.
"Write SMF Records" on page 16-21	`RECORD/NORECORD`	Specifies whether measured data is to be written to the SMF data set.
"Create a Hardcopy Report" on page 16-18	`REPORT{(REALTIME)}/(NOREPORT)` `      {(DEFER)    }`	Specifies production of printed interval reports of measured data.
"Create a Hardcopy Report" on page 16-18	`SYSOUT(class)`	SYSOUT class to which the formatted printed reports are directed.
"Define Session Options" on page 16-15	`MEMBER (list)`	Parmlib member, or list of members, containing Monitor II background session options.

Figure 16-12 (Page 2 of 2). Monitor II Background Session Options		
Task	**Background Session Option**	**Result**
"Display and Set Options" on page 16-13	`{OPTIONS}/{NOOPTIONS}` `{OPTN } {NOOPTN }`	Print an options list at the operator console at the start of the session.

Monitor II Session Tasks

Display the Menu

For an ISPF Session

Press PF4 or enter the RETURN command to return to the Primary Menu (see Figure 16-1 on page 16-3).

For TSO/E and Local 3270 Sessions

To display the menu of available reports, issue the menu command:

```
M
```

Figure 16-7 on page 16-7 shows the menu panel. The menu lists each report name, its PF key assignment, and a description. The IPS Parmlib member name appears in the upper right corner of the menu.

To display the menu for a local 3270 display session, you can also press the attention (PA1) key. To re-display the previous panel, press PA2.

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.

Here is 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.
 Domain        ===> ALL     Specify a valid domain number (1..128) or ALL
                            for all domains.
```

Figure 16-13. 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 TSO/E and Local 3270 Sessions

You get an overview about the default options for all commands by entering the command:

`MM`

The standard definition of the default options is shown in this figure:

```
                         RMF DISPLAY MENU
      NAME      PFK#  DEFAULT OPERANDS FOR GATHERER AND REPORTER
                                                           IPS = IEAIPST0
      ARD        1    A,A,    --- A,I,
      ASD        2    A,A,    --- A,I,
      ASRM       3    A,A,    --- A,I,
      CHANNEL    4
      DDMN       5    -----NOT APPLICABLE IN GOAL MODE-----
      DEV        6    DASD
      PGSP       7    PAGE
      SENQ       8    D --- S
      SENQR      9    ALLVSER
      SPAG       10
      SRCS       11
      TRX        12   -----NOT APPLICABLE IN GOAL MODE-----
      ARDJ
      ASDJ
      ASRMJ
      DEVV
      SDS
      IOQUEUE         DASD
      USER
```

Figure 16-14. 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.

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 by option:

```
MEMBER (list)
```

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 TSO/E and Local 3270 Sessions

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, TSO/E and Local 3270 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:

ISPF session syntax

```
        ┌─A or D─(1)─┐
►►──SORT─┼────────────┼──────────────────────────────────────►◄
        ├─A──────────┤
        └─D──────────┘
```

Note:

1 Numerical columns are sorted in descending order, and columns with character values are sorted in ascending order.

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 have to 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, TSO/E and Local 3270 Sessions

To set the delta mode for the session, use the delta command. The syntax is:

```
     ┌─ON──┐
►►──D─┤     ├──────────────────────────────────────────────►◄
     └─OFF─┘
```

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/3270 session).

For a Background Session

To set the delta mode for the session, use the delta command:

```
    ┌─NODELTA─┐
►►──┼─────────┼──────────────────────────────────────────────────────►◄
    └─DELTA───┘
```

Create a Hardcopy Report

For ISPF, TSO/E and Local 3270 Sessions

To create a report from the current session, you set the hardcopy mode on, using the hardcopy command:

```
        ┌─ON──┐
►►──H───┼─────┼──────────────────────────────────────────────────────►◄
        └─OFF─┘
```

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

ISPF and TSO/E Session

- Allocate the output data set before you start the Monitor II session:

  ```
  ALLOC F(RMFDMTSO) 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(RMFDMTSO) DS(data.set.name) MOD
  ```

Local 3270 Session

- The output data set has to be defined in the start-up JCL:

  ```
  //RMFDMxxx DD DSNAME=data.set.name,DISP=disp
  ```

  where **xxx** is the terminal address of the local 3270.

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/3270 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:

```
      ┌─REPORT(DEFER)──────────────────┐
►►────┼────────────────────────────────┼──────────────────────────────────────────►◄
      └──┬─NOREPORT────────────────────┤
         └─REPORT(─┬─REALTIME─┬─)──────┘
                   └─DEFER────┘
```

whether you want to create a report.

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

```
//RMFxxnnn 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 class A. You can modify the output class using the option

```
SYSOUT(class)
```

Refresh a Report Automatically

For an ISPF Session

To refresh a report automatically, enter on any report panel the command:

```
          ┌─4──────┐
►►──GO────┼────────┼──────────────────────────────────────────────────────►◄
          └──n─────┘
```

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 TSO/E and Local 3270 Sessions

To update reports automatically, use the timed update command. The syntax is:

```
        ┌─10──,──4──────┐
►►──T───┼───────────────┼─────────────────────────────────────────────────►◄
        └──m──,──n──────┘
```

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.

Note: If you are running a local 3270 display session, choose your options for the timed update command carefully; you cannot end the update by pressing the attention (PA1) key. The update ends when RMF finishes executing the timed update command.

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:

```
     ┌─SINTV(30S)────────┐
►►───┼───────────────────┼──────────────────────────────────────────►◄
     └─SINTV(value─┬───┬─)─┘
                   └─S─┘
```

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 TSO/E and Local 3270 Sessions

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 16-7 on page 16-7. 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.

For a local 3270 session, pressing the attention key causes RMF to display the Display menu. Pressing PA2 causes RMF to re-display the previous screen and wait for input.

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 TSO/E and Local 3270 Sessions

To print the data currently displayed on the screen, use the print command:

```
P
```

See “Create a Hardcopy Report” on page 16-18 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 16-18.

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 overtype 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 TSO/E and Local 3270 Sessions

To end the session, enter:

```
►►──┬─Z────┬──────────────────────────────────────────────────►◄
    ├─END──┤
    ├─QUIT─┤
    ├─QQ───┤
    ├─STOP─┤
    └─X────┘
```

For a Background Session

You can specify the duration of a background session with the option:

```
    ┌─STOP(10M)──────────────────┐
►►──┼────────────────────────────┼────────────────────────────►◄
    │                ┌─M─┐       │
    ├───STOP(value───┼───┼──)────┤
    │                └─H─┘       │
    └─NOSTOP─────────────────────┘
```

Monitor II Report Commands - Overview

Figure 16-15 gives an overview of the available report commands.

Figure 16-15 (Page 1 of 2). Monitor II Report Commands

Background Session Syntax	Display Session Syntax	Report
`ARD [(class,status,domain)]/NOARD`	`ARD [class,status,domain]`	Address space resource data reporting. See page 16-24.
`ARDJ (jobname)/NOARDJ`	`ARDJ jobname`	Address space resource data reporting for a particular job. See page 16-24.
`ASD[(class,status,domain)]/NOASD`	`ASD [class,status,domain ]`	Address space state data reporting. See page 16-26.
`ASDJ(jobname)/NOASDJ`	`ASDJ jobname`	Address space state data reporting for a particular job. See page 16-27.
`ASRM[(class,status,domain)]/NOASRM`	`ASRM[class,status,domain]`	Address space SRM data reporting. See page 16-27.
`ASRMJ(jobname)/NOASRMJ`	`ASRMJ jobname`	Address space SRM data reporting for a particular job. See page 16-27.
`CHANNEL/NOCHANNEL`	`CHANNEL`	Channel path activity data reporting. See page 16-28.

Figure 16-15 (Page 2 of 2). Monitor II Report Commands

Background Session Syntax	Display Session Syntax	Report
DDMN/NODDMN	DDMN	Domain activity reporting (compatibility mode only). See page 16-28.
DEV [(type)]/NODEV	DEV [type]	Table reporting on I/O device activity. See page 16-29.
DEVV{(VOLSER(xxxxxx))}/NODEVV {(NUMBER(yyyy))}	DEVV {VOLSER(xxxxxx)} {NUMBER(yyyy) }	Row reporting on a specific direct access device. See page 16-31.
--	HFS [hfsname]	Table reporting on UNIX hierarchical file system statistics. See page 16-33.
--	ILOCK [ALL]	IRLM Long Lock detection. See page 16-33.
IOQUEUE[(type)]/NOIOQUEUE	IOQUEUE [type]	I/O request queuing reporting. See page 16-34.
--	LLI {LNK}{,A} {LPA} {APF}	Program library information listing. See page 16-35.
PGSP {(PAGE)}/NOPGSP {(SWAP)}	PGSP{PAGE} {SWAP}	Page/swap data set activity reporting. See page 16-38.
--	SDS	RMF Sysplex Data Server activity reporting. See page 16-40.
{(S)} {(D)} SENQ {(A,sysname)}/NOSENQ {(E,sysname)} {(majorname[,minorname])}	{S } {D } SENQ {A,sysname } {E,sysname } {majorname[,minorname]}	Enqueue contention activity reporting. See page 16-41.
SENQR {(ALLVSER)}/NOSENQR {(volser)}	SENQR {ALLVSER} {volser }	Reserve activity reporting. See page 16-43.
SPAG/NOSPAG	SPAG	System paging activity reporting. See page 16-45.
SRCS/NOSRCS	SRCS	Central storage/processor/SRM activity reporting. See page 16-45.
TRX {(ALLPGN)}/NOTRX {(sname,nnnn[,nnnn,nnnn:nnnn])}	TRX {ALLPGN } {sname,nnnn[,nnnn,nnnn:nnnn]}	Transaction activity reporting (compatibility mode only). See page 16-47.
USER/NOUSER	USER	Specifies whether or not a user-specified activity is to be measured. See page 16-49.

Details of Session and Report Commands

This section describes the Monitor II session and report commands in alphabetical order. Program defaults are underscored where appropriate.

The same report commands are available for Monitor II display and background sessions, but what you can specify varies from one type of session to the other. The background and display forms of each command are shown on the following pages.

Display Session

Specify commands either in the command or selection line of the ISPF panel or in the input area of a display 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 4-2 for the START session command syntax.
- The Monitor II background session Parmlib member ERBRMF01. See the explanation of the background session MEMBER option (“MEMBER” on page 16-36) The contents of ERBRMF01 is described in Storing Gatherer Options.

RMF uses a program default for any option you do not specify.

Chapter 5, How RMF Processes Session Options 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 4-5.

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

Background session syntax

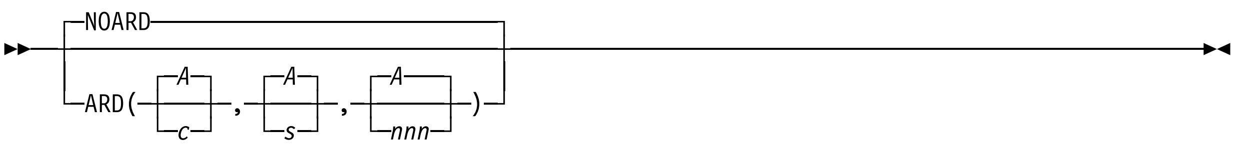

Display session syntax

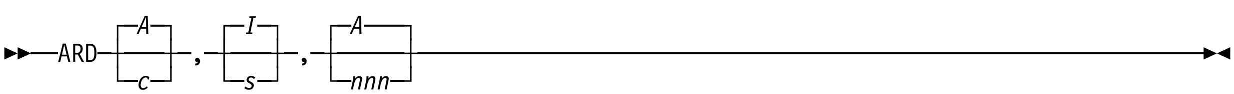

Specifies address space resource data reporting, where **class, status**, and **domain** specify the following selection criteria for the address spaces to be included:

class

A	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

I active address spaces; that is, those address spaces that are currently executing, non-swappable, or swapped out and eligible for swap-in

domain

This option is available only in compatibility mode. There is no equivalent in goal mode.

A all address spaces

nnn address spaces in the domain identified by **nnn**

Class, status, and **domain** are optional; if you omit them, the ARD report uses the default values, and includes all currently active address spaces in the system.

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 display 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 *RMF Programmer's Guide*.

Example

For a display session:

- To report on all address spaces that are currently active in domain 3, enter

```
ARD ,,3
```

- To report on all TSO address spaces that are currently active in any domain, enter

```
ARD T
```

- To report on all batch, started task, and mount task address spaces that are currently active in domain 3, enter

```
ARD B,,3
```

Example

For a background session:

- To select for measurement all address spaces that are currently active in domain 3, specify:

  `ARD(,,3)`

- To select for measurement all TSO address spaces that are currently active in any domain, specify:

  `ARD(T)`

- To select for measurement all batch, started task, and mount task address spaces that are currently active in domain 3, specify:

  `ARD(B,,3)`

ARDJ

Background session syntax

```
      ┌─NOARDJ─────────┐
►►────┤                ├──────────────────────────────────────────►◄
      └─ARDJ(jobname)──┘
```

Display session syntax

```
►►──ARDJ──jobname─────────────────────────────────────────────────►◄
```

Specifies address space resource data reporting for a particular job. Specify the job you want to measure under **jobname**.

ASD

Background session syntax

```
      ┌─ASD(A,A,A)──────────────────────┐
►►────┼─────────────────────────────────┼─────────────────────────►◄
      ├─NOASD───────────────────────────┤
      │        ┌─A─┐    ┌─A─┐    ┌─A───┐   │
      └─ASD(─┴─c─┴─,─┴─s─┴─,─┴─nnn─┴─)─┘
```

Display session syntax

```
         ┌─A─┐     ┌─I─┐     ┌─A───┐
►►──ASD──┴─c─┴──,──┴─s─┴──,──┴─nnn─┴──────────────────────────────►◄
```

Specifies address space state data reporting; **class, status,** and **domain** specify the address space selection criteria. For a description of **class, status,** and **domain**, see the ARD option.

ASDJ

Background session syntax

```
     ┌─NOASDJ──────────┐
►►───┤                 ├──────────────────────────────────────────────►◄
     └─ASDJ(jobname)───┘
```

Display session syntax

```
►►──ASDJ──jobname─────────────────────────────────────────────────────►◄
```

Specifies address space state data reporting for a particular job. Specify the job you want to measure in **jobname**.

ASRM

Background session syntax

```
     ┌─NOASRM──────────────────────────────┐
►►───┤                                     ├─────────────────────────►◄
     │        ┌─A─┐     ┌─A─┐     ┌─A───┐  │
     └─ASRM(──┴─c─┴──,──┴─s─┴──,──┴─nnn─┴─)┘
```

Display session syntax

```
         ┌─A─┐     ┌─I─┐     ┌─A───┐
►►──ASRM─┴─c─┴──,──┴─s─┴──,──┴─nnn─┴──────────────────────────────────►◄
```

Specifies address space SRM data reporting. **class, status,** and **domain** specify the selection criteria for the address spaces to be included. For a description of **class, status,** and **domain**, see the ARD option.

ASRMJ

Background session syntax

```
     ┌─NOASRMJ──────────┐
►►───┤                  ├─────────────────────────────────────────────►◄
     └─ASRMJ(jobname)───┘
```

Display session syntax

```
►►──ASRMJ──jobname────────────────────────────────────────────────────►◄
```

Specifies address space SRM data reporting for a particular job. Specify the job you want to measure in **jobname**.

CANCEL

ISPF session syntax

```
►►──CANCEL──────────────────────────────────────────────────────►◄
```

Ends the Report Options panel without making any changes.

CHANNEL

Background session syntax

```
     ┌─NOCHANNEL─┐
►►───┴─CHANNEL───┴──────────────────────────────────────────────►◄
```

Display session syntax

```
►►──CHANNEL─────────────────────────────────────────────────────►◄
```

Specifies channel path activity data reporting.

COMMANDS

ISPF session syntax

```
►►──COMMANDS────────────────────────────────────────────────────►◄
```

Displays all available commands.

DDMN

Background session syntax

```
►►──DDMN────────────────────────────────────────────────────────►◄
```

Display session syntax

```
     ┌─NODDMN─┐
►►───┴─DDMN───┴─────────────────────────────────────────────────►◄
```

Specifies domain activity reporting. This option is accepted only in compatibility mode.

DELTA

Background session syntax

```
     ┌─NODELTA─┐
►►───┤         ├──────────────────────────────────────────────────►◄
     └─DELTA───┘
```

Display session syntax

```
         ┌─ON──┐
►►───D───┤     ├──────────────────────────────────────────────────►◄
         └─OFF─┘
```

Specifies whether RMF should report either total values (NODELTA) or, after the first measurement, values that reflect changes since the previous measurement (DELTA).

Note: DELTA/NODELTA is a *reporting* option and has no impact on SMF recording in a background session.

For a complete description of a Monitor II display session DELTA option, see “Set Delta Mode” on page 16-17.

DEV

Background session syntax

```
     ┌─NODEV──────────────┐
►►───┤                    ├──────────────────────────────────────►◄
     │        ┌─(DASD)─┐  │
     └─DEV────┤        ├──┘
              └─(type)─┘
```

Display session syntax

```
          ┌─DASD─┐
►►───DEV──┤      ├─────────────────────────────────────────────────►◄
          └─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.

Type can be one of the following:

- A device class:

  DASD Direct access storage devices

  TAPE Magnetic tape devices

  COMM Communication equipment

CHRDR Character reader devices

UNITR Unit record devices

GRAPH Graphic devices

- One or more volume serial numbers:

```
         ({aaaaaa              })
{VOLSER} ({aaaaaa,bbbbbb:zzzzzz})
{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 aaaaa, a list of volumes, in the format aaaaa,bbbbb,...., or a range of volumes, in the format bbbbb:zzzzz where aaaaa is the first and bbbbb 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 of 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:

- To request a Device Activity report for all magnetic tape devices, specify:

  `DEV TAPE`

- To request a Device Activity report for volumes P50002, P50003, P50004, and P50007, specify:

  `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)`

- To request all storage groups, specify:

  `DEV SG`

Example

For a background session:

- To request a Device Activity report for all tape devices, specify:

  `DEV(TAPE)`

- To request a Device Activity report for tape device numbers 1580 through 1584, specify:

  `DEV(NUM(1580:1584))`

- To request a Device Activity report for storage group MANF01, and storage groups MANF05 through MANF15, specify:

  `DEV(SG(MANF01,MANF05:MANF15))`

DEVV

Background session syntax

```
    ┌─NODEVV─────────────────────┐
►►──┤                            ├──────────────────────────────────►◄
    └─DEVV(──┬─VOLSER(xxxxxx)─┬──)┘
             └─NUMBER(yyyy)───┘
```

Display session syntax

```
►►──DEVV──┬─VOLSER(xxxxxx)─┬────────────────────────────────────────►◄
          └─NUMBER(yyyy)───┘
```

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:

- To request a Device Activity report for tape device number 1580, specify:

```
DEVV N(1580)
```

- To request a Device Activity report for direct access device DASD01, specify:

```
DEVV V(DASD01)
```

Example

For a background session:

- To request a Device Activity report for tape device number 1580, specify:

```
DEVV(N(1580))
```

- To request a Device Activity report for tape TAPE01, specify:

```
DEVV(VOLSER(TAPE01))
```

FIND

ISPF session syntax

```
►►──FIND──xxxxxxx──────────────────────────────────────────────►◄
```

Searches for a text string in the report. Strings that include spaces must be enclosed in quotes.

GO

ISPF session syntax

```
        ┌─4───┐
►►──GO──┴─nnn─┴────────────────────────────────────────────────►◄
```

HFS

Display session syntax

```
►►──HFS──┬─────────┬──────────────────────────────────────────►◄
         └─hfsname─┘
```

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

Display session syntax

```
►►──ILOCK──┬─────┬────────────────────────────────────────────►◄
           └─ALL─┘
```

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, please refer to “Controlling the SMF Buffer” on page 3-4.

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. Please, refer to “Ensure Access to SMF Data for the Sysplex Data Services” on page 2-4 for details.

IOQUEUE

Background session syntax

```
   ┌─NOIOQUEUE───────────┐
►►─┤                     ├──────────────────────────────────────────►◄
   │         ┌─(DASD)─┐  │
   └─IOQUEUE─┤        ├──┘
             └─(type)─┘
```

Display session syntax

```
              ┌─DASD─┐
►►──IOQUEUE───┤      ├──────────────────────────────────────────────►◄
              └─type─┘
```

Specifies I/O request queuing data reporting, If you specify IOQUEUE, a Monitor I session must be measuring any I/O queuing activity. **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 an MVS 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:

- To request an I/O Queuing Activity report for LCUs representing all magnetic tape devices, specify:

  `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)`

Example

For a background session:

- To request an I/O Queuing Activity report for LCU numbers representing only the magnetic tape devices, specify:

  `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))`

LLI

Display session syntax

```
       ┌─LNK─┐
►►──LLI─┤     ├──┬────┬──────────────────────────────────────────────►◄
       ├─LPA─┤  └─,A─┘
       └─APF─┘
```

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 LNKLSTxx members in SYS1.PARMLIB.

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.

MEMBER

```
   ┌─MEMBER(01)──────────┐
►►─┤                     ├──────────────────────────────────────────────►◄
   │          ┌─,───┐    │
   └─MEMBER(──▼─xx(1)┴──)─┘
```

Note:
[1] You can specify up to five members

Specifies the Parmlib members that contain Monitor II background options for the session, where **(list)** 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 IEFRDER DD statement in the RMF cataloged procedure. See "Setting Up RMF Control Session including Monitor I and Monitor II" on page 2-5.

For more information on Parmlib members, including the contents of the Monitor II member ERBRMF01, see Storing Gatherer Options.

OPTIONS

```
   ┌─OPTIONS───┐
►►─┤           ├────────────────────────────────────────────────────────►◄
   ├─OPTNS─────┤
   ├─NOOPTNS───┤
   └─NOOPTIONS─┘
```

Specifies whether or not an options list for the session is to be printed at the operator console at the start of the background 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 16-16 shows each possible option followed by its input source.

Figure 16-16. Monitor I 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.
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 : NODDMN -- DEFAULT
ERB103I LS : NOSENQR -- DEFAULT
ERB103I LS : NOSENQ -- DEFAULT
ERB103I LS : NOTRX -- DEFAULT
ERB103I LS : NOASRMJ -- DEFAULT
ERB103I LS : NOASRM -- DEFAULT
ERB103I LS : NOARD -- DEFAULT
ERB103I LS : NOSRCS -- DEFAULT
ERB103I LS : NOSPAG -- DEFAULT
ERB103I LS : NOARDJ -- DEFAULT
ERB103I LS : NOASDJ -- DEFAULT
ERB103I LS : NOIOQUEUE -- DEFAULT
ERB103I LS : SYSOUT(A) -- MEMBER
ERB103I LS : OPTIONS -- MEMBER
ERB103I LS : REPORT(DEFER) -- MEMBER
ERB103I LS : RECORD -- MEMBER
ERB103I LS : STOP(30M) -- MEMBER
ERB103I LS : SINTV(30S) -- MEMBER
ERB103I LS : NODELTA -- MEMBER
ERB103I LS : NOUSER -- MEMBER
ERB103I LS : ASD -- MEMBER
```

PGSP

Background session syntax

```
   ┌─NOPGSP──────────────┐
►►─┤                     ├──────────────────────────────────────►◄
   │        ┌─(PAGE)─┐   │
   └─PGSP───┤        ├───┘
            └─(SWAP)─┘
```

Display session session syntax

```
            ┌─PAGE─┐
►►──PGSP────┤      ├────────────────────────────────────────────►◄
            └─SWAP─┘
```

Requests that data on page or swap data set activity be collected, where:

PAGE Indicates page data set activity

SWAP Indicates swap data set activity

A Monitor I session monitoring page/swap activity must be active when you specify the PGSP option.

PRINT

ISPF session syntax

```
►►──PRINT───────────────────────────────────────────────────────►◄
```

Writes the currently displayed screen to the ISPF list data set.

RECORD

```
   ┌─RECORD───┐
►►─┤          ├─────────────────────────────────────────────────►◄
   └─NORECORD─┘
```

Specifies whether or not measured data is to be written to the SMF data set. Make sure you follow the complementary SMF enabling procedures. These procedures are described in *OS/390 MVS System Management Facilities (SMF)*.

REPORT

```
     ┌─REPORT(DEFER)──────────────────┐
►►───┼────────────────────────────────┼──────────────────────────────────────────────────►◄
     └──┬─NOREPORT────────────────────┤
        └─REPORT(──┬─REALTIME─┬─)─┘
                   └─DEFER────┘
```

Specifies whether or not printed interval reports of the measured data are to be produced. When reports are to be produced (REPORT specified), the REALTIME or DEFER option indicates when the reports are to be printed.

When you omit the option, the default is REPORT(DEFER). If you 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 NOSPAG, any accumulated paging reports will be printed. If you change REPORT(REALTIME) to REPORT(DEFER) or NOREPORT, all accumulated reports will be printed.

DEFER indicates that the reports are to be printed after you stop RMF.

RESET

ISPF session syntax

```
►►──RESET──────────────────────────────────────────────────────────────────►◄
```

Resets all options on an ISPF Report Options panel to the RMF default values.

RETURN

ISPF session syntax

```
►►──RETURN──────────────────────────────────────────────────────────►◄
```

Returns to the Monitor II Primary menu.

RFIND

ISPF session syntax

```
►►──RFIND───────────────────────────────────────────────────────────►◄
```

Repeats a previously entered FIND command

RO

ISPF session syntax

```
►►──RO──────────────────────────────────────────────────────────────►◄
```

Displays the Report Options panel for the current report.

SDS

Display session syntax

```
►►──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

Background session syntax

```
     ┌─NOSENQ──────────────────────────────────────┐
►►───┤                                             ├──────────────────────────────────►◄
     └─SENQ──┬─(S)─────────────────────────────┬───┘
             ├─────────────────────────────────┤
             ├─(D)─────────────────────────────┤
             ├─(A,sysname)─────────────────────┤
             ├─(E,sysname)─────────────────────┤
             └─(majorname──┬───────────────┬─)─┘
                           └─,minorname────┘
```

Display session syntax

```
            ┌─S──────────────────────────────┐
►►──SENQ────┼────────────────────────────────┼──────────────────────────────────►◄
            ├─D──────────────────────────────┤
            ├─A,sysname──────────────────────┤
            ├─E,sysname──────────────────────┤
            └─majorname──┬───────────────┬───┘
                         └─,minorname────┘
```

Specifies whether or not enqueue contention activity is to be measured. 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:

- To obtain a summary report for all resources that have contention, enter:

  `SENQ S or SENQ`

- To obtain a detail report for all resources that have contention, enter:

  `SENQ D`

- To obtain a report for all resources that system C303 hold in a global resource serialization complex, enter:

  `SENQ A,C303`

- To obtain a detail report for all resources grouped under the major name of SYSCTLG, enter:

  `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`

Examples

For a background session:

- To obtain summary data for all resources that have contention, enter:

  `SENQ(S) or SENQ`

- To obtain detail data for all resources that have contention, enter:

  `SENQ(D)`

- To obtain a report for all resources that system C303 holds in a global resource serialization complex, enter:

  `SENQ (A,C303)`

- To obtain detail data for all resources grouped under the major name of SYSCTLG, enter:

  `SENQ(SYSCTLG)`

- To obtain detail data for all resources with the major name of SYSI and the minor name of OPENUADS, enter:

  `SENQ(SYSI,OPENUADS)`

SENQR

Background session syntax

```
       ┌─NOSENQR─────────────────┐
►►─────┤                         ├──────────────────────────────────────────►◄
       │         ┌─(ALLVSER)─┐   │
       └─SENQR───┤           ├───┘
                 └─(volser)──┘
```

Display session syntax

```
               ┌─ALLVSER─┐
►►──SENQR──────┤         ├──────────────────────────────────────────────────►◄
               └─volser──┘
```

Specifies whether or not reserve activity is to be measured. The operands describe the type of data you require. Only one operand can be specified. The meaning of each operand field is:

ALLVSER
requests data describing all reserve requests. If you specify SENQR without operands, ALLVSER is the default.

volser
requests data describing the reserve requests for a particular device, where **volser** is the one to six character volume serial number of the volume.

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(ALLVSER), 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 ALLVSER 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(ALLVSER) or SENQR`

- To obtain reserve activity data for the device on which the volume TSO200 is mounted, enter:

  `SENQR(TSO200)`

SINTV

```
    ┌─SINTV(30S)────────────┐
►►──┤                       ├──────────────────────────────────────────────►◄
    └─SINTV(value─┬───┬─)───┘
                  └─S─┘
```

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 as a reminder that the specification is in seconds.

When you specify a small SINTV value, RMF overhead increases and excessive swapping can result; if this should occur, you could make the RMF address space non-swappable to minimize this overhead.

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 16-49.

SORT

ISPF session syntax

```
        ┌─A or D─(1)─┐
►►──SORT──┼────────────┼──────────────────────────────────────────►◄
        ├─A──────────┤
        └─D──────────┘
```

Note:
1 Numerical columns are sorted in descending order, and columns with character values are sorted in ascending order.

Sorts the report by the column in which the cursor is located either 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.

SPAG

Background session syntax

```
    ┌─NOSPAG─┐
►►──┼────────┼──────────────────────────────────────────────────►◄
    └─SPAG───┘
```

Display session syntax

```
►►──SPAG────────────────────────────────────────────────────────►◄
```

Specifies whether or not paging activity is to be measured.

SRCS

Background session syntax

```
    ┌─NOSRCS─┐
►►──┼────────┼──────────────────────────────────────────────────►◄
    └─SRCS───┘
```

Display session syntax

```
►►──SRCS────────────────────────────────────────────────────────►◄
```

Specifies whether or not central storage/processor/SRM activity is to be measured.

STOP

```
   ┌─STOP(10M)──────────────────┐
►►─┼────────────────────────────┼──────────────────────────────────────────►◄
   │                ┌─M─┐       │
   ├───STOP(value───┼───┼──)────┤
   │                └─H─┘       │
   └─NOSTOP─────────────────────┘
```

Specifies the desired duration for the session in minutes (M) or hours (H). The range is from one minute to one week (168 hours or 10,080 minutes). The default value is ten minutes. If you specify a value outside the range, RMF substitutes the default value. If neither M nor H is specified, M (minutes) is assumed. NOSTOP means the session can be ended only by a STOP command.

Notes:

1. 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 16-49.

SYSOUT

```
   ┌─SYSOUT(A)─────────┐
►►─┼───────────────────┼───────────────────────────────────────────────────►◄
   └─SYSOUT(class)─────┘
```

Specifies the SYSOUT class to which the formatted interval reports are directed. Class A is the default. You cannot modify the SYSOUT option during session processing.

SYSTEM

Display session syntax

```
►►─┬───────────────────┬───────────────────────────────────────────────────►◄
   └─SYSTEM smf_id─────┘
```

Specifies for which system in the sysplex the requested report is to be generated. For *smf_id*, specify the identifier of the desired system.

The format is the same for all types of Monitor II display session, but in an ISPF session you have the additional alternative of overtyping the 'SYSTEM' field in the headers of the report panels.

TRX

Background session syntax

```
     ┌─NOTRX───────────────────────┐
►►───┤                             ├──────────────────────────────────────────►◄
     │     ┌─(ALLPGN)──────────┐   │
     └─TRX─┤                   ├───┘
           └─(─┤ Operands ├─)──┘
```

Display session syntax

```
           ┌─ALLPGN─────────┐
►►──TRX────┤                ├─────────────────────────────────────────────────►◄
           └─┤ Operands ├───┘
```

Operands:

```
├──sname,nnnn──┬──────────────────────┬───────────────────────────────────────────┤
               └─,nnnn,nnnn:nnnn──────┘
```

This option is available in compatibility mode only and requests the Transaction Activity report.

The meaning of each operand field is:

sname
: Requests data for all performance groups associated with the subsystem identified by **sname**, where **sname** is a 1-4 character subsystem name that is defined in the installation control specification. For example, you can specify ASCH to request all performance groups associated with the ASCH address spaces, or OMVS for the UNIX System services.

nnnn[,nnnn,nnnn:nnnn]
: Requests data for one or more performance groups, where **nnnn** is a 1-4 digit performance group number. Performance groups can be specified individually (nnnn) or as a range (nnnn:nnnn), where the lower and upper bounds are separated by a colon. Each number or range of numbers must be separated by a comma.

ALLPGN
: Requests data for all performance groups. ALLPGN is the default.

The operands are not positional; thus, an omitted operand does not require a comma in its place. More than one operand can be specified. While there is no reason to combine ALLPGN with any other operand, you can specify on the same request both one or more subsystem names and individual performance groups or ranges of performance groups. In this case, if you specify individually a performance group that is also associated with a subsystem, data for that performance group appears twice in the report.

The menu default for the maximum number of performance groups that can be reported in response to any one request is 504. You can change the default as described in the *RMF Programmer's Guide*. by changing the MAXRBS operand of the PICTURE macro. Because an SMF record cannot exceed 32K bytes, the maximum number of performance group periods that can be included in an SMF record is 504. If you change the default to exceed 504, the record written to SMF will be truncated.

Examples

For a display session:

- To request data for performance group 1 and performance groups 5 through 10, enter

  `TRX 1,5:10`

- To request data for all performance groups associated with the subsystem name of ABC, which is defined in the installation control specification, enter

  `TRX ABC`

- To request data for all performance groups associated with the subsystem name of ABC and performance group 50, enter

  `TRX ABC,50`

  Note: If performance group 50 is also associated with the subsystem name of ABC, the data for performance group 50 will appear twice in the report.

- To request data for all performance groups, enter

  `TRX`

Examples

For a background session:

- To request data for performance group 1 and performance groups 5 through 10, enter:

  `TRX(1,5:10)`

- To request data for all performance groups associated with the subsystem name of ABC, which is defined in the installation control specification, enter:

  `TRX(ABC)`

- To request data for all performance groups associated with the subsystem name of ABC and performance group 50, enter:

  `TRX(ABC,50)`

  Note: If performance group 50 is also associated with the subsystem name of ABC, the data for performance group 50 will appear twice in the report.

- To request data for all performance groups, enter:

  `TRX`

USER

```
      ┌─NOUSER─┐
►►────┤        ├──────────────────────────────────────────────────────►◄
      └─USER───┘
```

Specifies whether or not a user-specified activity is to be measured. Your installation must supply a corresponding data-gathering module and data-reporting module before USER can take effect. See the *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 5-1). 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)
STOP value specified is less than INTERVAL	Indicates session termination before obtaining any data	Set STOP value equal to INTERVAL value
REPORT(DEFER) and NOSTOP specified	SYSOUT becomes cluttered with unprinted reports	Change NOSTOP to STOP set equal to INTERVAL value

Chapter 17. Long-term Overview Reporting with the Postprocessor

How the Postprocessor Works

This chapter deals with the Postprocessor, which generates reports from data gathered by the RMF monitors.

It describes:

- The different types of reports
- How to handle an SMF record data set
- How to start the Postprocessor
- The Postprocessor data sets
- All available options

Sysplex Reporting

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 first 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 types of reports based on SMF data that have been gathered by Monitor I, Monitor II, or Monitor III. In addition, you can create records for further processing with spreadsheet applications on the workstation.

- Interval and duration reports

  Here, you use the options REPORTS and SYSRPTS to get single-system and sysplex reports, with the additional option DINTV you will create duration reports combining data from several measurement intervals into one report interval. In "REPORTS" on page 17-27 and "SYSRPTS" on page 17-37, you find a description of all suboptions that you can select to get the reports you are looking for.

- Summary report

  With the option SUMMARY, you get a Summary report providing key performance data for a single system. There is no way to modify the contents of the report, see the following paragraph for creating tailorable reports.

The following two report types are based internally on the same technology. They use the same tables (see "Overview and Exception Conditions" on page 17-39) as input, and they use the same option ETOD for selecting a reporting time range. As introduced in an earlier release of RMF, option EXCEPT can be used to create each of the two types of reports, the selection which one is created will be made through the OVERVIEW option.

Since OS/390 2.6 RMF, there is a new option OVW which should be used to create Overview reports (and records), although EXCEPT is still available for single-system reports.

- Overview report

  The OVW option offers you 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 17-39 give you the information about all data that is available for overview processing, related to the SMF data which is the source for these values.

  Furthermore, in *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.

- Exception report

  For exception reporting, you can use the same names as described for overview reporting. The difference is that you will specify thresholds for these values as criterion whether the data will be reported or not.

- Plot report

  This report gives you a character-graphic representation of a number of key system performance indicators.

If plot reporting does not satisfy you reporting requirements, you can use spreadsheet applications for creating appropriate reports based on Postprocessor data.

- Overview records

  In the same way, as you can create Overview reports, you can create Overview records, just be specifying an addition option OVERVIEW(RECORD). You can download manually the records to the workstation for further processing, or you can use the Collector function of the Spreadsheet Reporter to create and submit Postprocessor jobs without logging on explicitly to the host system. In addition, several spreadsheets are available to create a wide spectrum of reports (see Chapter 18, “RMF Spreadsheet Reporter (RMFPP)” on page 18-1 for a detailed description).

Preparing SMF Records for Postprocessing

The Postprocessor generates reports based on data gathered in SMF records. Because the Postprocessor does not process virtual storage access method (VSAM) data sets (and SMF produces VSAM data sets), you must put the SMF records into a non-VSAM data set. You should use the IFASMFDP program to dump the records into a non-VSAM data set. For more information about the IFASMFDP program, see *OS/390 MVS System Management Facilities (SMF)*. Using other utilities to copy SMF records often results in truncated or unusable records.

The SMF records **must** be arranged by RMF interval start date and interval start time in the data set. If you want to combine SMF records from several VSAM data sets, you must arrange the records in the order they were written, or sort them. If your SMF record data set combines SMF records created by more than one system, sort the records into one data stream, ordered by interval start date and interval start time. Records for the same interval must remain in the order in which they were written. The input data sets to the Postprocessor should not be compressed.

Note: RMF provides two SORT exits (ERBPPE15 and ERBPPE35) that should be used when running the SORT program. To guarantee correct processing when you use other SORT programs than DFSORT*, you should ensure that you do not use the old SORT exit ERBPPSRT.

Use the sample job supplied with RMF in SYS1.SAMPLIB(ERBSAMPP) for sorting Postprocessor input:

Example

```
//ERBSAMPP JOB (acct),'pgmrname',CLASS=A,REGION=0M,..
//RMFSORT  EXEC PGM=SORT
//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=(NEW,PASS),UNIT=SYSDA,SPACE=(TRK,(200,200))
//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=(ERBPPE15,500,,N),E35=(ERBPPE35,500,,N)
```

You can then start the Postprocessor as described in “Starting the Postprocessor.” If you want to start the Postprocessor with JCL, you can use the statements shown in the example under “Starting with JCL” on page 17-5 as a second step in the sort job.

Recommendation

If you have large data sets with SMF records of all types, you can make the sort faster by tailoring the input data. To do this, insert an INCLUDE control statement in the SYSIN data definition of the above example, as follows:

```
//SYSIN DD *
  INCLUDE COND=(06,1,BI,GE,X'46',AND,06,1 BI,LE,X'4F')
  SORT ... (as above)
  MODS ... (as above)
```

This ensures that only RMF records are selected for sorting.

Starting the Postprocessor

Enablement

RMF is an optional feature of OS/390. 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 2-3.) Just follow the instructions on the following pages.

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:

- A data set on disk or tape.

  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 “Ensure Access to SMF Data for the Sysplex Data Services” on page 2-4 for details.

Starting with JCL

To start processing from a data set, use the following sample JCL:

Example

```
//EXAMPLE  JOB  accounting information,REGION=0M
//POST     EXEC PGM=ERBRMFPP
//MFPINPUT DD   data set containing SMF records (tape or DASD)
//SYSIN    DD     *
  control statements
/*
```

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 on the previous page with the Postprocessor start step shown above 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 the following:

Example

```
//EXAMPLE  JOB  accounting information,REGION=0M
//POST     EXEC PGM=ERBRMFPP
//SYSIN    DD    *
  control statements
/*
```

Control statements specify the Postprocessor options, and are described in "Control Statements" on page 17-13, with examples of control statements in "Examples of Control Statements" on page 17-61.

When you specify control statements, please note the following:

- 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 17-5 as input on the panels.

Postprocessor Setup Panel

```
                       RMF - Postprocessor Setup              OS/390 2.7.0 RMF
Command ===>

 Input Data         ===> ________  DATASET or SDS (Sysplex Data Server Buffers)

 Output Data        ===> ___       YES or NO (NO to route output to SYSOUT)

 Report Profile     ===> ______________________________________________

 Edit generated JCL ===> ___       YES or NO

Job Statement Information:
 ===> //uidA     JOB (acct),'pgmrname',CLASS=A,REGION=0M
 ===> //*
 ===> //*
 ===> //*

Complete this panel and press ENTER to continue, or END to exit.
```

Figure 17-1. 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 (see Figure 17-2 on page 17-8) where you may enter up to 14 SMF data set 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 17-3 on page 17-8).

  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
  - PLOTS
  - 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 control statements have to be specified according to rules as described in "Starting with JCL" on page 17-5.

  The options that are generated in the Postprocessor Option panel (see Figure 17-4 on page 17-9) will be added to the control statements that you define in your profile data set. The values from the Option panel will be 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 Data Sets

```
                    RMF - Postprocessor Input
Command ===>

 SMF Data Sets   ===> ____________________________________________
                 ===> ____________________________________________
                 ===> ____________________________________________
                 ===> ____________________________________________
                 ===> ____________________________________________
                 ===> ____________________________________________
                 ===> ____________________________________________
                 ===> ____________________________________________
                 ===> ____________________________________________
                 ===> ____________________________________________
                 ===> ____________________________________________
                 ===> ____________________________________________
                 ===> ____________________________________________
                 ===> ____________________________________________

 Sort Input Data ===> ___ YES or NO

 Press ENTER to continue. To return to previous panel, press END.
 To return to RMF Primary Menu without saving input, enter CANCEL.
```

Figure 17-2. Postprocessor - Input Data Sets

This panel allows you to specify all SMF input data sets.

In addition, you can specify whether the data sets are 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 RMF Primary Menu without saving input, enter CANCEL.
```

Figure 17-3. 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 17-10.

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
Command ===>

 Reporting (DATE) Start ===> __ . ___        End ===> __ . ___        yy.ddd
              or Start ===> 01 / 01 / 1948 End ===> 12 / 31 / 2047 mm/dd/yyyy
 Duration  (DINTV)      ===> __ : __
 Plot      (PINTV)      ===> __ : __
 Exception (ETOD) Start ===> 00 : 00        End ===> 24 : 00  hh:mm
 Plot      (PTOD) Start ===> 00 : 00        End ===> 24 : 00  hh:mm
 Interval  (RTOD) Start ===> 00 : 00        End ===> 24 : 00  hh:mm
 Summary   (STOD) Start ===> 00 : 00        End ===> 24 : 00  hh:mm

 Summary   (SUMMARY)    ===> INT,TOT       NO, INT, TOT, or TOT,INT
 Overview  (OVERVIEW)   ===> _____________ RECORD, REPORT (or both)

 DELTA     ===> NO_  YES or NO      MAXPLEN   ===> 050  Max no. plotted lines
 EXITS     ===> NO_  YES or NO      SESSION   ===> __   Session ID Monitor II
 SYSOUT    ===> A    Sysout Class   SYSID     ===> ____ 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.
```

Figure 17-4. Postprocessor - Options Panel

Here, you can enter:

- The start and end dates of the reporting period

- The time value for the duration and plot reports
- 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:

Figure 17-5 (Page 1 of 3). Postprocessor ddnames

ddname	Contents	Allocations	Notes
MFEnnnnn	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.
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 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. For creating one output data set, use ddname PPRPTS.
PPRPTS	Combined interval report	One ddname for one data set to contain all reports for each interval during the session.	There is no dynamic allocation of this ddname, you have to 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.

Figure 17-5 (Page 2 of 3). Postprocessor ddnames

ddname	Contents	Allocations	Notes
MFXnnnnn	Exception interval report data	One ddname is generated and one data set created for each interval included in the exception report	**nnnnn** is a decimal number from 00001 to 99999. The first interval is assigned the ddname MFX00001. The second MFX00002, and each subsequent interval is assigned a ddname in ascending numerical order. If no exception interval reports are produced in a given interval, the data set for that interval is empty. For creating one output data set, use ddname PPXRPTS.
PPXRPTS	Combined exception report	One ddname for one data set to contain all exception reports for each interval during the session.	There is no dynamic allocation of this ddname, you have to 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 MFXnnnnn files will be created.
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.
PPSUMnnn	Summary report output	One ddname is generated and one data set is created for each system included in the input data set.	**nnn** is a decimal number from 001 to 999. The first system encountered is assigned the ddname PPSUM001, and each subsequent system is assigned a ddname in ascending numerical order.
PPPLTnnn	Plot report output	One ddname is generated and one data set is created for each system included in the input data set.	**nnn** is a decimal number from 001 to 999. The first system encountered is assigned the ddname PPPLT001, and each subsequent system is assigned a ddname in ascending numerical order. The data set the Postprocessor creates has a variable blocked (VBA) record format and a logical record length of 137, a length that allows for two Plot reports per page. If your printer has a longer line length, you can preallocate this data set with a logical record length of 193 to allow three plots per page. A logical record length of 193 allows 186 characters for the three 62-character plots, two separator characters between the plots, one ASA control character, and four characters for the variable record header.
PPEXTnnn	Exception report output	One ddname is generated and one data set is created for each system included in the input data set.	**nnn** is a decimal number from 001 to 999. The first system encountered is assigned to ddname PPEXT001, and each subsequent system is assigned a ddname in ascending numerical order.
PPORPnnn	Overview report output	One ddname is generated and one data set is created for each system included in the input data set.	**nnn** is a decimal number from 001 to 999. The first system encountered is assigned to ddname PPORP001, and each subsequent system is assigned a ddname in ascending numerical order.

Figure 17-5 (Page 3 of 3). Postprocessor ddnames			
ddname	**Contents**	**Allocations**	**Notes**
PPXSRPTS	Sysplex report output	One ddname is generated and one data set is created for all sysplex reports.	All sysplex reports are written to this data set.

Notes:

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

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

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

2. Please consider that the number of dynamically allocated data sets is limited to 1635.

Figure 17-6. Postprocessor ddname - Record Output			
ddname	**Contents**	**Allocations**	**Notes**
PPOVWREC	Overview record output	Generate one ddname and create one data set.	**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)`

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, PTOD, 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 LE/370 messages issued by routines that the Postprocessor uses. Please see the *RMF Diagnosis 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

Control Statements

This chapter describes the Postprocessor control statements in alphabetical order. Figure 17-7 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.

Figure 17-7 (Page 1 of 2). Postprocessor Control Statement Summary

Processor Control Statement	Interval Report	Duration Report	Summary Report	Plot Report	Exception Report	Overview Report	Notes	See page
DATE	*	*	*	*	*	*	1	17-14
DELTA	*						1	17-15
DINTV		*					2,3	17-15
ETOD					*	*	1	17-18
EXCEPT					*	*	2	17-19

Figure 17-7 (Page 2 of 2). Postprocessor Control Statement Summary

Processor Control Statement	Interval Report	Duration Report	Summary Report	Plot Report	Exception Report	Overview Report	Notes	See page
EXITS	*	*	*	*	*	*	1	17-21
EXRPTS					*		2	17-21
MAXPLEN				*			1	17-22
OVERVIEW						*	2	17-23
OVW						*	2	17-24
PINTV				*			1	17-24
PLOTS				*			2	17-25
PTOD				*			1	17-26
REPORTS	*	*					2	17-27
RTOD	*	*					1	17-34
SESSION	*						1	17-35
STOD			*				1	17-35
SUMMARY			*				1	17-36
SYSID	*	*	*	*	*	*	1	17-36
SYSOUT	*	*	*	*	*	*	1	17-37
SYSRPTS	*	*					2	17-37

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, exception, and plot reports concurrently with either duration or interval reports.

Description of Control Statements

DATE

The DATE control statement specifies the range of dates of the reporting period for all reports.

The syntax of the statement is:

```
DATE{(yyddd,yyddd)}
    {(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 the year in full-century form, for example, `1999`

Note: 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 DATE(01011950,12312049); that is, all dates found in the input data set are reported.

When 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, 1999, specify

```
DATE(06071999,06071999)
```

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 *RMF Report Analysis*. The syntax of the statement is:

```
DELTA/NODELTA
```

DELTA indicates that the affected fields are to reflect changed values; that is, the reports are to be generated in delta mode. NODELTA indicates that the affected fields are to reflect total values. NODELTA is the default.

DINTV

The DINTV control statement specifies that the Postprocessor is to generate duration reports and indicates the length of the duration interval. The duration interval is the length of time each report can cover. The syntax of the statement is:

```
DINTV(hhmm)
```

Where `hh` is hours and `mm` is minutes. The maximum is `9960` which is equivalent to 100 hours.

The time specified is the length of the duration interval.

You **must** specify the DINTV statement if you want duration reports. If you omit the statement, the Postprocessor generates interval reports.

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 17-8 illustrates how the duration interval relates to the reporting period. Assume a reporting period that covers the twelve-hour block of time from 8:00 A.M. to 8:00 P.M. over a range of three days. As the figure shows, specifying DINTV(1200) causes the 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, would cause the Postprocessor to generate the following reports:

1. 8 A.M. to 6 P.M. on the first day of the reporting period.
2. 6 P.M. to 8 P.M. on the first day of the reporting period, and 8 A.M. to 4 P.M. on the second day.
3. 4 P.M. to 8 P.M. on the second day of the reporting period, and 8 A.M. to 2 P.M. on the third day.
4. 2 P.M. to 8 P.M. on the third day of the reporting period.

In this case, if you wanted to use the reports to compare system performance over the same hours on each day of the reporting period, it would be difficult because each report covers a different time range and some span two days and include a twelve-hour gap when no reporting was done.

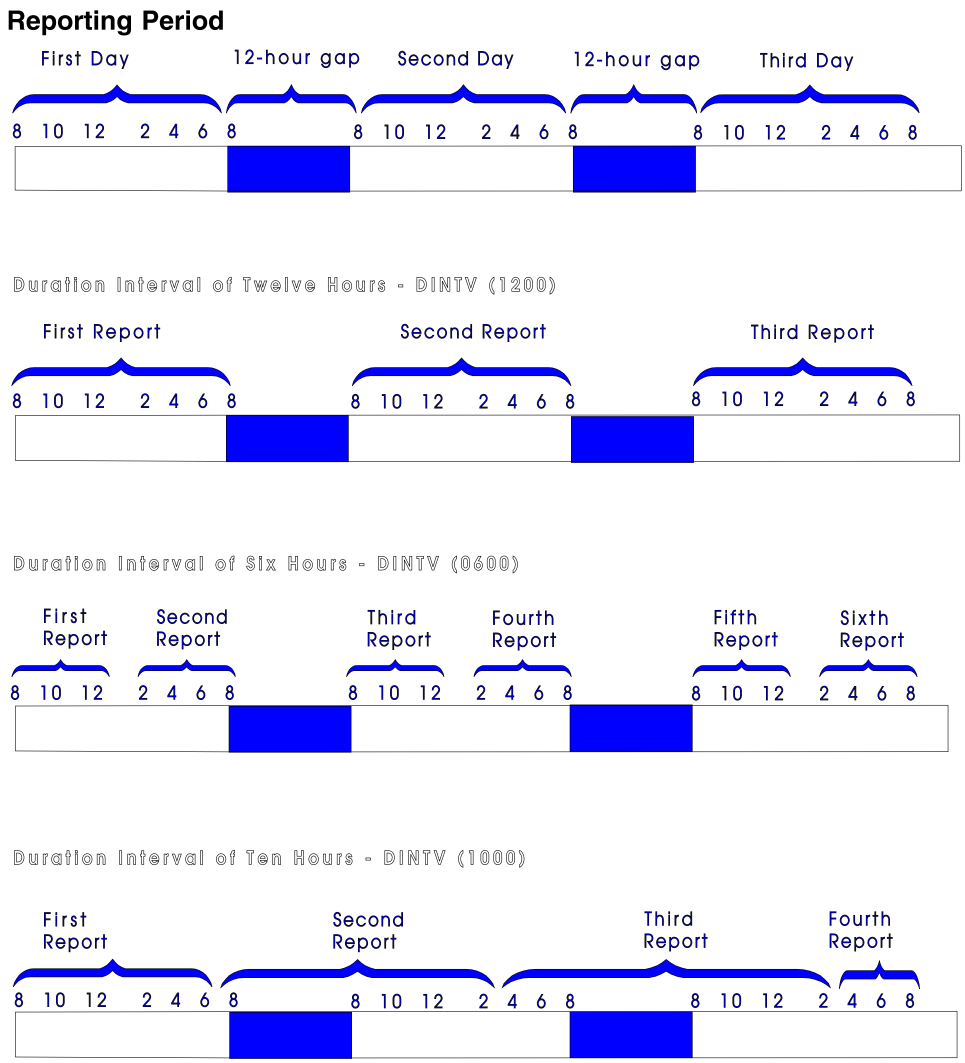

Figure 17-8. Relationship between Duration Interval and Reporting Period

The syntax of the control statement allows a maximum duration interval of 99 hours and 60 minutes.

However, care should be taken in selecting a duration interval. For most effective reporting, the duration interval should bear a relationship to the length of the reporting period.

Note: It is recommended to specify for duration reports not only the reporting interval but also the date (also if the SMF data set contains only records for those days you want to report on) because of performance reasons as well as of internal processing reasons.

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 A.M. to 8:00 P.M.) on June 7, 1999:

```
DATE(06071999,06071999)
RTOD(0800,2000)
DINTV(1200)
REPORTS(CHAN)
```

However, if channel path activity measurement did not begin until 10:00 A.M., and ended at 6:00 P.M., the duration report covers the time from 10:00 A.M. to 6:00 P.M. The start time of the first 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, plot, 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 is the hours and mm is the minutes on a 24-hour clock. Specify each time 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 must be later than the first, or a syntax error occurs.

When the ETOD statement is omitted, the default value is ETOD(0000,2400); that is, all times are reported. Thus, you would use this control statement when you want a reporting period for an Exception report that is different from the default value.

Because the range of values allowed is from 0000 to 2400, it is not possible to define a reporting period that consists of a single block of time 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 A.M. to 1:00 P.M. for the week beginning on January 3, 1999 and ending on January 9, 1999, the required DATE and ETOD statements would be:

Example

```
DATE(01031999,01091999)
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 – 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 17-39.

An exception consists of one condition or several conditions; when you group conditions into a single condition, all conditions must exceed their threshold values in order for RMF to recognize that an exception has occurred.

If you specify an OVERVIEW statement (see “OVERVIEW” on page 17-23) together with EXCEPT statements, RMF produces an Overview report based on the EXCEPT 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, but you specify as many EXCEPT statements as you require. The syntax of the EXCEPT statement is:

```
                                                   {LE}
EXCEPT([control-stmnt-name](condition-name[(qualifier)] ,{GE} ,threshold-value))
```

control-stmnt-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 17-39, 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 following tables also list the qualifiers.

LE or GE
Specifies the relational operator RMF is to use to determine if the exception has occurred.

LE indicates that any value in the SMF record that is less than or equal to the threshold value causes an exception.

GE indicates that any value in the SMF record that is greater than or equal to the threshold value causes an exception.

threshold-value
Specifies the value that RMF compares to a computed value from the appropriate SMF record fields. 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.

Examples

- The Postprocessor is to generate a line in the Exception report for each interval when the percent busy value for channel path 1 is greater than or equal to ten percent. Use the following control statement:

  ```
  EXCEPT((CHPBSY(01),GE,10))
  ```

  The absence of a control-statement-name indicates that this is a single condition exception and that no interval reports are generated.

- The Postprocessor is to generate a line in the Exception report for each interval when the I/O service rate for performance group number three is less than or equal to .01 and the device percent utilized value for device 06D8 is greater than or equal to 3. Use the following control statements:

  ```
  EXCEPT(EXCP1(IOSRV(0030),LE,.01))
  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 *RMF Programmer's Guide* for information on how to code a user exit routine for the Postprocessor.

NOEXITS indicates that no user exit routine is to be executed; this is the default.

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 or a blank. No continuation statements are permitted; however, you can use as many EXRPTS statements as you require. The syntax of the statement is:

```
EXRPTS (control-stmnt-name(report,report,..))
```

control-stmnt-name
Specifies a one to eight-character name starting with an alphabetic character that associates one or more EXCEPT statements with the EXRPTS statement.

The EXCEPT statement defines the exception to RMF; the EXRPTS statement defines the action that RMF is to take.

report
Specifies any one of the Monitor I session reports or the OMVS or XCF report that is acceptable on the REPORTS control statement. The following key-words can be specified for `report`. For their meaning, see “REPORTS” on page 17-27:

ALL
CACHE/NOCACHE
CHAN/NOCHAN
CPU/NOCPU
DEVICE(option,option ..)/NODEVICE
ENQ/NOENQ
IOQ/NOIOQ
OMVS/NOOMVS
PAGESP/NOPAGESP
PAGING/NOPAGING
TRACE/NOTRACE
VSTOR(option,list)/NOVSTOR
WKLD(list)/NOWKLD
XCF/NOXCF

Example

The Postprocessor is to generate a Channel Path Activity report for each interval when the I/O service rate for performance group number 3 is less than or equal to 100 service units per second. Use the following control statements:

```
EXCEPT(CHNRPT(IOSRV(0030),LE,100))
EXRPTS(CHNRPT(CHAN)))
```

MAXPLEN

The MAXPLEN control statement specifies the maximum number of lines of plotted data that can appear in a single Plot report. The syntax of the statement is:

```
MAXPLEN(nnn)
```

where **nnn** is the number of lines. The maximum value allowed is 999. When you omit the MAXPLEN statement, the default number of lines is 50. When you explicitly specify MAXPLEN, you must specify the number of lines.

When a Plot report exceeds one page, the plot lines run continuously from one page to the next to preserve the unity of the plotted data; there is no break or repetition of headings on each successive page.

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 statements (see "OVW" on page 17-24), 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 statements (see "EXCEPT" on page 17-19), 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 really do this, overview data is created using both types of statements, and you must ensure that no duplicate control statement names are being used.

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 17-10) as

```
//PPOVWREC DD DSNAME=data.set.name, ...
```

or in the Postprocessor Output Data Set panel (see "Postprocessor Output Data Sets" on page 17-8).

The data set should have these characteristics:

```
DCB=(RECFM=VB,LRECL=32756,BLKSIZE=32760)
```

You find a description of the records in *RMF Report Analysis*.

OVW

You can use the OVW control statement to define overview processing 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.

The syntax of the statement is:

```
OVW(control-stmnt-name(condition-name(qualifier)){,SYSTEMS})
                                                 {,NOSYSTEMS})
```

If you have performed in the past overview processing with EXCEPT statements, you should exchange them to OVW statements. Only the OVW statement enables you creating sysplex reports. The OVW syntax enforces that a control statement name is required, and it removes the necessity of specifying a threshold as required with the EXCEPT statement.

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 17-39.

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 A.M. and 2 P.M., and you want to create a report and 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)
```

PINTV

The PINTV control statement specifies the length of time that is to elapse between each line of plotted data, that is, the length of the plot interval. The syntax of the statement is:

```
PINTV(hhmm)
```

Where **hh** is the hour and **mm** is the minute on a 24-hour clock. Note that the time is the length of the plot interval, not the time of day when the plot reporting period begins.

When you omit the PINTV statement but indicate that Plot reports are required by specifying the PLOTS control statement, the Postprocessor generates a line in a Plot report for each RMF interval included in the reporting period. Thus, the default for the PINTV control statement is the length of the RMF measurement interval in effect during the reporting period. If you explicitly specify PINTV, you must explicitly specify a time value. Specifying PINTV without a time value causes a syntax error. Note that the PINTV statement is ignored when PLOTS is not specified.

PLOTS

The PLOTS control statement specifies that Plot reports are to be produced and identifies the plots that you require. The syntax of the statement is:

```
PLOTS(plot[,plot],..)
```

where **plot** can be any of the following:

ASCH	Maximum number of APPC address spaces.
BATCH	Maximum number of batch users produced for all channels on all processors.
CH(nn)	Channel path utilization, where **nn** is the hexadecimal channel path identifier. When CH is requested, the channel path identifier must be specified.
CPUID[(n)]	Processor busy percentage, where **n** is the processor identifier. When **n** is omitted, plots are produced for all processors.
CSAA	Maximum allocated CSA (below the 16-megabyte line).
CSAFP	Minimum number of free CSA pages (below the 16-megabyte line).
CUBDL(nnnn)	Average number of milliseconds that any I/O requests for the device were delayed because the control unit was busy, where **nnnn** is the hexadecimal device number of the device to be included in the report. When the control unit busy delay time plot is requested, **nnnn** must be specified to identify the device.
DAC(nnnn)	Device activity rate, where **nnnn** is the hexadecimal device number of the device to be included in the report. When the device activity plot is requested, **nnnn** must be specified to identify the device.
DBDL(nnnn)	Average number of milliseconds that I/O requests for the device were delayed because the device was busy, where **nnnn** is the hexadecimal device number of the device to be included in the report. When the device busy delay time plot is requested, **nnnn** must be specified to identify the device.
DPBDL(nnnn)	Director port busy delay percentage, where **nnnn** is the hexadecimal device number.
DCON(nnnn)	Device connect time, where **nnnn** is the hexadecimal device number. When device connect time is requested, **nnnn** must be specified to identify the device.
DRT(nnnn)	Device response time where **nnnn** is the hexadecimal device number. When device response time is requested, **nnnn** must be specified to identify the device.
ESMR	Average rate of page migration from expanded to auxiliary storage.
HUIC	Average high unreferenced interval count.

IOAC(nnnn)	Rate of requests successfully initiated for the logical control unit, where **nnnn** is the hexadecimal identifier of the logical control unit. When IOAC is requested, **nnnn** must be specified to identify the logical control unit.
OMVS	Maximum number of OMVS address spaces.
PAGE	Demand paging rate.
PTES	Average paging transfer rate to expanded storage.
SEV	Service rate.
SQAA	Maximum allocated SQA (below the 16-megabyte line).
SQAE	Amount of SQA expansion into CSA (below the 16-megabyte line).
SQAFP	Minimum number of free SQA pages (below the 16-megabyte line).
STC	Maximum number of started task and mount task users.
SWA	Swap rate SWA option.
TPAG	Total paging rate.
TRA	Transaction rate.
TSO	Maximum number of TSO sessions.

Each type of plot is described in more detail in *RMF Report Analysis*.

When you omit the PLOTS statement, no Plot reports can be produced. If you specify PLOTS, you must identify at least one plot type. No continuation statements are permitted for the PLOTS statement. When you require more plot types than can fit on one statement, specify the additional plot types by supplying additional PLOTS statements. There is no limit on the number of PLOTS statements. For example:

Example

```
PLOTS(DRT(1230),DCON(1230),CH(04))
PLOTS(DAC(1350))
PLOTS(STC,SEV)
PLOTS(PAGE,TPAG,SWA,CPUID)
```

Note: The order in which you specify the plot types does not affect the order in which RMF generates the requested Plot reports.

PTOD

The PTOD control statement specifies the starting and ending time of the reporting period for a Plot report for each day in the reporting period. The syntax of the statement is:

```
PTOD(hhmm,hhmm)
```

where **hh** is the hour and **mm** is the minute on a 24-hour clock. The first time specifies the beginning of the reporting period and the second time specifies the end of the reporting period. The second time you specify must be later than the first, or a syntax error occurs.

When the PTOD statement is omitted, the default value is PTOD(0000,2400); that is, all times are reported. Thus, you would use this control statement when you want a reporting period for Plot reports that is different from the default value.

Because the range of values allowed is from 0000 to 2400, it is not possible to define a reporting period that consists of a single block of time that spans more than one calendar day. Thus, you cannot define a reporting period that runs from 12 noon on one day to 12 noon on the next day. However, you can define a reporting period that consists of the same block of time over several days. For example, to produce a Plot report using data collected from 8:00 A.M. to 4:00 P.M. for the week beginning on January 3, 1999, and ending on January 9, 1999, the required DATE and PTOD statements would be:

Example

```
DATE(01031999,01091999)
PTOD(0800,1600)
```

REPORTS

The REPORTS control statement specifies the reports to be generated by the Postprocessor for a single system. (See "SYSRPTS" on page 17-37 for sysplex reporting.) This statement is required to generate duration reports or interval reports. If you have specified DINTV, duration reports are generated. If not, interval reports are generated. For enqueue activity (ENQ), no duration report is available.

Each report specified must be separated from any other reports by a comma or a blank. No continuation statements are permitted; however, you can use as many REPORTS statements as you require. The syntax of the statement is:

```
REPORTS(report[,report],..)
```

where **report** can be any of the reports listed here and described in detail later in this section:

```
ALL                  HFS]NOHFS
ARD()]NOARD          IOQ]NOIOQ
ARDJ()]NOARDJ        IOQUEUE]NOIOQUEUE
ASD()]NOASD          OMVS]NOOMVS
ASDJ()]NOASDJ        PAGESP]NOPAGESP
ASRM()]NOASRM        PAGING]NOPAGING
ASRMJ()]NOASRMJ      PGSP()]NOPGSP
CACHE()]NOCACHE      SENQ()]NOSENQ
CHAN]NOCHAN          SENQR()]NOSENQR
CHANNEL]NOCHANNEL    SRCS]NOSRCS
CPU]NOCPU            SPAG]NOSPAG
DDMN]NODDMN          TRACE]NOTRACE
DEV()]NODEV          TRX()]NOTRX
DEVICE()]NODEVICE    VSTOR()]NOVSTOR
DEVV()]NODEVV        WKLD()]NOWKLD
ENQ]NOENQ            XCF]NOXCF
```

ALL indicates that all of these reports are to be generated, if gathered data is available.

When you specify an option either alone or in combination with ALL, the reports generated depend on whether you have used the negative value or the positive value of the option. This concept can perhaps best be illustrated by the following examples:

- When you specify REPORTS(CPU), the only report generated is the CPU Activity report.
- When you specify either REPORTS(NOCPU) or REPORTS(ALL,NOCPU), the Postprocessor generates all reports except for processor activity.
- When you specify REPORTS(ALL,CPU) or REPORTS(ALL), the Postprocessor generates all reports.

Notes:

1. There is no additional overhead introduced when you specify REPORTS(ALL) when some system activities were not measured during the reporting period.
2. When the DINTV statement is present, only the options ALL, CPU, CHAN, DEVICE, WKLD, PAGESP, IOQ, PAGING, VSTOR, OMVS, or XCF are accepted. The other options are described in the following list.
3. Regardless of the specification on the REPORTS statement, the Postprocessor can generate a report on a system activity only when that activity has been measured during the reporting period.

ALL

Specifies that all reports are to be generated. Any user-supplied Monitor II background session reports are also included when ALL is specified. ALL can be combined with explicit specifications of other options. For example, specifying:

Example

```
REPORTS(ALL,NOENQ,DEVICE(NOUNITR,NOCOMM))
```

causes all of the reports to be generated, with the exception of enqueue activity and device activity for unit record devices and communication equipment.

When you specify ALL on the REPORTS statement and the DINTV statement is present, duration reports are produced for processor, channel path, device, workload, page/swap data set, I/O queuing, paging activity, OMVS activity, and XCF activity.

ARD[(class,status,domain)]] NOARD

Specifies the Monitor II Address Space Resource Data report. **Class, status,** and **domain** specify the selection conditions for the address spaces to be included.

ARDJ(jobname)] NOARDJ

Specifies the Monitor II Address Space Resource Data by Jobname report, where **jobname** identifies a specific job for the report. If you specify ADRJ for the Postprocessor, make sure you have specified the Monitor II ARDJ option. You cannot run a Postprocessor ARDJ report with data collected by the Monitor II ARD option.

ASD[(class,status,domain)]] NOASD
Specifies the Monitor II Address Space State Data report, where **class, status**, and **domain** are as described under ARD.

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.

ASRM[(class,status,domain)]] NOASRM
Specifies the Monitor II Address Space SRM Data report, where **class, status,** and **domain** are as described above under ARD.

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.

CACHE([SSID(*list*)][,EXSSID(*list*)][,DEVICE] SUBSYS][,SUMMARY])] NOCACHE
Specifies the Monitor I Cache Subsystem Activity report.

SSID(*list*)
Specifies a list of storage subsystem identifiers (SSIDs), identifying the control units to be included in the report. You can specify as many SSIDs as you like. 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 as many SSIDs as you like. 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.

If EXSSID is not specified, all control units in the SSID option list are reported on, or, if the SSID option has not been specified, all control units are selected.

DEVICE]SUBSYS
Specify DEVICE to create a report on device level and additionally a report on subsystem level for each control unit selected.

Specify SUBSYS to create reports on subsystem level only.

SUMMARY
Specify SUMMARY to create a Summary report. You can specify this parameter in addition to the other parameters, if you want to get the Subsystem report or the Device report, too. IF SUMMARY is the only parameter, you get *only* the Summary report.

CHAN] NOCHAN
Specifies the Monitor I Channel Path Activity report.

CHANNEL] NOCHANNEL
Specifies the Monitor II Channel Path Activity report.

CPU] NOCPU
Specifies the Monitor I CPU Activity report.

DDMN] NODDMN
Specifies the Monitor II Domain 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.

DEVICE(option,option..)] 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 options 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:

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

- Storage groups in the form SG (aaa,bbb) where aaa and bbb are 1 to 8 character names. 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. f 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:

Example

```
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:

Examples

```
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 a specific device number. RMF reports a storage group name for each volume assigned to one.

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.

HFS/NOHFS
Specifies an HFS Statistics report.

IOQ/NOIOQ
Specifies a Monitor I I/O Queuing Activity report.

IOQUEUE] NOIOQUEUE
Specifies a Monitor II I/O Queuing Activity report.

OMVS] NOOMVS
Specifies an OMVS Kernel Activity report.

PAGESP] NOPAGESP
Specifies the Monitor I Page/Swap Data Set Activity report.

PAGING] NOPAGING
Specifies the Monitor I Paging Activity report.

PGSP] PGSP(option)] NOPGSP
Specifies the Monitor II Page/Swap Data Set Activity report. The following options are possible:

PAGE/SWAP Get Page or Swap data set activity

Specifying PGSP is equivalent with PGSP(PAGE).

SENQ] SENQ(option)] NOSENQ

Specifies the Monitor II Enqueue Activity 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. Each operand is described in Details of Session and Report Commands(page “SENQ” on page 16-41).

Notes:

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 Reserve Activity report. The operands describe the type of data you require:

ALLVSER	Report on all volumes
volser	Report on a specific volume

Specifying SENQR is equivalent with SENQR(ALLVSER).

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.

SRCS] NOSRCS

Specifies the Monitor II Central Storage/Processor/SRM report.

SPAG] NOSPAG

Specifies the Monitor II Paging Activity report.

TRACE] NOTRACE

Specifies the Monitor I Trace Activity report.

TRX [(sname,nnn[,nnn,nnn:nnn],<u>ALLPGN</u>)]] <u>NOTRX</u>

Specifies the Monitor II Transaction Activity 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 *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(option)] NOVSTOR

Specifies the 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)

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, indeed, 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.

WKLD(list)] NOWKLD

Specifies whether the Monitor I Workload Activity report for compatibility mode data is to be generated and indicates the types of sub-reports to be included. When the DINTV statement is present, a duration report is produced; otherwise, an interval report is produced.

When WKLD is specified, **list** must indicate the types of sub-reports to be included. Any or all of the following can be specified in the list:

PERIOD/NOPERIOD Requests reporting by performance group period

GROUP/NOGROUP Requests reporting by performance group

RANGE/NORANGE Requests reporting by performance objective, domain, and performance group in a range of performance group numbers

DOMAIN/NODOMAIN Requests reporting by domain number

SYSTEM/NOSYSTEM Requests reporting by system

TIME/NOTIME Requests reporting by time slice group

When you omit WKLD and specify ALL, the underscored reports are generated. When you specify WKLD, you must include **list**. When you specify a report type in the **list** field, the reports generated depend on whether you have used the negative value or the positive value of the option. If you use the negative value of the option, you get all the Workload Activity reports except the options you specify. For example, WKLD(NODOMAIN) causes the Postprocessor to generate all the Workload Activity reports except the reports by domain number. If you use the positive value of the option, you get only the report for that option. For example, WKLD(DOMAIN) causes the Postprocessor to generate only the domain sub-report; no other reports are printed.

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 that 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 A.M. to 1:00 P.M. for the week beginning on January 3, 1999, and ending on January 9, 1999, the required DATE and RTOD statements would be:

```
DATE(01031999,01091999)
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

A.M. to 1:00 P.M. for the week beginning on January 3, 1999, and ending on January 9, 1999, the required DATE and STOD statements would be:

Example

```
DATE(01031999,01091999)
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 **SYNC(SMF)** option for all systems. 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. There is no duration report for coupling facility activity.

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 goal mode version of 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)]]<u>NOPOLICY</u>

Specify policy names in `namelist`. For each policy specified, the Postprocessor issues a summary report.

WGROUP[(namelist)]]<u>NOWGROUP</u>

Specify workload group names. For each workload group specified, the Postprocessor issues a summary report.

SCLASS[(namelist)]]<u>NOSCLASS</u>

Specify service class names in `namelist`. For each service class, the Postprocessor issues a summary report.

SCPER[(namelist)]]<u>NOSCPER</u>

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.

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.

RCLASS[(namelist)]]NORCLASS
Specify report class names in `namelist`. The Postprocessor issues reports of the specified classes.

SYSNAM[(namelist)]]NOSYSNAM
Specify system names in `namelist`. The Postprocessor combines data from all the specified systems in one report.

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

One qualifier is possible:

cpuid A one-digit processor identifier. If the qualifier is omitted, the values represents the average of all processors.

Figure 17-9 (Page 1 of 2). CPU Activity - Conditions Based on SMF Record Type 70

Condition	Condition Name	Qualifier	Source	Algorithm
Percent processor busy	CPUBSY	cpuid	SMF70WAT SMF70CID SMF70INT SMF70INT SMF70CPN	If a processor identifier is specified, RMF: 1. Finds the processor section for the specified processor. 2. Calculates the percent busy from the processor wait time. 3. Compares the percent busy against the exception threshold. If a processor identifier is not specified, RMF: 1. Examines all processor data sections. 2. Calculates the percent busy for each processor from the processor wait time. 3. Calculates the average percent busy for all processors. 4. Compares the average percent busy against the threshold value.
Maximum number of batch users	MXBATCH	none	SMF70BMM	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

Figure 17-9 (Page 2 of 2). CPU Activity - Conditions Based on SMF Record Type 70

Condition	Condition Name	Qualifier	Source	Algorithm
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
Percent MVS busy	MVSBSY	cpuid	Same as for CPUBSY	Same as for CPUBSY
Number of processors at least partially online during the reporting interval	NUMPROC	none	SMF70CPN	Value or comparison
Percent of reporting interval at least one job could not be dispatched	OCPU1	none	SMF70CPN SMF70R02 to SMF70R15	Summarize SMF70Rnn to SMF70R15 for nn = SMF70CPN+1 Restriction: only valid if all CPUs were online during the reporting period.
Percent of reporting interval at least two or three jobs could not be dispatched	OCPU2 OCPU3	none	See OCPU1	See OCPU1 for nn=SMF70CPN+2 and nn=SMF70CPN+3
Average number of out and wait users	AVGUOWT	none	SMF70WTT	WTT/SAM
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

Paging Activity - SMF Record Type 71

Figure 17-10 (Page 1 of 4). 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
Swap out due to long wait time	PSOCLW	none	SMF71AXD(3) SMF71LES(3) SMF71LAX(3) SMF71ESD(3) SMF71MIG(3) SMF71INT	(AXD(3))+LES(3)+LAX(3) +ESD(3)+MIG(3))/INT

Figure 17-10 (Page 2 of 4). Paging Activity - Conditions Based on SMF Record Type 71

Condition	Condition Name	Qualifier	Source	Algorithm
Swap-out due to detected wait	PSOCDW	none	SMF71AXD(6) SMF71LES(6) SMF71LAX(6) SMF71ESD(6) SMF71MIG(6) SMF71INT	(AXD(6)+LES(6)+LAX(6) +ESD(6)+MIG(6))/INT
Number of unilateral swap-outs	PSOCU	none	SMF71AXD(10) SMF71LES(10) SMF71LAX(10) SMF71ESD(10) SMF71MIG(10) SMF71INT	(AXD(10)+LES(10)+LAX(10) +ESD(10)+MIG(10))/INT
Swap-out on recommendation value	PSOCEORV	none	SMF71AXD(9) SMF71LES(9) SMF71LAX(9) SMF71ESD(9) SMF71MIG(9) SMF71INT	(AXD(9)+LES(9)+LAX(9) +ESD(9)+MIG(9))/INT
Swap-out on enqueue exchange	PSOCENQE	none	SMF71AXD(8) SMF71LES(8) SMF71LAX(8) SMF71ESD(8) SMF71MIG(8) SMF71INT	(AXD(8)+LES(8)+LAX(8) +ESD(8)+MIG(8))/INT
Number of requested swap-outs	PSOCREQ	none	SMF71AXD(7) SMF71LES(7) SMF71LAX(7) SMF71ESD(7) SMF71MIG(7) SMF71INT	(AXD(7)+LES(7)+LAX(7) +ESD(7)+MIG(7))/INT
Swap-out due to auxiliary storage shortage	PSOCAXSS	none	SMF71AXD(4) SMF71LES(4) SMF71LAX(4) SMF71ESD(4) SMF71MIG(4) SMF71INT	(AXD(4)+LES(4)+LAX(4) +ESD(4)+MIG(4))/INT
Swap-out due to transition to non-swappable	PSOCTONS	none	SMF71AXD(11) SMF71LES(11) SMF71LAX(11) SMF71ESD(11) SMF71MIG(11) SMF71INT	(AXD(11)+LES(11)+LAX(11) +ESD(11)+MIG(11))/INT
Number of page faults per second	PAGERT	none	SMF71PIN SMF71INT	PIN/INT
Demand paging per second	DPAGRT	none	SMF71PIN SMF71POT SMF71INT	(PIN+POT)/INT
Swap rate	SWART	none	SMF71SSQ SMF71INT	SSQ/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 frames	MXSQA	none	SMF71MXQ	Value or comparison
Average number of SQA frames	AVGSQA	none	SMF71AVQ	Value or comparison
Maximum number of CSA fixed-frames	MXCSAF	none	SMF71MXC	Value or comparison
Maximum number of total CSA frames	MXCSAT	none	SMF71MXP	Value or comparison
Average number of total CSA frames	AVGCSAT	none	SMF71AVP	Value or comparison
Average number of CSA fixed frames	AVGCSAF	none	SMF71AVC	Value or comparison
Average number of VIO allocated slots	AVGVIOF	none	SMF71AVV	Value or comparison

Figure 17-10 (Page 3 of 4). Paging Activity - Conditions Based on SMF Record Type 71				
Condition	**Condition Name**	**Qualifier**	**Source**	**Algorithm**
Maximum number of allocated VIO slots	MAXVIOF	none	SMF71MXV	Value or comparison
Swap-out terminal wait	PSOCTW	none	SMF71AXD(1) SMF71LES(1) SMF71LAX(1) SMF71ESD(1) SMF71MIG(1) SMF71AXD(2) SMF71LES(2) SMF71LAX(2) SMF71ESD(2) SMF71MIG(2) SMF71INT	((AXD(1)+LES(1)+LAX(1) +ESD(1)+MIG(1)+AXD(2) +LES(2)+LAX(2)+ESD(2) +MIG(2))/INT
Page move rate	PGMVRT	none	SMF71PMV SMF71INT	PMV/INT
Swap-out due to central storage shortage	PSOCRPSS	none	SMF71RSS SMF71INT	RSS/INT
Average high unreferenced interval count for central storage frames	AVGHUIC	none	SMF71ACA	ACA*100
Maximum high unreferenced interval count for central storage frames	MXHUIC	none	SMF71HIC	HIC*1000
Page movement rate to expanded storage	PTES	none	SMF71PES	PES/INT
Migration rate from expanded to auxiliary storage	ESMR	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 ES frames	ESTORAVM	none	SMF71EAM	Value or comparison
Maximum number of available ES frames	ESTORAVX	none	SMF71EAX	Value or comparison
Average number of available ES frames	ESTORAVA	none	SMF71EAA	Value or comparison
Minimum number of low-impact ES frames	ESTORLIM	none	SMF71ELM	Value or comparison
Maximum number of low-impact ES frames	ESTORLIX	none	SMF71ELX	Value or comparison

Condition	Condition Name	Qualifier	Source	Algorithm
Average number of low-impact ES frames	ESTORLIA	none	SMF71ELA	Value or comparison
Minimum number of medium-impact ES frames	ESTORMIM	none	SMF71EMM	Value or comparison
Maximum number of medium-impact ES frames	ESTORMIX	none	SMF71EMX	Value or comparison
Average number of medium-impact ES frames	ESTORMIA	none	SMF71EMA	Value or comparison
Minimum number of high-impact ES frames	ESTORHIM	none	SMF71EHM	Value or comparison
Maximum number of high-impact ES frames	ESTORHIX	none	SMF71EHX	Value or comparison
Average number of high-impact ES frames	ESTORHIA	none	SMF71EHA	Value or comparison

Figure 17-10 (Page 4 of 4). Paging Activity - Conditions Based on SMF Record Type 71

Workload Activity (Compatibility Mode) - SMF Record Type 72-1

One qualifier **type** is possible:

group A number in the form nnnnp where nnnn is a three-digit or four-digit performance group number or report performance group number in the range of 0000 through 9999, and p is a one-digit performance group period in the range 1 through 8, which should be 0 or 1 for report performance groups.

For exception processing only:

When you specify a performance group period of zero, the threshold applies to the sum of all performance group periods of the performance group. When you specify a performance group period of 1 through 8, the threshold applies only to the specified performance group period for the performance group. If the qualifier is omitted, the threshold applies to the sum of all performance group periods.

subsys A one- to four-byte subsystem name.

dmn A one- to three-digit domain number in the range 0 through 128.

Figure 17-11 (Page 1 of 3). Workload Activity (Compatibility Mode) - Conditions Based on SMF Record Type 72-1

Condition	Condition Name	Qualifier	Source	Algorithm
Total service per second	TOTSRV	type	SMF72SER SMF72INT	SER/INT
I/O service rate per second	IOSRV	type	SMF72ITS SMF72INT	ITS/INT
Processor service per second	CPUSRV	type	SMF72CTS SMF72INT	CTS/INT
Total SRB service per second	SRBSRV	type	SMF72STS SMF72INT	STS/INT
Central storage service per second	MSOSRV	type	SMF72MTS SMF72INT	MTS/INT
Number of transactions per second	TRANS	type	SMF72TTX SMF72INT	TTX/INT
Transaction time in seconds	RTIME	type	SMF72TTM SMF72TTX	TTM/TTX

Figure 17-11 (Page 2 of 3). Workload Activity (Compatibility Mode) - Conditions Based on SMF Record Type 72-1

Condition	Condition Name	Qualifier	Source	Algorithm
Number of swaps per transaction	SPERTRA	type	SMF72TTX SMF72SPP	SPP/TTX
Absorption rate	ABSRPTN	type	SMF72SER SMF72TAT	SER/TAT
Transaction service	TRXSERV	type	SMF72SER SMF72ACT	SER/ACT
Execution velocity	EXVEL	type	SMF72TOU SMF72TOT	TOU/(TOU+TOT)
TCB execution time in seconds	TCBSEC	type	SMF72CTS SMF72CSD	CTS/CSD
SRB execution time in seconds	SRBSEC	type	SMF72STS SMF72SSD	STS/SSD
Region control task time in seconds	RCTSEC	type	SMF72RCT	Value or comparison
I/O interrupt time in seconds	IITSEC	type	SMF72IIT	Value or comparison
Hyperspace service time in seconds	HSTSEC	type	SMF72HST	Value or comparison
Application execution time in seconds	APPLSEC	type	SMF72CTS SMF72STS SMF72RCT SMF72IIT SMF72HST SMF72CSD SMF72SSD	(CTS/CSD) + (STS/SSD) + RCT + IIT + HST
TCB execution time in percent of interval	TCBPER	type	SMF72CTS SMF72CSD	((CTS/CSD)/INT)*100
SRB execution time in percent of interval	SRBPER	type	SMF72STS SMF72SSD	((STS/SSD)/INT)*100
Application execution time in percent of interval	APPLPER	type	SMF72CTS SMF72STS SMF72RCT SMF72IIT SMF72HST SMF72CSD SMF72SSD	(((CTS/CSD) + (STS/SSD) + RCT + IIT + HST)/INT) * 100
Rate of demand page-ins from auxiliary storage while the transactions are resident in central storage	SINGLE	type	SMF72PIN SMF72TAT	PIN/TAT
Rate of block page-ins from auxiliary storage while the transactions are resident in central storage	BLOCK	type	SMF72BPI SMF72TAT	BPI/TAT
Rate of demand page-moves of single pages from expanded storage	EXPSNGL	type	SMF72PIE SMF72TAT	PIE/TAT
Rate of blocked page-moves from expanded storage	EXPBLK	type	SMF72BPE SMF72TAT	BPE/TAT
Rate of standard hiperspace pages read into central storage from auxiliary storage	HSP	type	SMF72HIN SMF72TAT	HIN/TAT
Rate of expanded-storage-only hiperspace read misses	HSPMISS	type	SMF72HRM SMF72TAT	HRM/TAT
Rate of shared storage page-ins	SHARED	type	SMF72SPA SMF72TAT	SPA/TAT
Rate of shared storage page-ins from expanded storage	EXPSHR	type	SMF72SPE SMF72TAT	SPE/TAT
Total number of EXCPs	EXCP	type	SMF72ITS SMF72ISD	ITS/ISD
Rate of EXCPs	EXCPRT	type	SMF72ITS SMF72ISD	(ITS/ISD)/INT

Figure 17-11 (Page 3 of 3). Workload Activity (Compatibility Mode) - Conditions Based on SMF Record Type 72-1

Condition	Condition Name	Qualifier	Source	Algorithm
Total number of central storage frames allocated to resident ASIDs	STOCEN	type	SMF72FT1 SMF72FT2 =SMF72FRS	Value or comparison
Total number of expanded storage frames allocated to resident ASIDs	STOEXP	type	SMF72ER1 SMF72ER2 =SMF72ERS	Value or comparison
Total number of shared storage frames allocated to resident ASIDs	STOSHR	type	SMF72SRS	Value or comparison
Total number of storage frames allocated to resident ASIDs	STOTOT	type	SMF72FRS SMF72ERS	FRS+ERS **Note:** For TSO and Batch, this does not reflect the Address space working set size, so it should not be used for storage capacity estimates
Number of transactions that ended during the interval	TRANSTOT	type	SMF72TTX	Value or comparison
Average number of transactions resident in central storage during the interval	TRANSMPL	type	SMF72TAT	TAT/INT
Average number of active transactions during the interval	TRANSAVG	type	SMF72ACT	ACT/INT
Total transaction time (queue and execution time)	RTIMETOT	type	SMF72TST SMF72TTX	TST/TTX
Average queue time a job waited on a JES or APPC queue or during TSO logon	RTIMEQUE	type	SMF72TTM SMF72TST SMF72TTX	(TST-TTM)/TTX
Transaction ineligible queue time	TRANSIQT	type	SMF72IQT SMF72TTX	IQT/TTX
Transaction R/S affinity delay time	TRANSADT	type	SMF72ADT SMF72TTX	ADT/TTX
Transaction JCL conversion time	TRANSCVT	type	SMF72CVT SMF72TTX	CVT/TTX
Start subchannel rate	SSCHRT	type	SMF72IRC INTERVAL	IRC/Interval
Average DASD response time	RESP	type	SMF72ICT SMF72IWT SMF72IDT SMF72IOT SMF72IRC	(IRC+IWT+IDT+IOT) / IRC
Average DASD connect time	CONN	type	SMF72ICT SMF72IRC	ICT/IRC
Average DASD disconnect time	DISC	type	SMF72IDT SMF72IRC	IDT/IRC
Average DASD pending time	QPEND	type	SMF72IWT SMF72IRC	IWT/IRC
Average DASD IOS queue time	IOSQ	type	SMF72IOT SMF72IRC	IOT/IRC

Workload Activity (Goal Mode) - SMF Record Type 72-3

The following table is valid only for overview processing, not for exception reporting. Dependend 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`	Report class
`W.wname`	Workload
`POLICY`	Policy

period Mandatory - service class period or report class

Figure 17-12 (Page 1 of 3). 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 R723MCPU R723MADJ	Sum(R723CCPU / (R723MCPU * R723MADJ))
SRB seconds	SRBSEC	type	R723CSRB R723MSRB R723MADJ	Sum(R723CSRB / (R723MSRB * R723MADJ))
Region Control Task (RCT) seconds	RCTSEC	type	R723CRCT	Sum(R723CRCT)
I/O interrupt (IIT) seconds	IITSEC	type	R723CIIT	Sum(R723CIIT)
Hiperspace service (HST) seconds	HSTSEC	type	R723CHST	Sum(R723CHST)
Application execution time	APPLSEC	type	R723CCPU R723CSRB R723CRCT R723CIIT R723CHST R723MCPU R723MSRB R723MADJ	Sum(R723CCPU / (R723MCPU * R723MADJ) + R723CSRB / (R723MSRB * R723MADJ) + R723CRCT + R723CIIT + R723CHST)
TCB second percentage	TCBPER	type	R723CCPU R723MCPU R723MADJ Interval	Sum(R723CCPU / (R723MCPU * R723MADJ)) / Interval * 100
SRB second percentage	SRBPER	type	R723CSRB R723MSRB R723MADJ Interval	Sum(R723CSRB / (R723MSRB * R723MADJ)) / Interval * 100

Figure 17-12 (Page 2 of 3). Workload Activity (Goal Mode) - Conditions Based on SMF Record Type 72-3				
Condition	**Condition Name**	**Qualifier**	**Source**	**Algorithm**
Application execution time percentage	APPLPER	type	R723CCPU R723CSRB R723CRCT R723CIIT R723CHST R723MCPU R723MSRB R723MADJ Interval	Sum(R723CCPU / (R723MCPU * R723MADJ) + R723CSRB / (R723MSRB * R723MADJ) + R723CRCT + R723CIIT + R723CHST) / Interval * 100
Page-in rate from auxiliary storage	SINGLE	type	R723CPIR R723CTRR	Sum(R723CPIR) / Sum(R723CTRR)
Block page-in rate from auxiliary storage	BLOCK	type	R723CBPI R723CTRR	Sum(R723CBPI) / Sum(R723CTRR)
Page-in rate from expanded storage	EXPSNGL	type	R723CPIE R723CTRR	Sum(R723CPIE) / Sum(R723CTRR)
Block page-in rate from expanded storage	EXPBLK	type	R723CBPE R723CTRR	Sum(R723CBPE) / Sum(R723CTRR)
Hiperspace page-in rate	HSP	type	R723CHPI R723CTRR	Sum(R723CHPI) / Sum(R723CTRR)
ESO-hiperspace read miss rate	HSPMISS	type	R723CCRM R723CTRR	Sum(R723CCRM) / Sum(R723CTRR)
Shared storage page-in rate from auxiliary storage	SHARED	type	R723CSPA R723CTRR	Sum(R723CSPA) / Sum(R723CTRR)
Shared storage page-in rate from expanded storage	EXPSHR	type	R723CSPE R723CTRR	Sum(R723CSPE) / Sum(R723CTRR)
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) / Interval
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

Figure 17-12 (Page 3 of 3). Workload Activity (Goal Mode) - Conditions Based on SMF Record Type 72-3				
Condition	**Condition Name**	**Qualifier**	**Source**	**Algorithm**
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 *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
I/O Using %	IOUSGP	period	R723CIOU R723CTSA	Sum(R723CIOU) / Sum(R723CTSA) * 100
I/O Delay %	IODLYP	period	R723CIOD R723CTSA	Sum(R723CIOD) / Sum(R723CTSA) * 100
Swap-in delay %	SWINP	period	R723CSWI R723CTSA	Sum(R723CSWI) / Sum(R723CTSA) * 100
MPL delay %	MPLP	period	R723CMPL R723CTSA	Sum(R723CMPL) / Sum(R723CTSA) * 100
Queue %	QUEUEP	period	R723CQ R723CTSA	Sum(R723CQ) / Sum(R723CTSA) * 100
Capping %	CAPP	period	R723CCCA R723CTSA	Sum(R723CCCA) / Sum(R723CTSA) * 100
Storage %	STOP	period	R723CAPR R723CACO R723CAXM R723CVIO R723CHSP R723CCHS R723CTSA	Sum(R723CAPR + R723CACO + R723CAXM + R723CVIO + R723CHSP + R723CCHS) / Sum(R723CTSA) * 100
Server delay %	SERVP	period	R723CSPV R723CSVI R723CSHS R723CSMP R723CSSW R723CTSA	Sum(R723CSPV + R723CSVI + R723CSHS + R723CSMP + R723CSSW) / Sum(R723CTSA) * 100
Idle %	IDLEP	period	R723CIDL R723CTSA	Sum(R723CIDL) / Sum(R723CTSA) * 100
Unknown %	UNKP	period	R723CUNK R723CTSA	Sum(R723CUNK) / Sum(R723CTSA) * 100

Channel Path Activity - SMF Record Type 73

One qualifier is possible:

cpid A two-digit hexadecimal number that identifies a channel path.

For overview processing, the qualifier is required. If it is omitted for exception reporting, the threshold applies to all channel paths in the SMF record.

Figure 17-13. Channel Path Activity - Conditions Based on SMF Record Type 73

Condition	Condition Name	Qualifier	Source	Algorithm
Percent channel busy	CHPBSY	cpid	SMF73BSY SMF73SMP	(BSY*100)/SMP

Device Activity - SMF Record Type 74-1

One qualifier is possible:

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, and you must use one of them. If you do not use a qualifier, a syntax error occurs and RMF will not process the condition. If you have selected a shared device in the sysplex, you will receive a value which reflects the sysplex view (not possible for DNOTRDY).

All times reported are in milliseconds, unless otherwise noted.

Figure 17-14 (Page 1 of 2). 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
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

Figure 17-14 (Page 2 of 2). Device Activity - Conditions Based on SMF Record Type 74-1				
Condition	**Condition Name**	**Qualifier**	**Source**	**Algorithm**
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 control unit busy delay time	CUBDL	devnmbr, volser, stg grp, or class	SMF74CUB SMF74MEC	CUB/MEC
Average director port busy delay time	DPBDL	devnmbr, volser, stg grp, or class	SMF74DPB SMF74MEC	DBP/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.

Figure 17-15 (Page 1 of 2). 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	SYNCRT	struct	R744SSRC Interval	Sum(R744SSRC) / Interval	B
Average service time of ASYNC operations	ASYNCST	struct	R744SATM R744SARC	Sum(R744SATM) / Sum(R744SARC)	B
Ended ASYNC operation rate	ASYNCRT	struct	R744SARC Interval	Sum(R744SARC) / Interval	B

Figure 17-15 (Page 2 of 2). Coupling Facility Activity - Conditions Based on SMF Record Type 74-4

Condition	Condition Name	Qualifier	Source	Algorithm	Scope
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 R744SQRC	(Sum(R744FSCC) + Sum(R744SQRC)) / (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	CWRITERT	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(R744SFCN)	B
False contention on lock structure	LCKFCONT	struct	R744SFCN	Sum(R744SFCN)	B

Cache Activity - SMF Record Type 74-5

One qualifier is possible:

ssid SSID number

devn Device number

To define a subsystem-related exception, you specify `SSID(ssid)`
To define a device-related exception, you specify `SSID(ssid),DEVN(devn)`

For exception processing only:

The conditions CASSC, CADSC, and CASSNVS can be used only with the operator EQ, not with LE or GE.

Figure 17-16 (Page 1 of 6). 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
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

Figure 17-16 (Page 2 of 6). Cache Activity - Conditions Based on SMF Record Type 74-5

Condition	Condition Name	Qualifier	Source	Algorithm
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
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

Figure 17-16 (Page 3 of 6). Cache Activity - Conditions Based on SMF Record Type 74-5

Condition	Condition Name	Qualifier	Source	Algorithm
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

Figure 17-16 (Page 4 of 6). Cache Activity - Conditions Based on SMF Record Type 74-5

Condition	Condition Name	Qualifier	Source	Algorithm
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
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

Figure 17-16 (Page 5 of 6). Cache Activity - Conditions Based on SMF Record Type 74-5

Condition	Condition Name	Qualifier	Source	Algorithm
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
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

Figure 17-16 (Page 6 of 6). Cache Activity - Conditions Based on SMF Record Type 74-5

Condition	Condition Name	Qualifier	Source	Algorithm
Device Overview: I/O RATE	CADT	DEVN(devn)	R745CINT TOTAL = R745DRCR + R745DRSR + R745DRNR + R745DWRC + R745DWSR + R745DWNR + R745DICL + R745DBCR	TOTAL / R745CINT
Device Overview: DASD I/O RATE STAGE	CADSTG	DEVN(devn)	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 - DFW_BYPASS) / R745CINT

Figure 17-17 (Page 1 of 2). 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 + R745DWNR + R745DICL + R745DBCR COTOT = TOTAL for CACHE-OFF ALLTOT = TOTAL for ALL I/Os	COTOT * 100 / ALLTOT
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

Figure 17-17 (Page 2 of 2). Cache Activity - Conditions Based on SMF Record Type 74-5

Condition	Categorie	Condition Name	Qualifier	Source	Algorithm
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 - DFW_BYPASS) / R745CINT

Page/Swap Data Set Activity - SMF Record Type 75

One qualifier is possible:

pagename The name of a page or swap data set.

For exception processing, the data set name is optional. If this qualifier is omitted, the threshold value applies to each page or swap data set.

Figure 17-18. Page/Swap 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
Average service time per SSCH (swap data sets only)	PSAVG	pagename	SMF75REQ SMF75SAM SMF75SIO SMF75INT	((REQ*INT)/SAM)/SIO
Page transfer time (page data sets only)	PSPTT	pagename	SMF75REQ SMF75SAM SMF75PGX SMF75INT	((REQ*INT)/SAM)/PGX
Pages transferred per second	PSPT	pagename	SMF75INT SMF75PGX	PGX/INT
Page data set activity rate per SECOND	PSART	pagename	SMF75INT SMF75SIO	SIO/INT
Average slots used	PSAVGSL	pagename	SMF75AVU	Value or comparison
Number of bad slots	PSBADS	pagename	SMF75BDS	Value or comparison
Average swap sets used	PSAVGSE	pagename	SMF75AVU	Value or comparison
Bad swap sets	PSBADSE	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.

Figure 17-19. 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
Average contention time	ENQAVG	major/minor	SMF77WTT SMF77EVT	WTT/EVT
Maximum contention time	ENQMAX	major/minor	SMF77WTX	Value or comparison

I/O Queuing Activity (4381 Processor only) - SMF Record Type 78-1

One qualifier is possible:

lcuid A four-digit hexadecimal number that identifies a logical control unit. For exception processing, the qualifier is optional. If it is omitted, the threshold applies to all logical control units in the SMF record.

Figure 17-20. I/O Queuing Activity (4381 Processor only) - Conditions Based on SMF Record Type 78-1

Condition	Condition Name	Qualifier	Source	Algorithm
Logical control unit activity rate	IOART	lcuid	R781SIS R781INT	SIS/NT
Average logical control unit queue length	IOQLEN	lcuid	R781QUE R781TIS	QUETIS
Percentage of requests deferred	IOD	lcuid	R781TIS R781SIS	100*((TIS-SIS)/TIS)
Percentage of requests deferred (CU busy)	IOCUB	lcuid	R781CUB R781TIS	100*(CUB/TIS)
Percentage of requests deferred (device busy)	IODB	lcuid	R781DVB R781TIS	100*(DVBTIS)
Percentage all channel paths busy	IOCHPB	lcuid	R781ABY	ABY/SAM

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).

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 8

MIN Applies to exception operator LE, and specifies the sum of each channel path taken, where i represents channel path 0 to channel path 8

CHPID(i) Is calculated as R783PB(i)/R783GSAM.

Figure 17-21. 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 R783IQSM	IQSM/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
Percentage of all channel path busy	IOCHPB	lcuid	R783PB R783GSAM	CHPID(0) BUSY * CHPID(1) BUSY * CHPID(2) BUSY * CHPID(3) BUSY

Virtual Storage Activity - SMF Record Type 78-2

Figure 17-22. Virtual Storage Activity - Conditions Based on SMF Record Type 78-2

Condition	Condition Name	Qualifier	Source	Algorithm
Maximum amount of SQA expansion - into CSA in K bytes	SQAE	none	R782SQEX	Value or comparison
Minimum size of largest free block - CSA in K bytes	CSAFB	none	R782CSLF	Value or comparison
Minimum size of largest free block - SQA in K bytes	SQAFB	none	R782SQLF	Value or comparison

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, PTOD, 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 A.M. to 12 noon for the five-day business week from June 6, 1999 to June 10, 1999. Use the following control statements:

Example

```
DATE(06061999,06101999)
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 A.M. to 6:00 P.M. on June 24, 1999. The duration interval is six hours, causing two duration reports to be produced for each specified activity. Use the following control statements:

Example

```
DATE(06241999,06241999)
DINTV(0600)
REPORTS(CPU,CHAN)
REPORTS(DEVICE(TAPE,DASD,COMM))
RTOD(0600,1800)
NOSUMMARY
```

Sysplex Report

The Postprocessor is to generate sysplex reports. The reporting period runs from 8:00 A.M. to 6:00 P.M. for the five-day business week from June 6, 1999 to June 10, 1999. Use the following control statements:

Example

```
DATE(06061999,06101999)
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 A.M. to 4:00 P.M. for the week beginning June 13, 1999 and ending June 17, 1999:

Example

```
DATE(06131999,06171999)
ETOD(0800,1600)
```

A line in the **Exception report** when the I/O service rate for performance group number 3 is less than or equal to 100 service units per second and the percent device utilization for device 06D8 is greater than or equal to 3.

A **Channel Path Activity report** and a Device Activity report for the DASD device class if both conditions are met:

Example

```
EXCEPT(IORATE(IOSRV(0030),LE,100))
EXCEPT(IORATE(DVUTL(06D8),GE,3))
EXRPTS(IORATE(CHAN,DEVICE(DASD)))
```

A line in the **Exception report** if the percent busy for channel path 01 is greater than or equal to ten percent:

Example

```
EXCEPT((CHPBSY(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 A.M. and 2 P.M. 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 `TOTSRV` specifies the total service units, the qualifier `S.TSOSERV` refers to service class TSOSERV, and suboption `NOSYSTEMS` defines sysplex reporting:

```
OVW(SERVUNIT(TOTSRV(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 P.M. and 10:00 P.M. on May 13, 1999, then you define

Example

```
DATE(05131999,05131999)
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 have seen in the list of the top-20 devices that volumes DATA01 on address 06F3, DB2PRD on 0722, and CICS14 on 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(05131999,05141999)
OVERVIEW(RECORD,REPORT)
OVW(DATA01(CADRT(06F3)))
OVW(DATA01(CADRHN(06F3)))
OVW(DATA01(CADSTG(06F3)))
OVW(DB2PRD(CADRT(0722)))
OVW(DB2PRD(CADRHN(0722)))
OVW(DB2PRD(CADSTG(0722)))
OVW(CICS14(CADRT(0734)))
OVW(CICS14(CADRHN(0734)))
OVW(CICS14(CADSTG(0734)))
```

In this example, the Overview report will provide information about the total I/O rate, the read cache hit rate, and the DASD staging I/O rate.

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.

One is the **RMF Spreadsheet Converter (RMF2SC)**. The files containing this spreadsheet application are shipped along with RMF, and just have to be downloaded to your workstation before you can start. With it, you can handle data from the RMF report displays or from report data sets using techniques familiar to every spreadsheet user. Please refer to the *RMF Programmer's Guide* for more details.

The **RMF Spreadsheet Reporter (RMFPP)** is an enhancement of the Spreadsheet Converter to provide a more usable support for converting Postprocessor listings and Overview records into spreadsheets. In addition, it provides sample macros to help you in presenting and analyzing performance data at a glance.

Performance Monitoring of OS/390 (PM of OS/390) provides an interface between the OS/2 workstation and the OS/390 sysplex that gives you the flexibility to create unique scenarios that monitor the performance of your system. You can collect real-time data in graphic and text mode, combine data from different collection types, or even from different applications, and group resources together.

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 MVS systems in your network that have an APPC or TCP/IP connection configured to your PWS. This way, RMFCS combines the advantages of a single point of control for MVS performance management with a state-of-the-art user front end.

Chapter 18. RMF Spreadsheet Reporter (RMFPP)

Spreadsheet Reporting with RMF

This chapter covers the following topics:

- Overview of RMFPP
- Installing RMFPP
- Preparing data with RMF
- Converting data to spreadsheet format with RMFPP
- Using RMFPP Macros

Overview - About the Spreadsheet Reporter

This chapter introduces the Spreadsheet Reporter which allows you to convert RMF data to spreadsheet format and which provides a practical approach how to use spreadsheet macros for converted reports and Overview records. The Spreadsheet Reporter is an extension of the Spreadsheet Converter and enhances its capability and flexibility. The spreadsheet macros contained in the Spreadsheet Reporter are samples to demonstrate how you can use spreadsheets to process RMF data.

Items discussed in this chapter are:

- Exploiting the Spreadsheet Reporter
- Concepts of the Spreadsheet Reporter
- Functions of the Spreadsheet Reporter
- What is Different to the Spreadsheet Converter

Just as an introduction, here you see a sample report that you can create with the Spreadsheet Reporter, it shows the CPU consumption of a selected set of performance groups for one day.

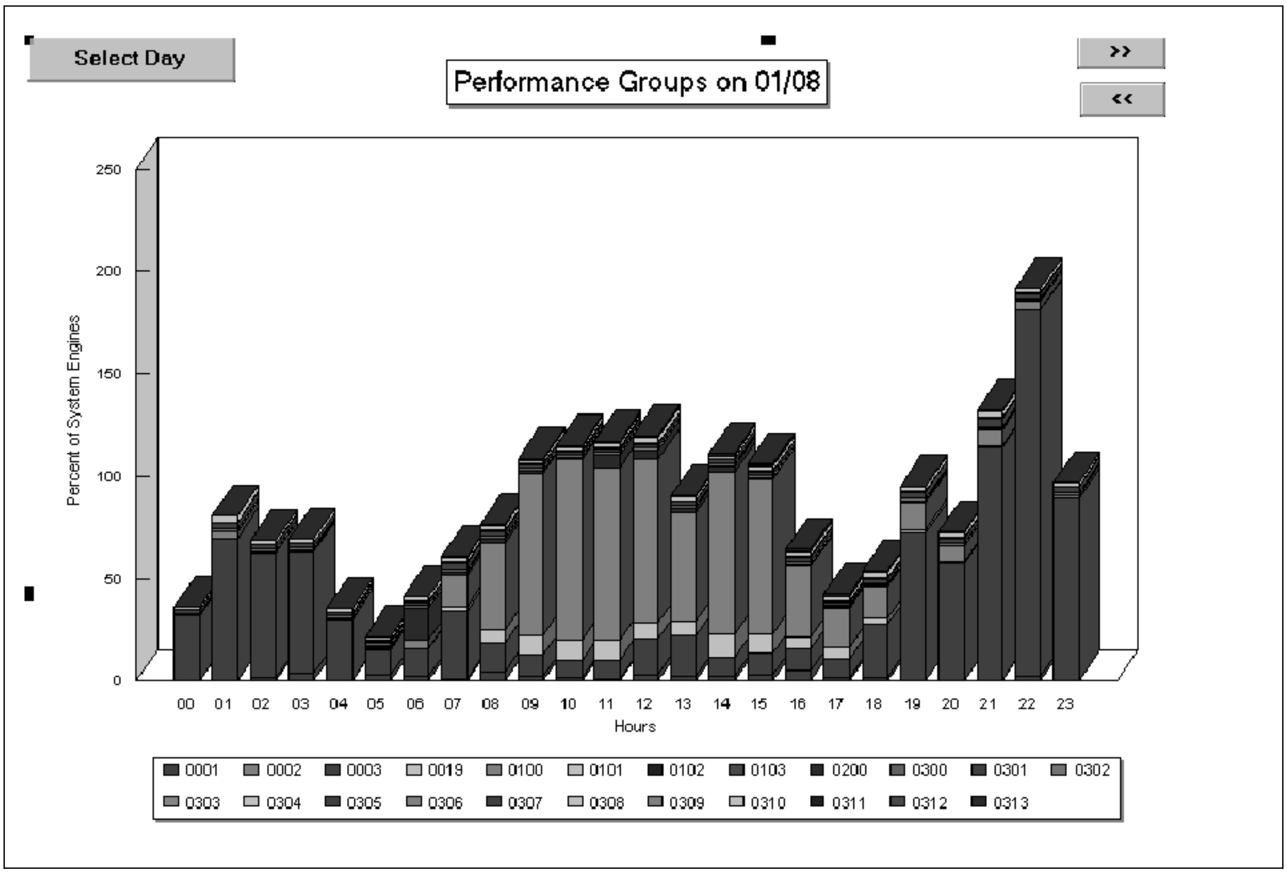

Figure 18-1. Workload Activity Report

Exploiting the Spreadsheet Reporter

Getting More out of Postprocessor Listings

The handling of Postprocessor listings can be quiet difficult. To analyze their content, you have either to print them, use an editor or write a data reduction routine to obtain the most important data out of it. With the introduction of the Spreadsheet Converter, it became possible to convert RMF reports to spreadsheet format, and to use spreadsheet applications for further processing.

While the Spreadsheet Converter can handle single RMF reports, it cannot completely handle a Postprocessor listing data set containing several reports. To get this done was the main reason for the introduction of the Spreadsheet Reporter.

The Spreadsheet Reporter is a set of workstation-based functions to give you a mechanism for converting parts or all of RMF listings to report worksheets in Lotus 1-2-3 WK1 or Excel XLS format.

Becoming Familiar with Overview Records

Probably, Postprocessor reports are the most often used media to analyze RMF SMF data, but they are not the most efficient media. Converting RMF report data to spreadsheet format and using spreadsheet macros to analyze their content, significantly improves the capability of historical reporting on RMF data. Nevertheless, RMF already provides more flexible mechanisms to create reports for historical performance analysis.

This can be done with Overview reports and records. You can generate data tables, similar to the RMF Summary report, with the performance data reflecting your own requirements. The Overview report and record are based on Postprocessor exception control statements, thus providing a huge set of possible data items from which you can choose. Overview records were especially designed to be used for further processing, for example by spreadsheet applications.

Concepts of Overview Records: You can generate Overview reports and records to create

- Installation-specific reports
- Customized Summary reports
- Improved Exception reports

The following sections discuss the possibilities to create installation-specific and extended Summary reports. In *RMF Report Analysis* you find additional examples and information for improved summary and exception reporting.

The **Overview report** has the same layout as the Summary report. By this, the report provides performance data that summarizes system activity for each interval within the reporting period. In deviation to the Summary report, the Overview report does not have a fixed layout. You can define, by using exception control statements, which performance data should be included in the report. This gives you a high flexibility because the Postprocessor supports more than 130 different exceptions.

Overview records are generated in the same way as the Overview report. They are intended to be used by additional applications to process RMF data. RMF provides a record mapping if you want to process the records with your own applications. In addition, the Spreadsheet Converter supports Overview records and creates a spreadsheet file from all records contained in a data set.

Note: The Spreadsheet Converter and the Spreadsheet Reporter **do not convert** the Overview report. The intention of the Overview report is to be used as a hardcopy report. You should use Overview records to generate input for spreadsheet applications.

Concepts of the Spreadsheet Reporter

Working with Listing Data Sets

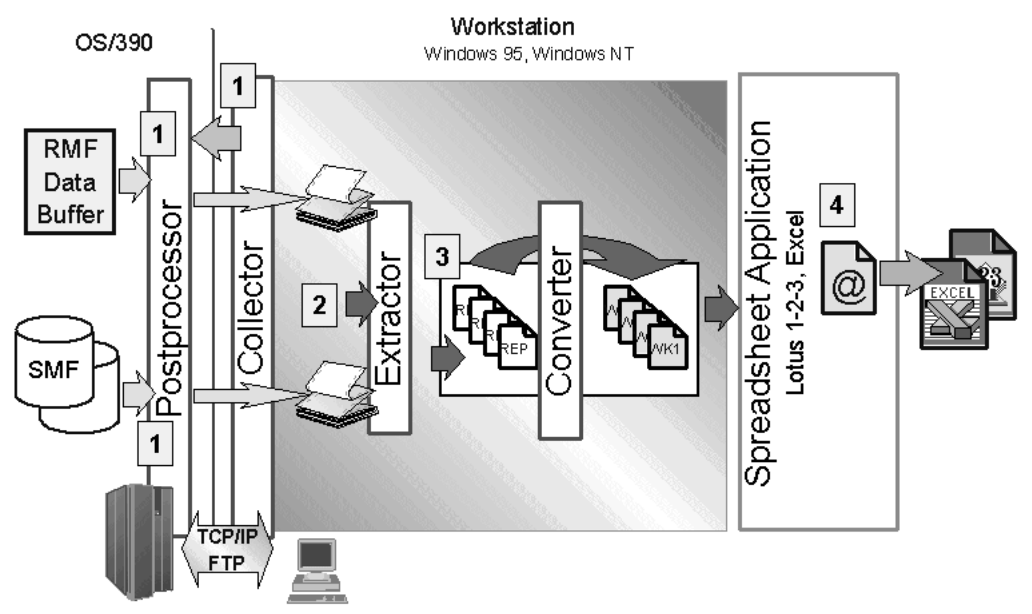

Figure 18-2. Spreadsheet Reporter - Working with Listing Data Sets

You have to perform the following steps from creating an RMF listing with the Postprocessor until you can print and present the results created with your spreadsheet application:

1. Use the **Collector** function to process the RMF SMF data on the OS/390 system with the Postprocessor, and to download the listings to your workstation.
2. Use the **Extractor** function to create a Report-Work-Set based on the downloaded listing data set.
3. Use the **Converter** function to convert the reports to spreadsheet format.
4. Use spreadsheet macros to process the converted data and to create analysis summaries and graphics which can be printed and used for presentations.

Report-Work-Set: The functions of the Spreadsheet Reporter ensure that you can use the data of your listing consistent across all processing steps. In order to guarantee this consistent data access, the Spreadsheet Reporter introduces a new element, the **Report-Work-Set**. A Report-Work-Set is the collection of RMF reports which have been extracted from an RMF listing plus the converted RMF reports in spreadsheet format. Each Report-Work-Set is identified by a **description** or **name** which you specify on the Extractor dialog. Each RMF report contained in a Report-Work-Set is identified by its

- SMF or sysplex id
- Creation date
- Creation time
- RMF report type

The Extractor function creates an index containing the above information for all reports which have been extracted from the source listing. This index is used by the Converter to display the content of a Report-Work-Set, and it allows you to select the reports you want to convert to spreadsheet format.

The Converter function creates indexes for all converted reports. These indexes are written as files in Lotus WK1 or Excel XLS format, and used by the spreadsheet macros to access and process the converted reports. Whenever a spreadsheet macro requests input for data access from you, it displays the information for the Report-Work-Set (its description) and for the reports as described above.

Supported RMF Reports: The Spreadsheet Reporter supports the subset of Postprocessor reports as listed in the following table.

Supported Reports	**Postprocessor Options**
Cache Subsystem Activity Report	REPORTS(CACHE)
Channel Activity Report	REPORTS(CHAN)
Coupling Facility Activity Report	SYSRPTS(CF)
CPU Activity Report	REPORTS(CPU)
DASD Activity Report	REPORTS(DEVICE(DASD)) SYSRPTS(SDEVICE(DASD))
I/O Queuing Activity Report	REPORTS(IOQ)
Paging Activity Report	REPORTS(PAGING)
Partition Data Report	REPORTS(CPU)
Summary Report	SUMMARY(INT,TOT)
Tape Activity Report	REPORTS(DEVICE(TAPE))
Workload Activity Report	REPORTS(WKLD) SYSRPTS(WLMGL)

All of the above listed single-system reports (REPORTS) will also be generated when you specify REPORTS(ALL) and SUMMARY(INT,TOT). All syplex reports (SYSRPTS) need to be requested by the described parameters.

Working with Overview Records

If you plan to use Overview records as base for the spreadsheet reporting, you have to perform similar steps than described in Working with Listing Data Sets:

1. Use the **Collector** function to process the RMF SMF data on the OS/390 system with the Postprocessor, and to download the listings to your workstation.

2. Use the **RecConvert** function to convert the records to spreadsheet format.

Note: There is no Report-Work-Set for Overview records.

3. Use the spreadsheet macros to process the converted data and to create analysis summaries and graphics which can be printed and used for presentations.

Nearly all Postprocessor reports provide data for Overview records. You can either see a list of data based on the SMF record type (see "Overview and Exception Conditions" on page 17-39), or you find a table of available data for each Postprocessor report in *RMF Report Analysis*.

Analyzing Data with Sample Macros

In addition, the application contains sample spreadsheet macros to automatically generate graphical reports and to analyze your report data. The macros use data that you have converted from Postprocessor listing or from Overview records. These macros are provided to demonstrate new, more enhanced ways of using RMF data for reporting, analysis and planning.

DISCLAIMER OF WARRANTIES

All spreadsheet macros are sample code created by IBM Corporation. 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.

The code is provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of such sample code, even if you have been advised of the possibility of such damage.

Functions of the Spreadsheet Reporter

The installation procedure creates a program group named *RMF Spreadsheet Reporter* on your desktop. You can start the Spreadsheet Reporter by clicking on the ***launch*** object in the program group. The LaunchPad starts and gives you access to all functions to process RMF listings and records:

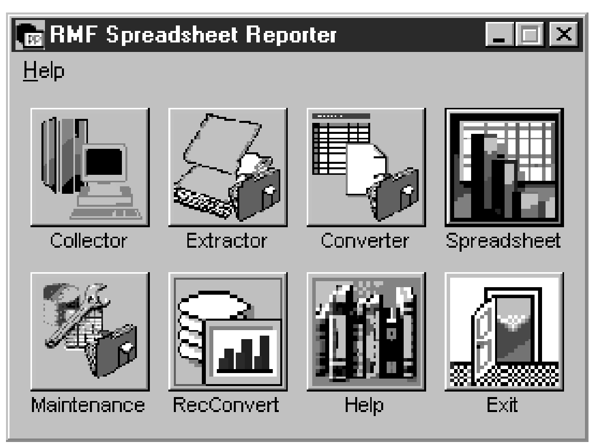

The main functions of the LaunchPad to process RMF listings and records are:

Collector Prepare a Postprocessor job, submit it to the OS/390 host systems, and store the results (listings and records) on the workstation.

Extractor Extract reports from RMF listings and create a Report-Work-Set on your workstation.

Converter View the content of a Report-Work-Set and convert reports to spreadsheet format.

RecConvert Convert RMF Overview records to spreadsheet format.

Spreadsheet Open your spreadsheet application and process converted reports and records with spreadsheet macros.

In addition, you find the following functions which complement the capabilities of the Spreadsheet Reporter:

Maintenance Rename and remove Report-Work-Sets.

Help Read the Spreadsheet Reporter User's Guide.

Exit Close the LaunchPad.

Besides launching the different functions, the LaunchPad guarantees that the functions are synchronized. Creating an Report-Work-Set and converting its reports requires to maintain several indexes to ensure that subsequent functions can view the content of the Report-Work-Set correctly. Therefore, the LaunchPad only allows you to start the next function when you have completed the preceding function. If you try to start a function while another function is active that modifies Report-Work-Sets, the function will not start. The following functions will disable the LaunchPad when they are active:

- Extractor
- Converter
- RecConvert
- Maintenance

Disabling the LaunchPad means that no other function can be started. Opening the spreadsheet macro folder or opening the User's Guide (Help) do not disable the other functions.

Maintenance

Creating an Report-Work-Set requires that you can rename and remove it at any time. This is exactly the purpose of the Maintenance dialog which appears like a *Setting* notebook with three pages:

1. The **Info** page shows the level of the RMF Spreadsheet Converter you are using, and the RMF version, release and PTF number for which the Spreadsheet Converter became available. The Date entry tells you when the Spreadsheet Converter has been packaged with the RMF Spreadsheet Reporter function.

   Note: The Spreadsheet Reporter uses the Spreadsheet Converter to convert RMF reports. The module for the Spreadsheet Converter can be replaced at any time without reinstalling the complete application. In addition, the Spreadsheet Converter and the Spreadsheet Reporter can be used in general for all RMF releases starting with RMF 5.1.

2. The **Rename** page allows you to rename a Report-Work-Set. The information displayed on the page is the same as in the **Select Report-Work-Set** group-box of the Converter function.
3. On the **Delete** page, you can remove a Report-Work-Set from your workstation. You should **always** use this function to delete Report-Work-Sets, because it is the only way that guarantees that the Spreadsheet Reporter function and spreadsheet macros will work properly.

Help

There are several ways to obtain assistance while using the Spreadsheet Reporter:

- On the **LaunchPad**, you find a **Help** button. This button starts a view facility to browse documentation (files with .INF and .HLP extension).

  Selecting .INF files leads you to the **User's Guide (UsersGuide.INF)**.
- All functions provide online help which is sensitive to the current focus on the dialog. You can start online help
  - by clicking on **Help** ,
  - by selecting **general help** from the Help item in the Menu bar,
  - by clicking on **F1** when you work with the dialog.

What is Different to the Spreadsheet Converter

There are two functions that you can use for preparing your data for spreadsheet processing:

RMF2SC RMF Spreadsheet Converter

RMFPP RMF Spreadsheet Reporter

This raises the

Questions

- When should I use the Spreadsheet Converter?
- When should I use the Spreadsheet Reporter?

Both functions have different software environments on the workstation, and both functions provide different capabilities.

Software Environment

Spreadsheet Converter

- OS/2 Version 2.0 or higher
- Any DOS Version 5.0 or higher with Windows** 3.0 or higher

Spreadsheet Reporter

- Windows 95
- Windows NT Version 4

Depending on the operating system on your workstation, you see which function is available for you.

Note: For OS/2, the Spreadsheet Reporter is available as tool in the Internet. Please, refer to “Prerequisites” on page 18-10 for details.

Capability

Spreadsheet Converter

You can convert Postprocessor reports as well as data from interactive Monitor II and Monitor III sessions. Please, refer to the *RMF Programmer's Guide* for details.

Spreadsheet Reporter

This function includes the capability of the Spreadsheet Converter, but for Postprocessor data only, certainly with significantly improved usability. Conversion is much more easier if you want to convert several Postprocessor reports in one step.

In addition, powerful spreadsheet macros are available to display the converted data.

Installing the Spreadsheet Reporter

The Spreadsheet Reporter is installed on the host along with the rest of the MVS components of RMF. The deliverable includes

- Spreadsheet Reporter application files
- Spreadsheet Converter application files
- Spreadsheet Macros for Lotus 1-2-3 for Windows Version 5 (US Version)
- Spreadsheet Macros for Microsoft** Excel Version 5
- Spreadsheet Macros for Microsoft Excel 97
- A sample RMF listing data set

You start the installation by downloading and extracting the files on your workstation. Use the installation procedure described below.

Note

The installation of RMFPP includes the installation of the Spreadsheet Converter (RMF2SC) — it not required to do this separately.

The installation procedure creates the following directories on your workstation:

`d:\rmfpp`	Main directory
`d:\rmfpp\Progs`	Directory for Spreadsheet Reporter application files
`d:\rmfpp\Rmf2sc`	Directory for RMF Spreadsheet Converter application files
`d:\rmfpp\Listing`	Default directory for RMF listing data sets
`d:\rmfpp\Overview`	Default directory for RMF overview data sets
`d:\rmfpp\Macros`	Main directory spreadsheet macros (only if at least one set of spreadsheet macros has been installed)
`d:\rmfpp\Macros\L123`	Directory for Lotus 1-2-3 spreadsheet macros (only if Lotus 1-2-3 spreadsheet macros have been installed)

`d:\rmfpp\Macros\Xcel`	Directory for Microsoft Excel V5/V7 spreadsheet macros (only if Excel spreadsheet macros have been installed)
`d:\rmfpp\Macros\Xcel97`	Directory for Microsoft Excel 97 spreadsheet macros (only if Excel spreadsheet macros have been installed)

`d:` is the drive letter and `rmfpp` the path on your workstation where you have installed the Spreadsheet Reporter. The installation also creates a program group (default name *Spreadsheet Reporter*) with a program icon *launch*.

In addition, some directories will be defined for sample data.

Prerequisites

Hardware Requirements

- A workstation with an i486** processor (66MHz), and at least 16MB of memory
- Disk space required is 10MB

  However, to use the Spreadsheet Reporter efficiently, we recommend that you install the application on a disk with at least 30MB of free space. This is because of RMF listings which have to be downloaded to your workstation and which can be quiet large.

Software Requirements

- Operating System
  - Windows 95
  - Windows NT Version 4
  - For earlier versions of Windows NT (3.5 or 3.51) please read “Known Problems” on page 18-41

  Note: A down-level version of the Spreadsheet Reporter to run on OS/2 is available in the Internet. You can get access to it from the RMF home page:

  http://www.ibm.com/s390/rmf/

  You find information about the Spreadsheet Reporter in the **Tools** page, which you can access from the home page directly.

- Spreadsheet Program

  You can use any spreadsheet program that can read Lotus files with an extension of **.WK1**. However, in order to use the spreadsheet macros shipped with the Spreadsheet Reporter one of the following products is required:
  - Lotus 1-2-3 Version 5 (International English Version)
  - Microsoft Excel Version 5.0c or Version 7
  - Microsoft Excel 97

Installation Steps

The installation on Windows 95 and Windows NT is done with InstallShield**. InstallShield provides an interactive dialog which guides you through the installation process.

The code of RMFPP (in the self-extracting ZIP file **ERB9R2SW.EXE**) is distributed as member ERB9R2SW of the SERBPWSV distribution library. Ask your system programmer for the name of the library in which this file is stored, and install it on your PWS by following this procedure:

1. Download ERB9R2SW.EXE to this directory. Use the command

   ```
   RECEIVE ERB9R2SW.EXE h:'hlq.SERBPWSV(ERB9R2SW)'
   ```

   where `h:` is your host session name, and `hlq` is the high-level qualifier of the MVS distribution library. Make sure your MVS host session is in TSO Ready mode.

   Of course, you can use any other program you know to perform the download.

2. Now, you can start the installation by executing the installation file:

   ```
   ERB9R2SW
   ```

   You can also use another suitable program, for example

   a. Open the Windows Explorer.
   b. Open the folder which contains the installation file.
   c. Double-click on installation file symbol in the *Contents of ...* panel.

3. InstallShield prepares the installation setup dialog which guides you through the installation process. The installation process displays some information panels which you should read carefully. You can go forward and backward in the installation process by clicking on **NEXT** and **BACK**. The process requests information from you to specify the installation directory, the installation type, and the name of the Spreadsheet Reporter program group.

   You will be asked on which drive and path you want to install the Spreadsheet Reporter. The default is

   ```
   <WININST>\Rmfpp
   ```

   `<WININST>` is the drive where your Windows system has been installed (e.g. C:). When you select a new drive, please make sure that you also specify a directory name.

   It is **highly recommended** to use `Rmfpp` as directory name to avoid problems in the further processing. In any case, a blank character is not allowed as part of the path name as it is supported in general by Windows.

4. The installation offers you three types:

   Typical Install all application files and all spreadsheets macros for Microsoft Excel. This is the default (11 MB required).

   Compact Install only the application files, without any spreadsheet macros (7 MB required)

   Custom You choose whether you want to install

   - the application files
   - the Lotus 1-2-3 spreadsheet macros
   - the Excel spreadsheet macros

   You should use **Custom** installation for two reasons:

   - If you only want to use one spreadsheet application, you can deselect the spreadsheet application you do not want to use.
   - If you want to reinstall spreadsheet macros for any of your spreadsheet applications, you can deselect the application files and the spreadsheet macros you do not want to reinstall.

Please ensure that you keep the same drive and path for the macro reinstallation.

5. The installation routine asks for the name of the Spreadsheet Reporter program group. The default name is *RMF Spreadsheet Reporter*.
6. The installation routine will inform you when the installation process has been completed. Then, you exit the installation by clicking on **Finish**.
7. Now, you can erase ERB9R2SW.EXE if you want to save space on your disk, you do not need it anymore.

After having installed the Spreadsheet Reporter, you can immediately start it by double-clicking on the launch icon in the RMF Spreadsheet Reporter program group.

Preparing Host Data for Conversion

The easiest way to prepare the host data for conversion is to use the Collector function. Of course, if you prefer for any reason to create the reports and records directly on the host and to download them manually to the workstation, this is possible, too.

Collecting Data on the Workstation

You can use the Collector function

- to specify all parameters for the Postprocessor job
- to submit the Postprocessor job to run on a OS/390 host system
- to receive the results (either a listing data set, or Overview records)

Collector

The Collector function has several notebook pages to specify and to store the parameters you need for running a Postprocessor job.

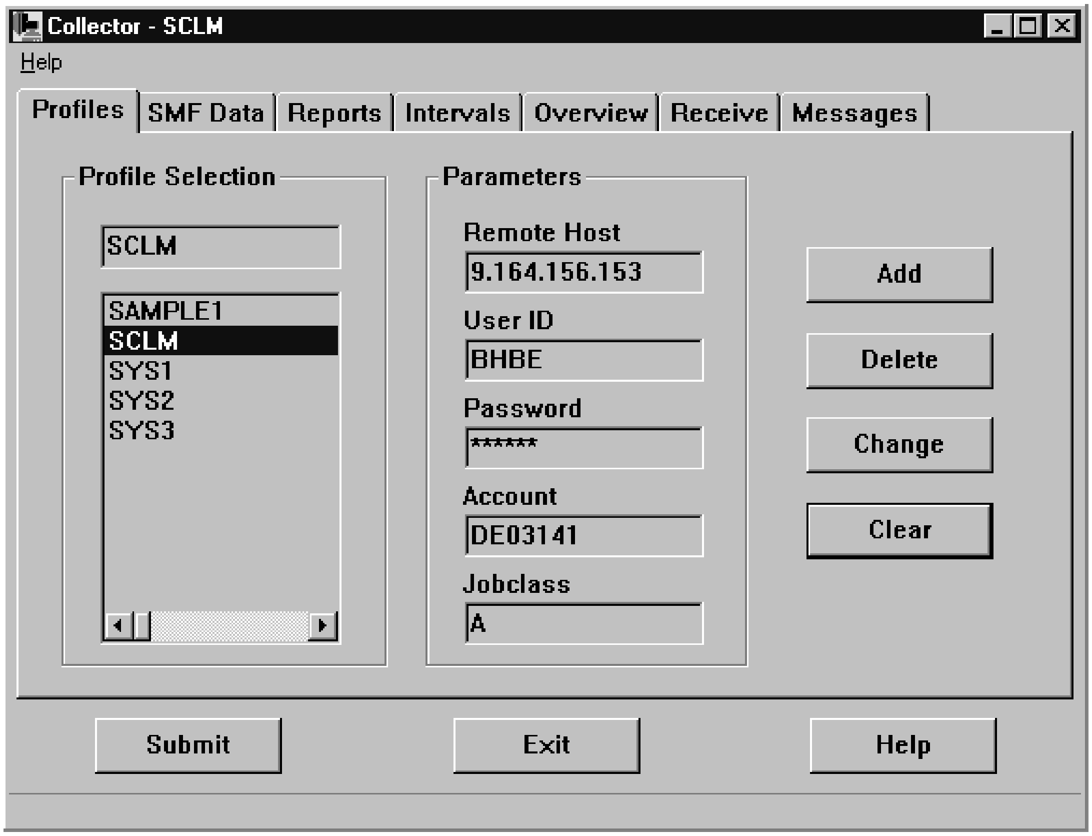

In the first panel, you define your profile. For each remote system that you have defined in the Profile Selection list, it includes the following parameters:

Remote Host Hostname or IP address

User ID Your TSO user id on that system

Password Your logon password - if you specify an invalid or expired password, your will receive a message that the JCL failed.

Account A valid account number

Jobclass The jobclass for the job to run in

These parameters will be used together with a JCL skeleton to create the actual JCL for the Postprocessor job. If you need additional parameters or records (for example, a //STEPLIB record), you can modify the skeleton according to your requirements. It is stored in

```
d:\rmfpp\progs\rmfpp.jcl
```

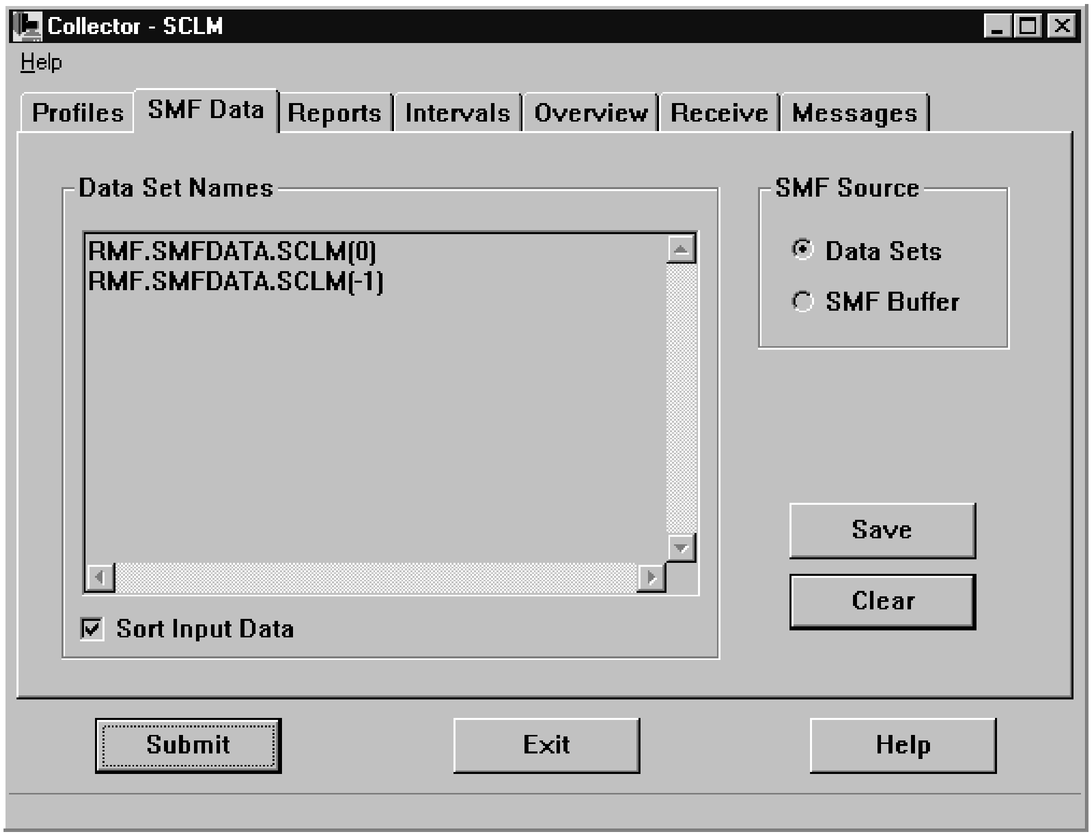

If you want to process SMF records that are stored in data sets, you can specify the names (fully qualified, without quotes) in the Data Set Names list. This requires **Data Sets** in the SMF Source selection. If you specify **SMF Buffer**, the data sets will be ignored, and the Postprocessor takes the current SMF buffer of the system where the job will be running. In addition, you can specify whether the SMF data should be sorted.

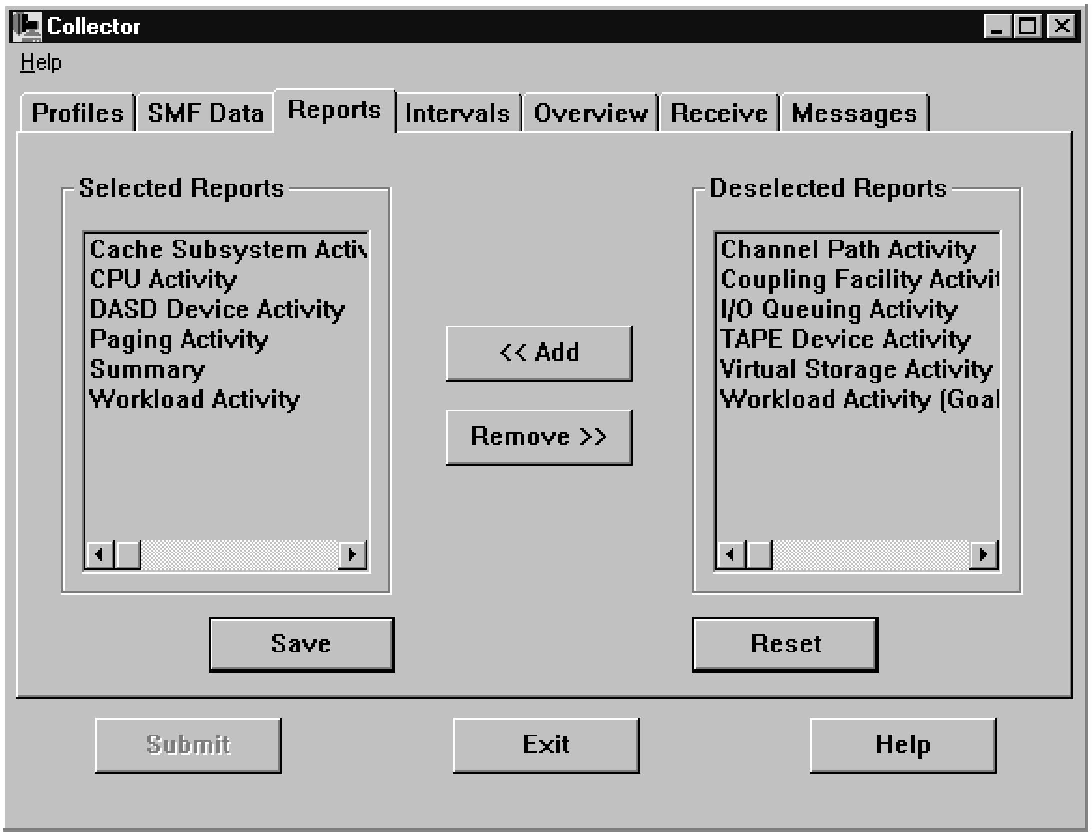

In the **Reports** panel, you can select the report types that you want to be created. By using the buttons **Add** and **Remove**, you can tailor your list of Selected Reports.

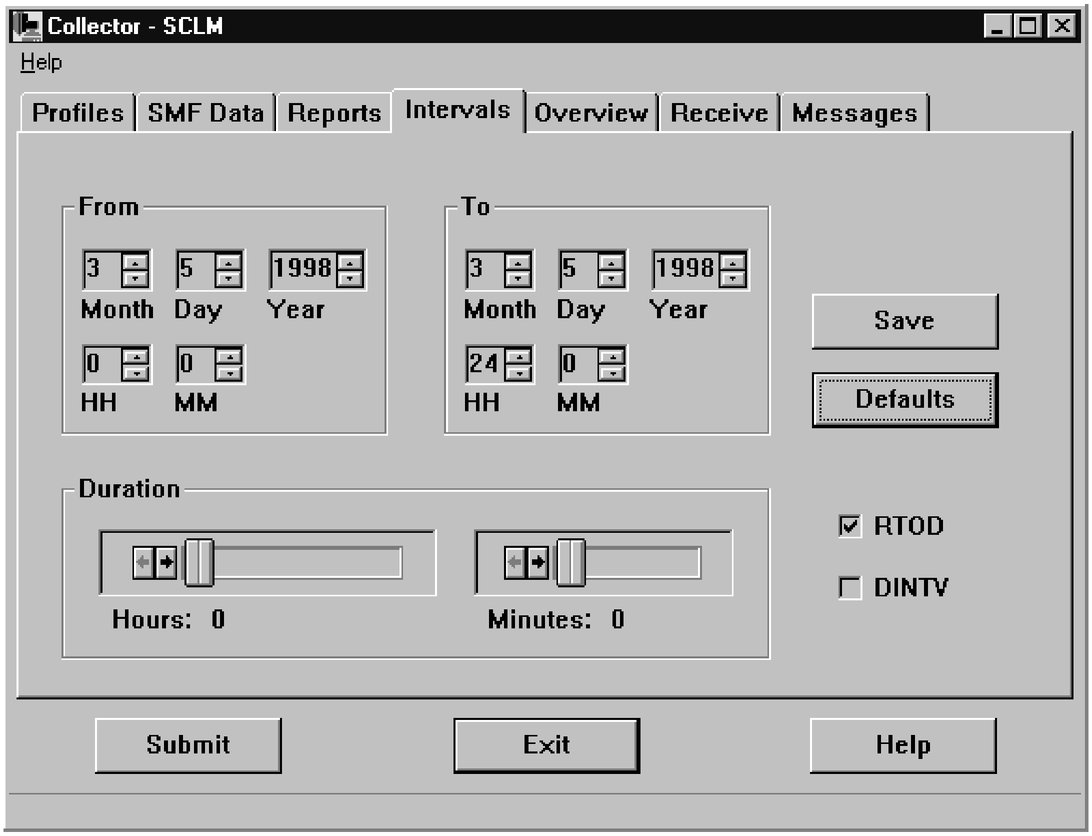

The Interval panel is the place to define the report interval. You have to select **RTOD** if you want to activate the selected interval, otherwise all SMF records found in the input will be used as source for the reports. By selecting **DINTV** you can define that duration reports will be created.

If you want to create Overview records (or reports), you can define the control statements in the Overview Control Statements list. You can either enter them manually, or you can use **Import** to insert a file with statements, for example, those which have been generated automatically by a spreadsheet macro (see page 18-19).

The creation of Overview records or reports is exclusive to the creation of other reports. Therefore, if you select **Record** or **Report** in this panel, all selections in the **Reports** panel will be ignored.

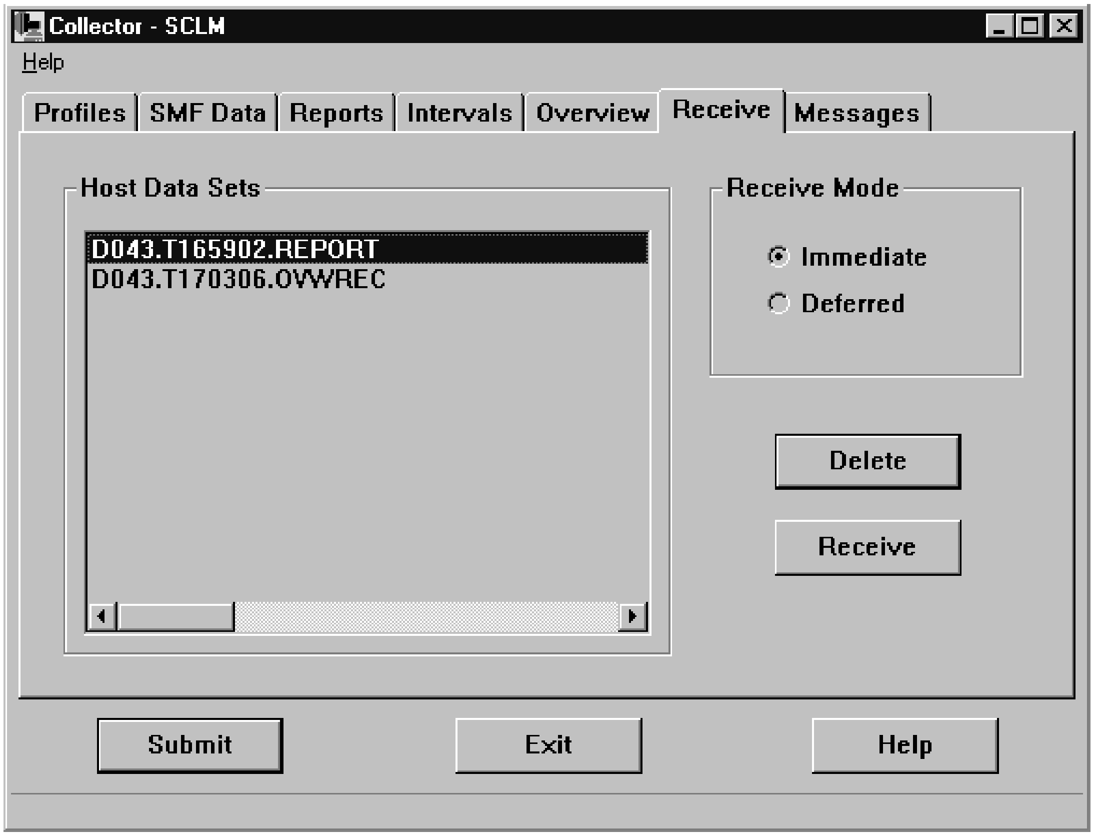

You can receive the results of the Postprocessor job either immediately, or you can specify that they should be stored in a data set for a later transfer to the workstation. This will be defined by selecting **Deferred**. After the completion of the job, you will see the name of the data set which contains the results. You can transfer the data set with **Receive** , this function also deletes the host data set. With **Delete** you can just remove an entry from the Host Data Sets list.

Done - Submit

Now, everything should be prepared for submitting the job to the selected host system. Independent of the notebook page that you have currently open, you click on **Submit** , and the file transfer to the host system will be started.

Note: If **Submit** is not enabled for processing, please check whether you have selected a system on the **Profiles** page.

Hopefully, everything works correctly with the job, and you get the message that the data have been stored on the workstation (in the path `d:\rmfpp\`). You can see all messages being generated by the job in the next panel, and you can select between JES messages and RMF messages.

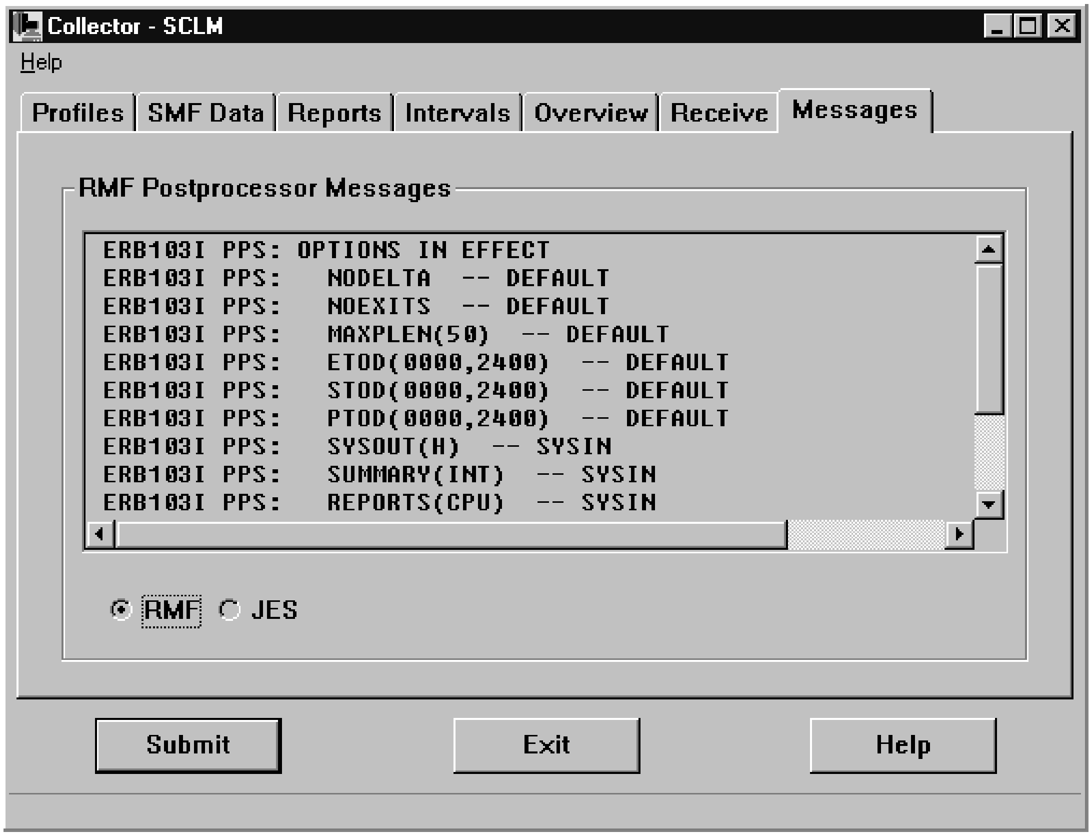

If the job was successful, you can use the results either as input for the Extractor (in case of listings) or for the RecConvert function (in case of Overview records).

And if the jobs was not successful - what can you do?

These are some points you might check to find the reason for the problem:

- If your job did run on the host system, but not successfully, you might check the output directly by logging on to the system (verify correct SMF data sets or sufficient RACF authorization to access the SMF data buffer)
- You can find the JCL that was submitted in the data set `userid.RMFPP.JES`. In case that you cannot find the job (or the JCL), it might have happened that the job could not be sent to the host system. Please, verify your profile (correct userid and host password), and check the TCP/IP connection between the workstation and the host system using the PING command:

  ```
  ping hostname
  ```

- If you get an error message from FTP, you might check the error log `ftperr.log` (in the path `d:\rmfpp\progs\`).

Manual Downloading

Of course, you can create the RMF listings and Overview records directly by running the Postprocessor job on the host system, and then you download the data to the workstation. You have to perform the following steps:

1. Process the SMF data with the Postprocessor and save the results into a data set (RMF listing or Overview records), see “Starting the Postprocessor” on page 17-4 for details.
2. Download the data set to your workstation.
   - Specify **text format or ASCII conversion** for the RMF listing. You may use the default directory for listings on your workstation (`d:\rmfpp\Listing`).
   - Specify **binary conversion** for Overview records. You may use the default directory for records on your workstation (`d:\rmfpp\Overview`).

Generating Overview Records

Overview records will be generated by a set of overview control statements OVW or exception control statements EXCEPT together with the option OVERVIEW(RECORD). Each job control to generate Overview records must contain:

- A PPOVWREC DD statement to define the output data set for the Overview records.
- The Postprocessor option OVERVIEW(RECORD)

  If you specify OVERVIEW(RECORD,REPORT), you get in addition a report where you can see the data that you are going to send to the workstation.
- OVW or EXCEPT statements to define the data columns of the output record. The data columns appear in the same order as you define the control statements.

For details about OVW, EXCEPT, and the overview conditions, please refer to “OVW” on page 17-24, “EXCEPT” on page 17-19, and “Overview and Exception Conditions” on page 17-39.

Converting Data to Spreadsheets

Converting Listing Data Sets

Extractor

The Extractor function uses the RMF listing, which has been downloaded to your workstation, to create a **Report-Work-Set**. The dialog allows you to select the RMF source listing and the report types you want to extract from the listing, and it requires to describe (name) the Report-Work-Set:

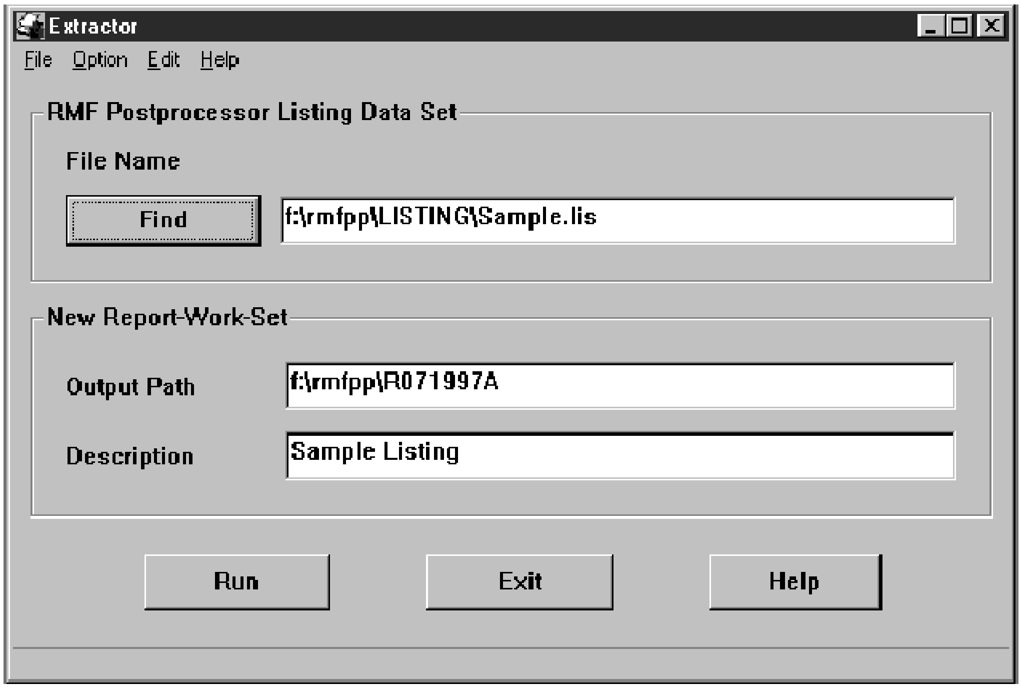

The dialog consists of four areas:

- The **Menu bar** has four entries:

  File allows you to select an RMF source listing. It also saves the last three directories you used to open RMF listings.

  Option allows you to open a dialog on which you can specify the RMF report types to be extracted from the source listing.

  Edit lets you manipulate the **Description** text-box.

  Help lets you open the online help for the extractor.

- The **RMF Postprocessor Listing Dataset** group-box contains a **Find** button and a text-entry field to specify the RMF source listing. **Find** always opens a **File Selection** menu for the default LISTING directory of the Spreadsheet Reporter. If you choose another directory, the path information is saved in the menu.

- The **New Report-Work-Set** group-box displays the name of the directory (it will be created automatically based on the current date and a suffix, for example, `d:\rmfpp\R060299A`, if you run the program on June 2nd, 1999) for the new Report-Work-Set, and lets you describe the new Report-Work-Set. The description is **mandatory**, because it is used by all other functions to identify the Report-Work-Set. With the **Maintenance** function, you can change this description.

- The **Control Button** area contains three buttons:

  Run - to start the Extractor
  Exit - to terminate the Extractor
  Help - to select the online help

  You can only start (**Run** enabled) the Extractor when you have selected an RMF source listing and have entered a description for the Report-Work-Set. After having started the Extractor, you see the progress on the **Create Report-Work-Set** dialog. It provides an overview on the converted reports (description, type, and count), and will be completed by clicking on **Done** .

Usage Notes: The following list describes the behavior of the Extractor dialog for exceptional situations:

- If the RMF source listing does not exist, you will see the **Create Report-Work-Set** dialog and a warning message that the source listing does not exist.

  → No Report-Work-Set can be created.

- If no report is found in the specified RMF listing, you will see the **Create Report-Work-Set** dialog, but no report types being extracted.

  → No Report-Work-Set can be created, because no valid report has been found. This error most often occurs when you download the listing in binary format, or when your installation truncates the first column (containing print-control characters) from the listing.

- If your disk is full, the **Create Report-Work-Set** dialog will issue an error message and will stop the extraction process. The error message shows you the name of the Report-Work-Set directory and asks you to remove this directory manually. The Report-Work-Set has not been completely created, that means you will not find the description listed in the Converter or Maintenance dialog, but some already extracted reports are saved to your disk.

See also “Usage Notes” on page 18-23 for the Converter function.

Converter

The Converter function lists the content of all Report-Work-Sets on your workstation, and you can select the reports you want to convert to spreadsheet format. You can do this multiple times. Newly converted RMF reports will be added to the corresponding Report-Work-Set:

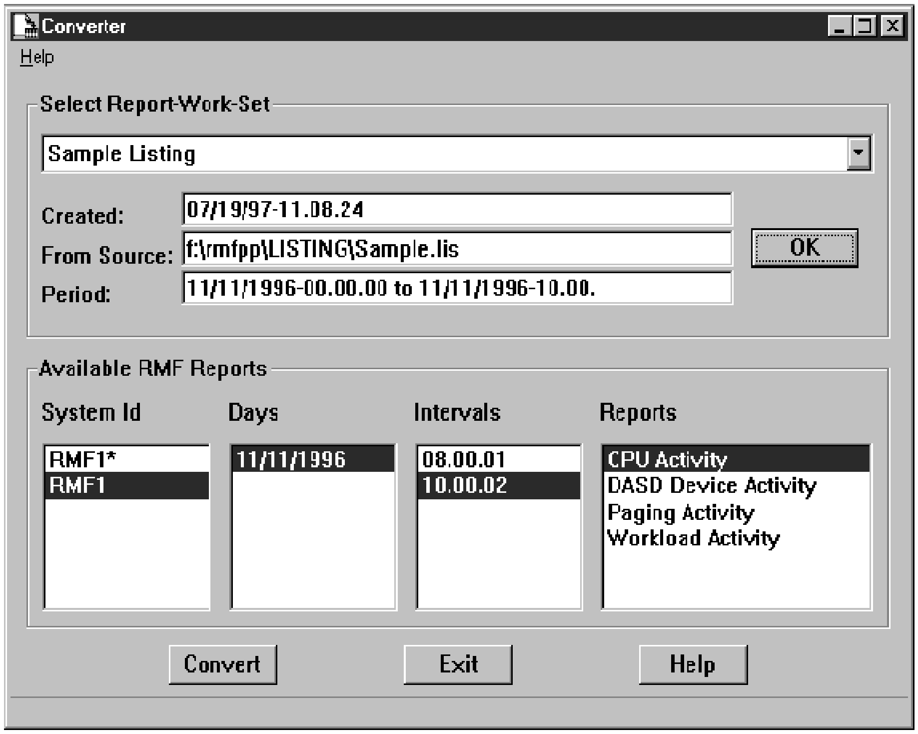

The dialog consists of a **Menu bar** with one entry **Help** to access the online help for the Converter, of three buttons Convert (to start conversion), Exit, and Help, and of two group-boxes:

- In the **Select Report-Work-Set** group-box, you select the Report-Work-Set you want to work with in a scrollable text-box showing the description line of each available Report-Work-Set. Some details (creation date, source listing name, and reporting range contained in the listing) will be displayed to help you making your selection.

  You have to click on OK for the selected Report-Work-Set in order to browse the reports contained in the Report-Work-Set and to select reports which should be converted to spreadsheet format.

- The **Available RMF Reports** group-box displays the SMF-ids or Sysplex-ids of all reports in the currently selected Report-Work-Set. If you make a selection, you get the list of all days, then of all intervals, and then of all report types. You can convert multiple RMF reports with one selection to spreadsheet format.

  Each SMF-id with a trailing star displays the RMF Summary report for this system. This report has been separated because it is not tied to specific interval or duration, like the other RMF reports.

After having started the Converter, you see the progress on the **Conversion Progress** dialog. It provides an overview on the number of converted reports, and will be completed by clicking on Exit.

Usage Notes: The following list describes the behavior of the Converter dialog for exceptional situations:

- When you select a very big Report-Work-Set (more then 20MB) on a small workstation (slow i486, slow drive), you may have to wait a couple of seconds before you see list-boxes filled with entries.

- If there is not enough space on your disk to convert reports to spreadsheet format, you will get an error message. Also the index for the converted report might not be generated. In this case you are no longer able to access the converted reports with the spreadsheet macros, but you have the following option:
  - Remove files from your disk which you no longer need and convert one report for the Report-Work-Set for which the error occurred. Converting one report will recreate all spreadsheet indexes completely.
- If you experience that fewer reports were converted to spreadsheet format as you have selected, most likely one of the following happened:
  - You selected multiple days or intervals for one or multiple report types.
  - The list-boxes always display the content for the top-most selected item of the list-box in the higher hierarchy. That means in case you select two days, you will see the Interval time-stamps in the Interval list-box only for the first selected day.

    Nevertheless, the Converter tries to convert all reports for all selected days and intervals.
  - If the time-stamps on preceding days do not match, the Converter cannot find a corresponding report for the selected interval on the next day, thus it does not convert any report for the missing interval.

    This situation most likely occurs under the following circumstances:
    - You have created a listing manually by appending reports from separate listings.
    - It is possible that the interval start times do not match for RMF interval report for several reasons. This usually does not happen for RMF duration reports, for which you specify the duration period. The described problem is one reason why RMF duration reports are better to be used than RMF interval reports.

See also “Usage Notes” on page 18-22 for the Extractor function.

Converting Overview Records

Overview Record Converter

Converting Overview records is one of the easiest tasks: just by clicking on **RecConvert** , you open the function to convert the records to spreadsheet format.

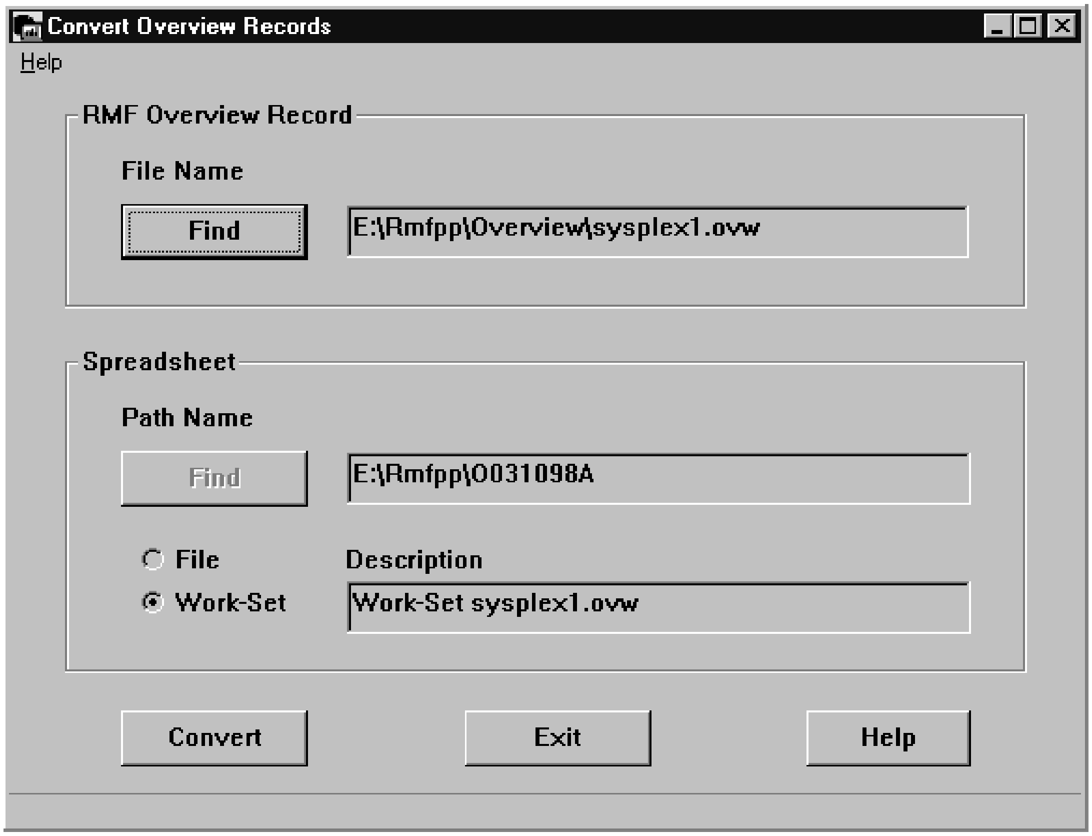

Convert Overview Records consists of two group-boxes:

1. **RMF Overview Record** to select a workstation file with Overview records.
2. **Spreadsheet** to specify the name for the spreadsheet output file in .WK1 format (if you have records only for one system), or you can store the converted records as a Work-Set (if you have records for several systems or a sysplex).

Convert starts the record conversion.

From this process you see that there is no selection of a report type, a date or time frame — all records will be converted into a file with spreadsheet format, and there is no Report-Work-Set. When you use the spreadsheet macros, you refer to the data just by the file name and not, as with listing-based data, to the description of the Report-Work-Set.

Using RMF Spreadsheet Macros

The objective of the Spreadsheet Reporter is to provide a complete solution of enhanced reporting capabilities for RMF reports. The previous chapter discussed how you can convert RMF listing data sets, which consist of RMF reports of multiple intervals, and how you can create and convert RMF Overview records. The purpose of this chapter is to assist you in using RMF spreadsheet macros to process converted RMF reports and Overview records.

Concepts of RMF Spreadsheet Support

The RMF spreadsheet support has been created in the same way for Lotus 1-2-3 and Microsoft Excel. The general behavior of the macros is the same for both spreadsheet applications.

You start the RMF spreadsheet support by clicking on **Spreadsheet** on the Spreadsheet Reporter LaunchPad. A folder opens and displays the **RMF Spreadsheet Management Macro(s)** of your spreadsheet application(s).

In addition, you see additional folder(s) named L123 and XCel which contain all other RMF macros.

It is recommended that you always start the spreadsheet support by double clicking on the spreadsheet management macro. The name of the macro is

RMFSCMN.WK4	Lotus 1-2-3
RMFSCMN.XLS	Excel V5/V7
RMFR9MN.XLS	Excel 97

The purpose of this macro is to give you access to all other RMF spreadsheet macros and to perform operations which guarantee that the macros can access the converted reports. In general, the macros behave very similar like small applications running on your workstation:

- All macros provide **buttons** which you should use to start a function, a processing step, or perform an activity.
- In most cases the spreadsheets provide **dialogs** to ask for your input.
- All macros which process data from Report-Work-Sets display the description of the Report-Work-Set you specified by using the Extractor function. In addition, the reference to all reports is done in the same way as on the Converter dialog, by using the system id, report creation date and report type.
- All macros which process converted RMF reports or Overview records contain sample data. That means you can examine the type of reports the macro provides before you start to process your own data.
- All macros are structured in the same way and provide the same look and feel in order to ease their usage.

Structure of RMF Spreadsheet Support

The following picture depicts the structure of the RMF spreadsheet support. When you double-click on the **RMF Spreadsheet Management Macro**, the **RMF Spreadsheet Main Dialog** is started and the management macro is automatically loaded. This assumes that the spreadsheet application which is required by the macro has been installed on your workstation.

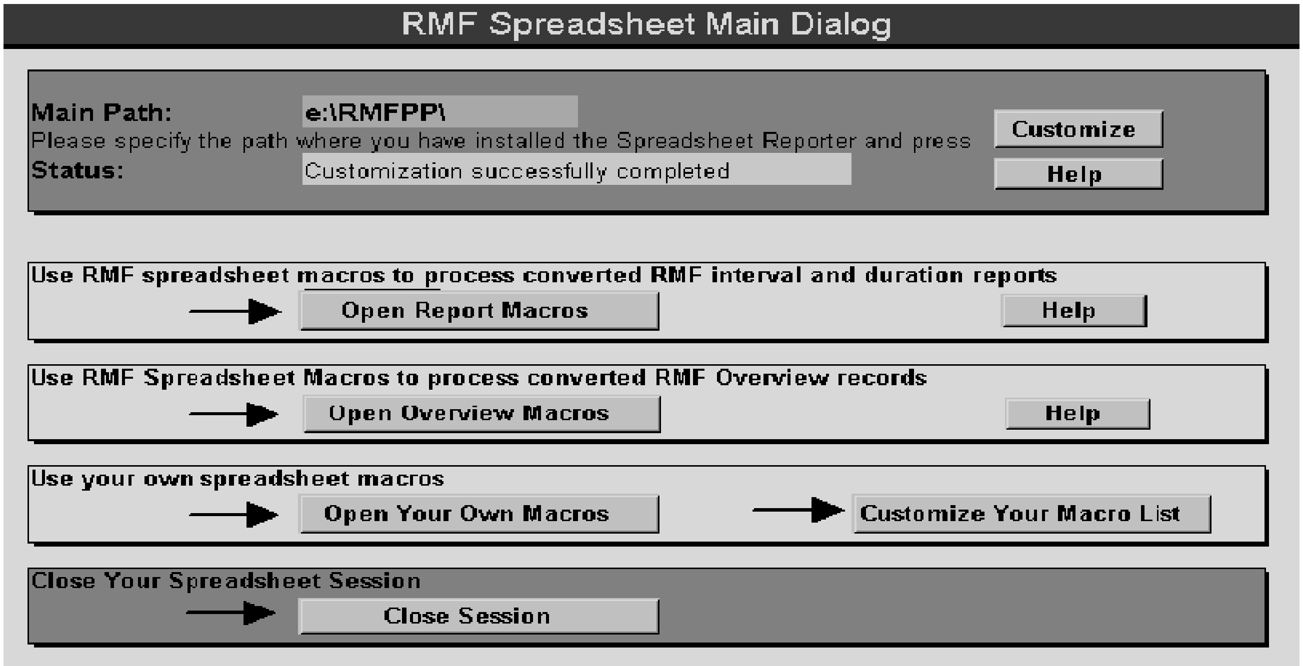

Figure 18-3. RMF Spreadsheet Main Dialog - RMFSCMN

This dialog is the entry point for working with the RMF spreadsheets and offers four functions:

- Customization

  This step has to be performed only once. After having installed the Spreadsheet Reporter, it is necessary to tell all macros which access reports from a Report-Work-Set the drive and path where the Spreadsheet Reporter has been installed. You must enter the correct path information on the Spreadsheet Management Macro, and do not forget to **enter a backslash** at the end of the path definition. Then click on **Customize** to propagate the correct information into the RMF report macros and your User macros which also access the Report-Work-Sets.

  Please, save the macro after a successful customization, otherwise you have to perform this step again when you call the macro the next time.

  Note: This step is not required if you use **RMFR9MN** (Excel 97) which has another layout and another structure.

- Open Report Macros

  Use RMF spreadsheet macros to process converted RMF interval and duration reports from Report-Work-Sets.

- Open Overview Macros

  Use RMF spreadsheet macros to process converted RMF Overview records.

- Open Your Own Macros

  Use your own spreadsheet macros.

All macros provide the capability to create a **Final Report**. This is a macro which contains your RMF data and which is saved with a new name on your workstation. Thus, you can use the data and reports in the spreadsheet at any time and use the original macro to process another RMF report or Overview record.

Available RMF Spreadsheet Macros

Report Macros

Macros which process converted RMF Reports contained in Report-Work-Sets.

Macro	Lotus	Excel V5/V7	Excel 97
Open RMF Report Spreadsheets You can open any converted RMF spreadsheet from a Report-Work-Set.	RMFSCOPN	RMFSCOPN	RMFR9OPN
Summary Report The macro processes a Summary report and creates analysis summaries and graphics from its data.	RMFSCSUM	RMFSCSUM	RMFR9SUM
DASD Activity Report The macro analyzes a converted DASD Activity report and provides summaries for the most heavily used LCUs and DASDs in your installation.	RMFSCDAS	RMFSCDAS (*)	RMFR9DAS (*)
Combined Workload-CPU-Paging Report (Compatibility mode) The macro combines three RMF reports of one reporting interval or duration, calculates performance reports and analyzes your system's behavior.	RMFSCWCP	RMFSCWCP (*)	RMFR9WCP (*)
Workload Activity Trend Report (Goal mode) The macro calculates performance reports and analyzes your system's behavior in goal mode.	—	RMFSCWLM	RMFR9WLM
CF Activity Trend Report The macro provides reports about activities in your coupling facilities.	—	RMFSCCF	RMFR9CF
Cache Subsystem Report The macro calculates provides reports about activities in your cache subsystems.	—	RMFSCCAC (*)	RMFR9CAC (*)
Multi-System Device Report The macro analyzes DASD Activity reports from several systems and provides summaries for the most heavily used LCUs and DASDs in your sysplex.	—	RMFSCMDV	RMFR9MDV
LPAR Trend Report The macro analyzes Partition reports and provides information about the active partitions in your PR/SM environment.	—	RMFSCLP	RMFR9LP

Macro	Lotus	Excel V5/V7	Excel 97
Tape Mount Report This macro displays the tape mounts and the tape activities for one or several systems.	—	RMFSCTAP	RMFR9TAP
Summary Trend Report This macro creates trend reports based on several Summary reports.	—	—	RMFN9SUM

All macros marked by (*) offer the capability to generate control statements to create Overview records. You can download these statements either to your host system where you run the Postprocessor job, or you can use them directly (by cut and paste) in the Collector function.

Overview Record Macros

Macro	Lotus	Excel V5/V7	Excel 97
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 and in every detail you wish. The macro expects Overview records as described in "Extended Summary Report" on page 18-31. **Note:** For Lotus only — if you have data for one day, use RMFOVSUM, otherwise use RMFOVSYS.	RMFOVSUM RMFOVSYS	RMFOVSUM	RMFY9OVW
Workload Trend Report The macro creates summaries and graphics for a set of selected performance groups or domains of your installation. The macro expects Overview records as described in "Workload Trend Report" on page 18-33.	RMFOVWKL	RMFOVWKL	RMFY9WKL
Device Trend Report The macro creates a trend report for selected devices of your installation. The macro expects Overview records as described in "Device Trend Report" on page 18-35.	RMFOVDEV	RMFOVDEV	RMFX9DEV

Macro	Lotus	Excel V5/V7	Excel 97
Cache Statistics Trend Report The macro creates a trend report for selected cache subsystems of your installation. The macro expects Overview records as described in "Cache Statistics Trend Report" on page 18-35.	—	RMFOVCAC	RMFX9CAC
Create Overview Control Statements The macro creates OVW and EXCEPT statement that can be used to generate data for the above described Overview macros.	—	—	RMFX9MAK

DISCLAIMER OF WARRANTIES

All spreadsheet macros are sample code created by IBM Corporation. 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 home page on the Internet for enhancements:

- **http://www.ibm.com/s390/rmf/**

You find information about the Spreadsheet Reporter in the **Tools** page, which you can access from the home page directly.

Data Contained in RMF Spreadsheet Macros

All macros, both for Lotus and Excel, have a similar behavior. They contain the following sheets:

Main
: This is the first sheet in the macro (it has different names on the Lotus spreadsheets).

  It contains buttons to select a report or Overview record and a button to create a final report. When you create a final report, this sheet will be removed.

Disclaimer
: 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 the following one to three sheets. In Lotus they are color-coded in red.

In general, these are tabular reports which summarize and analyze the processed RMF data.

Charts The following sheets until the About or first Help sheet.

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 All spreadsheet macros should contain sufficient Help information to use them and 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 author.

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.

Sample Overview Record Control Sets

You can define for nearly all fields in Postprocessor reports overview conditions. This gives you an enormous flexibility to define your own customized reports. On the other hand, some examples of meaningful Overview Record control sets might be helpful to use this new feature.

As described earlier, you can use OVW and EXCEPT statements for defining the conditions, and it is recommended to use OVW. Nevertheless, you will find in this chapter some samples and templates with EXCEPT statements because they refer to statements that can be created through some spreadsheets which are available in the Excel version of the Spreadsheet Reporter. This implementation ensures that the Spreadsheet Reporter can be used also for preceding releases of OS/390 2.6 RMF.

Of course, you can use OVW in these templates if you create the statements directly.

This section describes an “Extended Summary Report” and two “Installation-Specific Reports” on page 18-32 which you can use to extend the existing capabilities of the Postprocessor.

Extended Summary 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 exception control statements directly without modification in every installation running its system in compatibility mode.

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 will be derived from RMF SMF record types 70, 71, and 72 subtype 1. You can find most of the data in the CPU, Paging, and Workload Activity reports.

```
Control Statements            ] SMF Record  Description
OVERVIEW(RECORD)              ]             Necessary to create the records
OVW(CPUBUSY(CPUBSY))          ]   70        CPU utilization data
OVW(MVSBUSY(MVSBSY))          ]             ]
OVW(APPLPER(APPLPER))         ]   72.1      V  Workload utilization
OVW(NUMPROC(NUMPROC))         ]   70        ----
OVW(EXCP(EXCP))               ]   72.1      Total EXCPs
OVW(EXCPRT(EXCPRT))           ]             EXCP rate
OVW(OCPU1(OCPU1))             ]   70        CPU contention
OVW(OCPU2(OCPU2))             ]             V
OVW(OCPU3(OCPU3))             ]             ----
OVW(INREADY(AVGUIRDY))        ]   70        SRM queues
OVW(OUTREADY(AVGUORDY))       ]             ]
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(AVGAPPC))         ]             V
OVW(AVGOMVS(AVGOMVS))         ]             ----
OVW(STOTOT(STOTOT))           ]  72.1       Workload storage
OVW(STOCEN(STOCEN))           ]             ]   and paging
OVW(STOEXP(STOEXP))           ]             ]   for all
OVW(SINGLE(SINGLE))           ]             ]   system workloads
OVW(BLOCK(BLOCK))             ]             ]
OVW(EXPSNGL(EXPSNGL))         ]             V
OVW(EXPBLK(EXPBLK))           ]             ----
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(AVGCSA))           ]  71         Common storage
OVW(AVGSQA(AVGSQA))           ]                 utilization
```

You find the exception control set for the Extended Summary report in file `d:\rmfpp\SAMPLES\JCL.TXT`.

The macros **System Week-Summary Report** and **System Overview Report** process converted Overview records which have been created based on this control set.

Installation-Specific Reports

These reports are based on exception control statements which require an installation-dependent qualification.

Excel version only: The exception control sets for the Installation-Specific reports can be generated from the associated report macros which process the Workload Activity report and the DASD Activity report. The macros will present you a list of performance groups or volume serial numbers from which you can select the entities for the trend reports.

If you want to report on the CPU utilization for your TSO performance groups, you can use the Postprocessor exception APPLPER and add a qualifier for your TSO performance group, for example:

```
EXCEPT(CPUTSO(APPLPER(00140),GE,0)
```

This statement could be taken if performance group 14 is used to control your TSO users.

Workload Trend Report

The idea for the Workload Trend report is to monitor the CPU utilization for your most important performance groups.

The template for the Overview records contains the exception control statements for workload utilization (APPLPER) and EXCP rate (EXCPRT) for one performance group. You have to define these statements for each selected performance group. 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 performance groups.

Template:

```
                                 -----------------------
EXCEPT(CPUnnnn(APPLPER(nnnn0),GE,0))      This part must be repeated for each
EXCEPT(EXPnnnn(EXCPRT(nnnn0),GE,0))       performance group, replace nnnn by PGN
                                 -----------------------
EXCEPT(NUMPROC(NUMPROC,GE,0))             This part is fixed and must be added
EXCEPT(CPUBUSY(CPUBSY,GE,0))              at the end of the report
EXCEPT(APPLPER(APPLPER,GE,0))
EXCEPT(EXCPRT(EXCPRT,GE,0))
```

If you want to use the template for performance group 5 and 21 in your installation, you have to change the template as follows:

```
OVERVIEW(RECORD)
EXCEPT(CPU0005(APPLPER(00050),GE,0))
EXCEPT(EXP0005(EXCPRT(00050),GE,0))
EXCEPT(CPU0021(APPLPER(00210),GE,0))
EXCEPT(EXP0021(EXCPRT(00210),GE,0))
EXCEPT(NUMPROC(NUMPROC,GE,0))
EXCEPT(CPUBUSY(CPUBSY,GE,0))
EXCEPT(APPLPER(APPLPER,GE,0))
EXCEPT(EXCPRT(EXCPRT,GE,0))
```

The macro **Workload Trend Report** processes converted Overview records which have been created based on this control set.

Special Version for Excel

If you are using Excel Version 5, Version 7, or Excel 97, you can get additional information by using these Overview records:

Template:

```
                                  ----------------------
EXCEPT(CPUnnnn(APPLPER(nnnn0),GE,0))      This part must be repeated for each
EXCEPT(EXPnnnn(EXCPRT(nnnn0),GE,0))       performance group, replace nnnn by PGN

EXCEPT(MPLnnnn(TRANSMPL(nnnn0),GE,0))     These exceptions can be added optionally
EXCEPT(TRXnnnn(TRANS(nnnn0),GE,0))
EXCEPT(RTMnnnn(RTIME(nnnn0),GE,0))
EXCEPT(EVLnnnn(EXVEL(nnnn0),GE,0))
                                  ----------------------
EXCEPT(NUMPROC(NUMPROC,GE,0))             This part is fixed and must be added
EXCEPT(CPUBUSY(CPUBSY,GE,0))              at the end of the report
EXCEPT(APPLPER(APPLPER,GE,0))
EXCEPT(EXCPRT(EXCPRT,GE,0))
```

The exception control set can be enhanced by up to four exceptions for each performance group. 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 functionallity of the macros is not restricted.

In order to use the additional exception control statements you must adhere to the following rules:

- Whenever you add additional exception control statements, you must add them for **all** performance groups 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:

  `TRANSMPL` Average number of resident transactions

  or

  `TRANSAVG` Average number of active transactions

  2. The following two exceptions have to be specified always together, they are meaningful typically for TSO transactions and batch jobs.

  `TRANS` Number of ended transactions per second
  `RTIME` Transaction execution time

  or

  `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 performance groups, or **any combination of the three groups** for all performance groups. Only keep in mind that you must always add the **same exceptions** for all performance groups.

  The spreadsheet macros (RMFOVWKL.XLS and RMFX9WKL.XLS) allow you to adjust the corresponding chart legends.

- You can use the macro **Create Workload-CPU-Paging Report** (RMFSCWCP.XLS and RMFR9WCP.XLS) to create the exception control set based on your performance groups, with and without the optional exception definitions. You can also specify whether you want TRANSMPL or TRANSAVG, TRANS or TRANSTOT, and RTIME and RTIMETOT.

- 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, even if you use a different exception criteria, as described in the bullet above.

Device Trend Report

The purpose of the Device Trend report is to observe the utilization and response time for devices of interest in your installation.

The template again consists of a variable part which must be repeated for each device included in the Overview record and a fixed part, which is necessary for the spreadsheet macro to determine the number of devices included in the report. For each device the response time components and the activity rate is included in the report.

Template:

```
                          -------------- Variable Part, replace vvvvvv with
                          -------------- the VolSer of the device
EXCEPT(QTvvvvvv(DQTAVG('vvvvvv'),GE,0))      Average IOS Queue Time
EXCEPT(PTvvvvvv(DPTAVG('vvvvvv'),GE,0))      Average Device Pending Time
EXCEPT(DTvvvvvv(DDTAVG('vvvvvv'),GE,0))      Average Device Disconnect Time
EXCEPT(CTvvvvvv(DCTAVG('vvvvvv'),GE,0))      Average Connect Time
EXCEPT(ARvvvvvv(DART('vvvvvv'),GE,0))        Device Activity Rate
                          ------------------------
EXCEPT(CPUBUSY(CPUBSY,GE,0))                 should be added at the end
```

If you want to use the template to monitor two devices with the volume serial numbers MVS215 and MVS217, you can generate the following exception control statements:

```
OVERVIEW(RECORD)
EXCEPT(QTMVS215(DQTAVG('MVS215'),GE,0))
EXCEPT(PTMVS215(DPTAVG('MVS215'),GE,0))
EXCEPT(DTMVS215(DDTAVG('MVS215'),GE,0))
EXCEPT(CTMVS215(DCTAVG('MVS215'),GE,0))
EXCEPT(ARMVS215(DART('MVS215'),GE,0))
EXCEPT(QTMVS217(DQTAVG('MVS217'),GE,0))
EXCEPT(PTMVS217(DPTAVG('MVS217'),GE,0))
EXCEPT(DTMVS217(DDTAVG('MVS217'),GE,0))
EXCEPT(CTMVS217(DCTAVG('MVS217'),GE,0))
EXCEPT(ARMVS217(DART('MVS217'),GE,0))
EXCEPT(CPUBUSY(CPUBSY,GE,0))
```

The macro **Device Trend Report** processes converted Overview records which have been created based on this control set.

Note: The macro **Create Device Report** allows you to create the exception control set based on your devices and the analysis performed on the RMF Device Activity Report (Excel Version only).

Cache Statistics Trend 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 again 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
EXCEPT(RHTcu(CASRHT(SSID(cu)),GE,0))         Total Read hits rate
EXCEPT(WHTcu(CASWHT(SSID(cu)),GE,0))         Total Write hits rate
EXCEPT(MRNcu(CASMRN(SSID(cu)),GE,0))         Normal Read miss rate
EXCEPT(MRScu(CASMRS(SSID(cu)),GE,0))         Sequential Read miss rate
EXCEPT(MRCcu(CASMRC(SSID(cu)),GE,0))         CWF Read miss rate
EXCEPT(MWNcu(CASMWN(SSID(cu)),GE,0))         Normal Write miss rate
EXCEPT(MWScu(CASMWS(SSID(cu)),GE,0))         Sequential Write miss rate
EXCEPT(MWCcu(CASMWC(SSID(cu)),GE,0))         CWF Write miss rate
EXCEPT(DFBcu(CASDFWB(SSID(cu)),GE,0))        DFW bypass rate
EXCEPT(CFBcu(CASCFWB(SSID(cu)),GE,0))        CFW bypass rate
EXCEPT(DFIcu(CASDFWI(SSID(cu)),GE,0))        DFW inhibit rate
EXCEPT(NCIcu(CASNCICL(SSID(cu)),GE,0))       Non-cache I/O ICL rate
EXCEPT(NCBcu(CASNCB(SSID(cu)),GE,0))         Non-cache I/O bypass rate
EXCEPT(ASYcu(CASASYNC(SSID(cu)),GE,0))       Async rate
                          ------------------------
EXCEPT(LASTRNG(CASRHT(SSID(cu)),GE,0))       should be added at the end
```

If you want to use the template to monitor two control units 0050 and 0068, you can generate the following exception control statements:

```
OVERVIEW(RECORD)
EXCEPT(RHT0050(CASRHT(SSID(0050)),GE,0))
EXCEPT(WHT0050(CASWHT(SSID(0050)),GE,0))
EXCEPT(MRN0050(CASMRN(SSID(0050)),GE,0))
EXCEPT(MRS0050(CASMRS(SSID(0050)),GE,0))
EXCEPT(MRC0050(CASMRC(SSID(0050)),GE,0))
EXCEPT(MWN0050(CASMWN(SSID(0050)),GE,0))
EXCEPT(MWS0050(CASMWS(SSID(0050)),GE,0))
EXCEPT(MWC0050(CASMWC(SSID(0050)),GE,0))
EXCEPT(DFB0050(CASDFWB(SSID(0050)),GE,0))
EXCEPT(CFB0050(CASCFWB(SSID(0050)),GE,0))
EXCEPT(DFI0050(CASDFWI(SSID(0050)),GE,0))
EXCEPT(NCI0050(CASNCICL(SSID(0050)),GE,0))
EXCEPT(NCB0050(CASNCB(SSID(0050)),GE,0))
EXCEPT(ASY0050(CASASYNC(SSID(0050)),GE,0))
EXCEPT(RHT0068(CASRHT(SSID(0068)),GE,0))
EXCEPT(WHT0068(CASWHT(SSID(0068)),GE,0))
EXCEPT(MRN0068(CASMRN(SSID(0068)),GE,0))
EXCEPT(MRS0068(CASMRS(SSID(0068)),GE,0))
EXCEPT(MRC0068(CASMRC(SSID(0068)),GE,0))
EXCEPT(MWN0068(CASMWN(SSID(0068)),GE,0))
EXCEPT(MWS0068(CASMWS(SSID(0068)),GE,0))
EXCEPT(MWC0068(CASMWC(SSID(0068)),GE,0))
EXCEPT(DFB0068(CASDFWB(SSID(0068)),GE,0))
EXCEPT(CFB0068(CASCFWB(SSID(0068)),GE,0))
EXCEPT(DFI0068(CASDFWI(SSID(0068)),GE,0))
EXCEPT(NCI0068(CASNCICL(SSID(0068)),GE,0))
EXCEPT(NCB0068(CASNCB(SSID(0068)),GE,0))
EXCEPT(ASY0068(CASASYNC(SSID(0068)),GE,0))
EXCEPT(LASTRNG(CASRHT(SSID(0068)),GE,0))
```

The macro **Cache Statistics Trend Report** processes converted Overview records which have been created based on this control set.

Note: The macro **Cache Subsystem Report** allows you to create the exception control set based on your devices and the analysis performed on the Cache Subsystem Activity Report (Excel Version only).

Using Example with Listing Data Set

The following example is intended for somebody who just installed the Spreadsheet Reporter and uses it the first time.

Please, use the sample RMF listing which is located in the LISTING directory (`d:\rmfpp\LISTING\Sample.Lis`), it contains one Summary report and two reports each of CPU, Paging, DASD and Workload Activity. The report duration is two hours.

1. Start the Spreadsheet Reporter by double-clicking on the **launch** object in the *Spreadsheet Reporter* program group.
2. Start the **Extractor** function from the LaunchPad (click on **Extractor**):
   a. Click on **Find** in the **Select Report-Work-Set** group-box of the Extractor function and select the listing named *Sample.Lis*.

      Click on **Open** , the file menu will close, the full path of the listing is displayed in the text-box right to **Find** , and the directory path for the new Report-Work-Set is displayed.
   b. Enter a description for the new Report-Work-Set in the **Description** text-entry field of the **New Report-Work-Set** group-box. Later, this description will be used to identify your Report-Work-Set in all other dialogs. You will also see that **Run** is enabled as soon as you enter the first character.
   c. Click on **Run** to start the extraction process. The **Create Report-Work-Set** dialog appears, and you can watch the extraction process. The progress indicator shows how much of the RMF listing has been processed while the **Statistics** group-box displays the report types and number of reports found for each type.

      You may observe that the extraction of the Workload Activity report takes much longer than for any other report type. This is typical for many installations. In this example, the two Workload Activity reports occupy about 80% of the complete listing.
   d. When the extraction process has been finished, click on **Done** and **Exit** on the Extractor dialog.
3. Start the **Converter** function from the LaunchPad (click on **Converter**):
   a. The Converter dialog displays all available Report-Work-Sets in the **Select Report-Work-Set** group-box. In general, newly created Report-Work-Sets will be appended to the drop-down list-box. You can select a Report-Work-Set from the drop-down list. The three text-boxes below the drop-down list-box always display the *creation date*, the *RMF source listing* and the *reporting range* for the selected Report-Work-Set in the list-box.

      In our example, there is probably only the Report-Work-Set which you just have created.

Click on **OK** to display the reports of the selected Report-Work-Set in the **Available RMF Reports** group-box. Now, the **System Id** list-box shows the available SMF-ids and Sysplex-ids contained in the Report-Work-Set.

In our example, you see two identifiers:

1) RMF1* is the SMF-id for system RMF1, and it is used for the Summary Report of the source listing.
2) RMF1 describes the same system, and it is used for **all** other report types.

b. Now, you can convert all reports contained in the listing for the 10.00 o'clock interval:

1) Select `RMF1` in the **System id** list-box
2) Select `11/11/1996` in the **Days** list-box
3) Select `10.00.02` in the **Intervals** list-box
4) Select all reports in the **Reports** list-box

Convert is enabled as soon as you select an item in the **Reports** list-box.

c. Click on **Convert** to start the conversion of the selected reports. The **Conversion Progress** dialog is displayed, and you can observe which report is converted. Usually, this process is very fast. When it has finished, **Exit** is enabled.

d. Click on **Exit** on the progress dialog and **Exit** on the Converter dialog. You can invoke the dialog at any time and convert additional reports of the same or other Report-Work-Sets.

4. Click on **Spreadsheet** on the LaunchPad to open the **Macros** folder. The following steps depend very much on your installation:

   - If you did not install any spreadsheet macro, the folder is empty and our demonstration is finished.
   - If you have used *Typical* installation, you will find two spreadsheet macros (RMFSCMN.WK4 and RMFSCMN.XLS) and two subfolders (L123 and XCel) in the folder.
   - If you have installed only one spreadsheet application, you will find only one macro and one subfolder (for example RMFSCMN.WK4 and L123 if you only have installed the Lotus macros).

   For the following steps, we assume:

   - You have installed Lotus 1-2-3 Version 5 (International English) on your workstation
   - You have installed the RMF Lotus 1-2-3 Spreadsheet Macros (you see the RMFSCMN.WK4 icon and the L123 subfolder).

   The example works identical for the RMF Excel Macros.

5. Double-click on the RMFSCMN.WK4 icon to start Lotus 1-2-3 for Windows which automatically loads the **RMF Spreadsheet Main Dialog** macro.

   The purpose of this macro is to start all other RMF macros and to customize the Spreadsheet Macros setup.

6. Because this is the first time you use the RMF spreadsheet macros, you have to *Customize* the macro setup:

a. In the yellow field, enter the path were you have installed the Spreadsheet Reporter, and ensure that you **enter a backslash** at the end of the path definition.

b. Click on Customize to propagate the correct path information to all other report spreadsheet macros. You will see the message *Customization Started* and when it finishes you will be informed by *Customization successfully completed*.

c. Save the changed content of the RMF Spreadsheet Main Dialog macro by selecting **Save** in the File Menu bar.

This step is only necessary once after having installed the Spreadsheet Reporter or the RMF spreadsheet macros.

Note: This step is not required for Excel 97.

7. Now, you can use the spreadsheet macros to process the converted reports.

   Click on Open Report Macros and select the **Open RMF Report Spreadsheets** macro.

   a. This is a very basic macro which does not process the RMF report data. Nevertheless, it demonstrates the general concept of the macros.

   b. The macro is opened immediately, and is placed on top of the management macro. Now, you can open converted RMF reports from Report-Work-Sets.

   c. Click on Select Report-Work-Set and Open RMF Reports on the first worksheet of the spreadsheet macro.

   d. Select a Report-Work-Set from the **Select Report-Work-Set** dialog. A second dialog **Open RMF Report Spreadsheets** opens which lists the content (all converted RMF reports) of the Report-Work-Set. Each report is shown by its *SMFid-creation date-creation-time-report type* in a similar way as in the Converter dialog.

   e. You have two options:

      1) Select reports and open then as individual worksheets (or workbooks).

      2) Select reports and create **one** spreadsheet (workbook) which contains all selected RMF reports. In this case, you must check the **Create combined worksheet** check-box at the bottom of the dialog.

   f. Select some reports, activate, for example, the **Create combined worksheet** in the check-box and click on OK . A new spreadsheet is created and a filemenu is displayed which requests a name for the new spreadsheet.

   g. Enter a name and click on OK . The new spreadsheet is opened. It contains all your selected reports as separate worksheets. You can view each report by selecting the spreadsheet tab. The spreadsheet tab contain basically the same information as displayed in the **Open RMF Report Spreadsheet** dialog.

8. The purpose of this demonstration is simply to demonstrate the basic handling of RMF spreadsheet macros. Now, you can work with the RMF reports you just opened.

9. If you want to open other RMF spreadsheet macros, return to the RMF Spreadsheet Main Dialog. You can select this macro (RMFSCMN.WK4) from the list in the Window menu in the Lotus menu bar.

All other macros process RMF data. You can either try them or just load them and view their reports. All macros which process RMF data come with sample data. This allows you to examine the type of report, analysis and graphic they provide without processing your own data.

10. When you have finished with the first try, you can close Lotus 1-2-3 (or Excel).
11. Click on **Exit** to leave the Spreadsheet Reporter.

It is now recommended that you perform this sequence of steps with your own SMF data. Therefore, you might start with the Collector to create a listing on your OS/390 system and to download the data.

What You Should Know about Spreadsheet Applications

The macros for Lotus and Excel are very similar but in general not exactly identical. Also the behavior is similar but differs sometimes. To use the macros successfully, you should read the following sections:

- "What You Should Know About Lotus 1-2-3"
- "What You Should Know About Excel"
- "Known Problems" on page 18-41

What You Should Know About Lotus 1-2-3

The **correct setup** is very important. Please verify that the **International** setting shows **1:9,999.99 @fn(x,y)** as *Punctuation* setting. This is very important because the macros do not work without this setting. You can examine the International setting by selecting **User Setup** in the **Tools** Menu. The Tools drop-down menu is displayed in the Lotus menu bar. The International setting dialog can then be displayed by clicking on **International** on the User Setup menu.

Macro **RMFSCWCP** does not work properly for RMF reports which have been generated by an RMF Postprocessor with a version of RMF 5.1 or older. The Paging and Storage Summary will contain **ERR** indications. All other sheets should show correct data.

The Lotus 1-2-3 macros only work with Lotus 1-2-3 Version 5 International English (or US) version.

What You Should Know About Excel

Error Dialogs

If you try to process a Report-Work-Set for which no report has been converted, Excel currently displays an Error dialog. Click on **End**, the error is not processed by the macro yet.

Language Setting

In case the macros do not work, select the **Options** dialog from the **Tools** drop-down menu. On the **Module general** page, select **English/United States** for Language.

Excel 97 Compatibility

All spreadsheet macros for Excel Version 5 run also with Excel Version 7, but not with Excel 97. If you plan to use Excel 97, you have to install the Excel 97 macros via Custom installation. In this case, it is highly recommended not to install the other Excel macros.

Excel Version 5 Compatibility

All macros for Excel Version 5 and Version 7 have been developed using Microsoft Excel Version 5.0c. Therefore, you may encounter problems when you use an earlier version of Microsoft Excel Version 5. You can check the version by selecting **Help** and then **About Microsoft Excel ...** from the menu bar. Version 5.0c is displayed as **Microsoft Excel Version 5.0c**.

If you run into problems while using an older version, we recommend that you select **Debug** from the error dialog and put a comment (quote) sign before the statement where the error occurs. All problems are related to certain property settings for objects in Visual Basic and the macros will work without this property settings, too.

Excel Version 7

We never tested the spreadsheet macros on this version, but feedback from users indicated that they work without any problems.

Known Problems

Windows NT 3.5.1 or earlier: The Spreadsheet button on the LaunchPad does not work. On Windows NT 4.0 or Windows 95, the button opens a folder which presents the available spreadsheet macros. Because folders do not exist on earlier versions of Windows NT, this button cannot be used. But, you can start the main spreadsheet macros directly from the ***RMF Spreadsheet Reporter*** program group.

Chapter 19. Performance Monitoring of OS/390

Performance Management on the PWS

Performance Monitoring of OS/390 (PM of OS/390) enables you to collect, analyze, and graphically display performance data from several OS/390 host environments directly from your workstation. It uses the standard PM Common Functions interface to access the necessary data.

This chapter covers the following topics:

- What is Performance Monitoring of OS/390?
- Installing PM of OS/390
- Running PM of OS/390
- Details about Working with PM of OS/390
- Available Performance Data

What is Performance Monitoring of OS/390?

Performance Monitoring (PM) of OS/390 enables you to collect, analyze, and graphically display performance data from one or more OS/390 host environments, using the services provided by PM Common Functions. For a full description of all the services provided by PM Common Functions and a description of interapplication processes, see *Using Performance Monitoring Common Functions*.

Overview

PM of OS/390 provides an interface between the OS/2 workstation and the OS/390 sysplex that gives you the flexibility to create unique scenarios that monitor the performance of your system. You can collect real-time data in graphic and text mode, 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*.

With PerfDesks, you create customized objects containing configuration and performance views of your OS/390 system, and you can reuse these views any time. When you save the configuration and performance windows in a PerfDesk, all the actions required to produce these windows are saved. You can simply load the PerfDesk and start it whenever you want to view performance data or to monitor configuration resources in your system again from the same angle.

Note: A saved PerfDesk does not contain any performance data. The PerfDesk works with new performance data each time it is called.

When you start PM of OS/390, the **PM of OS/390 Domain Configuration** is displayed. This shows a list of OS/390 domains in which each domain represents an OS/390 sysplex. From here, you can logon to any domain and can manage the resources in the corresponding sysplex.

PM of OS/390 takes its input data from RMF Monitor III. This data is suitable for monitoring and analyzing performance in real time and in the recent past. How far backward in time historical data can be provided, depends on the size of the Monitor III Data Gatherer buffer and the VSAM data sets you have allocated.

PM of OS/390 provides a selected subset of the information provided by the Monitor III gatherer: general performance data; performance data for jobs; and for systems running in goal mode workload-related performance data like

- WLM workloads
- WLM service classes
- WLM service class periods
- WLM report classes

In addition to PM of OS/390, a second PM application is available on OS/2 as PM of SNA. With this integrated approach, you can monitor and analyze performance data for both OS/390 host systems and SNA networks from one workstation.

Understanding Performance Monitoring Concepts

Let us take a look at the concepts behind PM Common Functions — what they are, what they do, and how they meet your performance management requirements.

PM Common Functions provides a common interface for collecting, viewing, logging, and analyzing performance in a multiprotocol, multiplatform, distributed environment. It is a flexible, expandable solution for integrated performance management.

Single Point of Control

PM Common Functions provides a way to monitor performance in a distributed environment from a single point of control, just as was possible with host systems. In addition, it provides a way to group data from various sources logically, and display it in such a way that it can be monitored integrally.

PM Common Functions consist of two components, which work together to provide a complete performance-management solution:

- Basic functions for managing performance
- User interface

PM Common Functions is the base that provides a general set of functions for managing performance. Applications are installed on PM Common Functions, and then use its functions to collect specific types of performance data. Applications allow you to view and access this data and access configuration objects using the PM Common Functions. PM of OS/390 is one of the applications that work with PM Common Functions.

Accessing PM Common Functions

PM Common Functions provides basic support for the major performance management tasks in an OS/390 host environment and network. You receive this support directly through the user interface, and indirectly through the applications.

The basic functions of PM Common Functions fall into the following categories:

- Performance-data collection and monitoring

  PM Common Functions provides mechanisms for collecting, displaying, monitoring, and logging real-time performance data of your system and network.

- Configuration access and management

  PM Common Functions provides functions for accessing the configuration of the sources whose data you are monitoring. You can display these resources hierarchically as a tree, and can display resources from different applications in the same window.

- Performance data logging

  With PM Common Functions, you can log real-time data in a file in “comma-separated-variables” (CSV) format, so that it can be read by spreadsheet applications.

User Interface

With the user interface, you can access and manipulate performance data and view it graphically, and even construct your own logical views, using:

- CfgViews (configuration views) for viewing and accessing configuration data
- DataViews for viewing and collecting performance data
- PerfDesks for organizing logical views of your systems and networks
- Toolboxes for managing and defining tasks
- Object Browser for managing system objects

What is a CfgView?

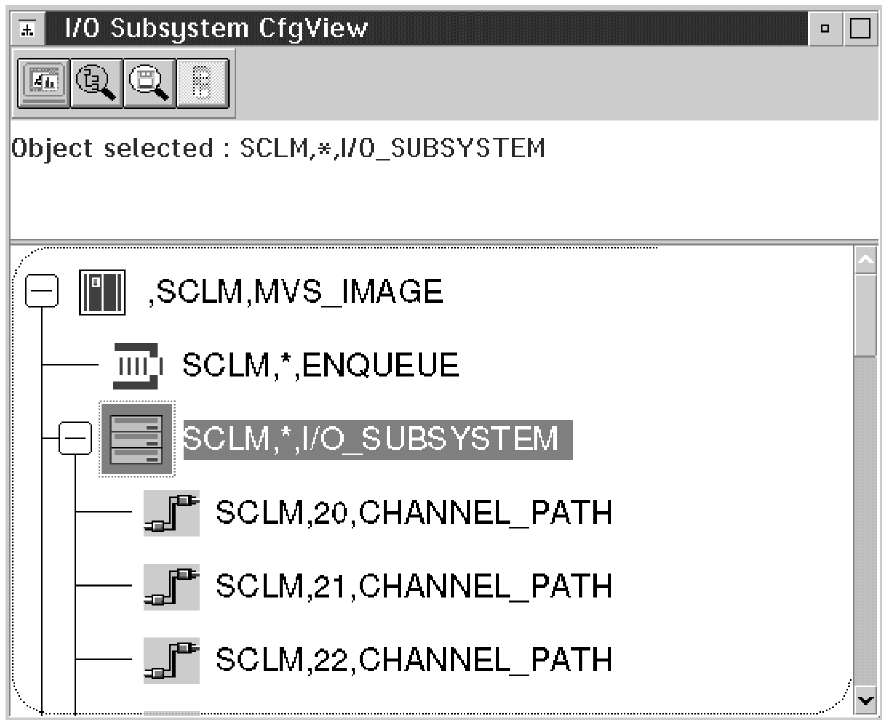

Figure 19-1. Example of CfgView Window

A CfgView is a window that displays configuration objects. Each object represents an actual resource, which can be a program, a device, a system, or a subsystem. These resources are represented in a CfgView by icons, text, or both, showing the hierarchical relationship of the resources by means of a tree structure.

For example, an application could show a hierarchical relationship between an I/O subsystem and a channel by showing the I/O subsystem as the parent and all the channels attached to it as the children. You can show or hide the children.

Each resource has a type associated with it. A set of resource types is defined by each application. For example, CHANNEL_PATH might be a valid resource type for one application, while another defines the same resource as CHANNEL. Each resource type has a set of counters, which are attributes that provide performance information about the resource. For example, the counters for a channel resource could include the utilization of the channel path. Although counters are not explicitly shown in a CfgView, they can be accessed for all resources.

Application CfgViews: You use CfgViews to show the configuration of resources for performance management applications. When you start an application, you see a list of domains that have agents as children. Agents pass the values of the counters for resources to the application, which displays this data in DataViews. The resources that are managed by the agent are shown as children, grandchildren, and so on, depending on the number of levels in the hierarchy. In PM of OS/390, each domain, which is an OS/390 sysplex, has exactly one agent.

More than one application can access and show the same physical resource. For example, two applications might show the same workstation in the application CfgView with different names or different counters.

CfgViews in PerfDesks: You can create your own CfgViews that show specific resources or groups of resources that are controlled by one or more applications. In other words, you can display, in one CfgView, resources from several protocols or platforms. To do this, you define a configuration specification that defines the resources you want to include in the CfgView. The configuration specification lets you identify the resource you want to see, and from which application the data are to be taken.

You can add these CfgViews to one or more PerfDesks, either as part of the PerfDesk or as an external object.

What is a DataView?

A DataView is a window that contains a tabular or graphical representation of performance data for one or more resources. A DataView can show the data as it occurs in real time by continually updating the representation at regular intervals. Data is available in form of single-valued or multi-valued counters (see “Counters” on page 19-12).

An example of a PM data view with two single-valued counters is shown in Figure 19-2.

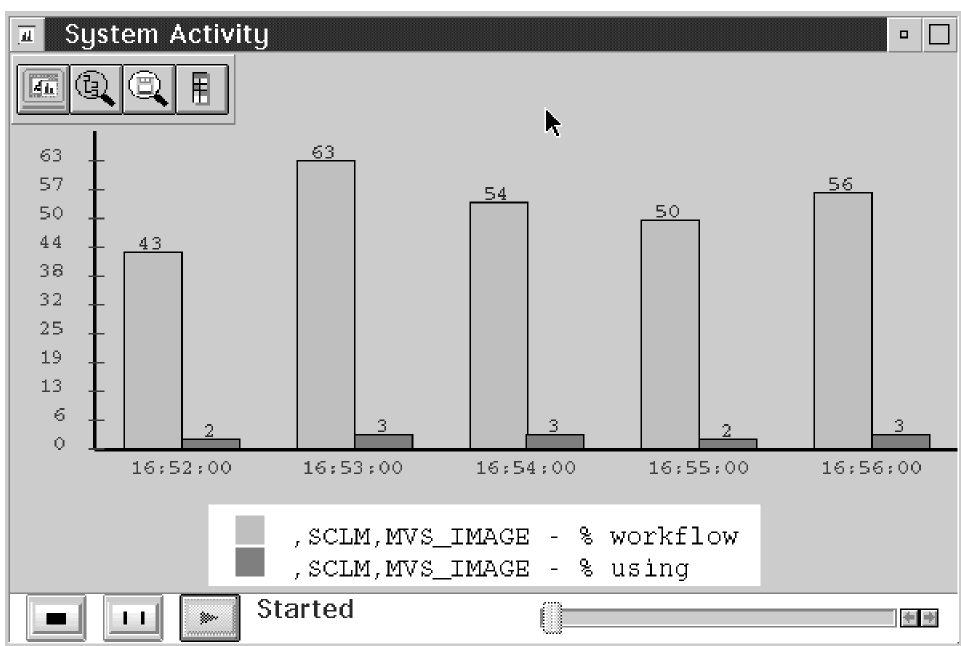

Figure 19-2. Example of a DataView with Single-valued Counters

Besides other graphical representations, there is also a textual form of the PM DataViews, as shown in Figure 19-3.

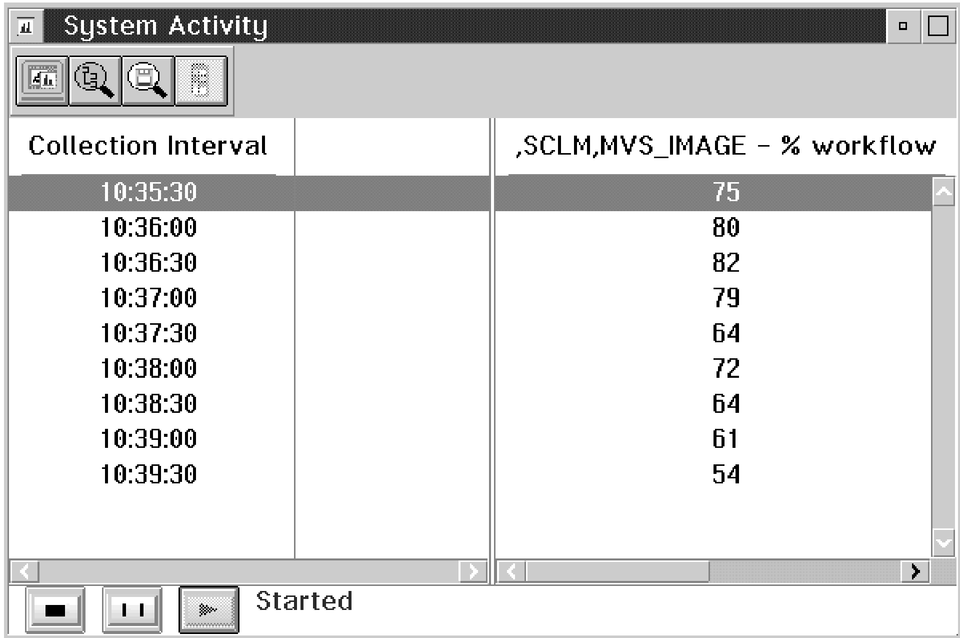

Figure 19-3. Example of a DataView with Single-valued Counters (Textual Format)

Multi-valued counters are also supported, and the resulting values can be filtered. An example of a PM data view with a multi-valued counter is shown in Figure 19-4.

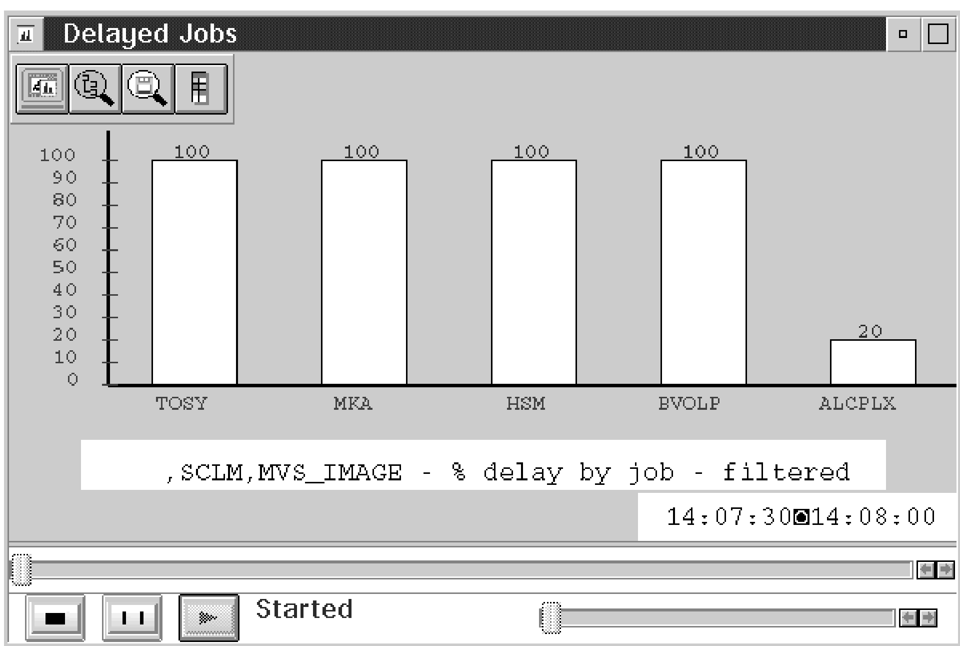

Figure 19-4. Example of a DataView with Multi-valued Counters

Real-Time Performance Data: In a DataView, you can collect and display data at regular intervals in real-time.

Real-time performance data reflects the current performance of the resource or resources being monitored. You obtain the real-time data by starting a DataView. The DataView contains values that represent a collection interval. For each collection interval, the new data is added to the DataView.

Resources and Counters in DataViews: In a DataView, you can view one or more counters for one resource, one counter for more than one resource, or many counters for many resources, and resources and counters from more than one application.

To add resources and counters to a DataView, you create resource and counter pairs called *series*. Each resource-and-counter pair appears as one metric, such as one bar on a bar chart, in the case of single-valued counters.

Example

You can create a DataView with one MVS image, such as SCLM, as the resource and `%workflow` and `%using` as the counters. Or you can create a DataView with `%workflow` and `%using` as the counters, and two MVS images, SCLM and TSOLAB, as the resources.

You can include all of this performance data in one DataView by creating the following resource and counter pairs:

```
SCLM - %using
SCLM - %workflow
TSOLAB - %using
TSOLAB - %workflow
```

If you display the data as a bar chart in the DataView, each of these resource-and-counter pairs would be a separate bar.

Summarization Intervals: You can set the time that you want to elapse between each update of the DataView. Data related to a specific counter is summarized after the time interval specified. The way in which the data is summarized depends on the counter. The summary could be, for example, the average or the sum of the data.

What is a PerfDesk?

PM Common Functions enables you to construct your own control panels to access and view performance data from sources across a sysplex. Each of these control panels, called PerfDesks, provides unique performance scenarios for your entire sysplex or any part of it. You can create as many PerfDesks as you want, and can use as many of them as your system resources allow.

In one PerfDesk, you can include one or more DataViews and CfgViews, so that you can simultaneously:

- Monitor current performance
- View the performance of resources across multiple sysplexes
- View and access resources from more than one application

You can save the PerfDesks and reload them later.

If you have more than one application running on PM Common Functions, you can access real-time data from more than one application and view it graphically in the same PerfDesk. This feature allows you to see all parts of your sysplex, regardless of the types of resource.

You can also share views of real-time performance data among PerfDesks. As the performance data is updated over a specified time interval, the view is simultaneously updated in all the PerfDesks.

You can access resources and view performance data by simply starting the PerfDesk, which also starts all its individual components. After it is started, a PerfDesk continues to collect and display performance data until you stop it or log off your application or PM Common Functions.

What is a Toolbox?

A Toolbox provides a graphical way of managing, configuring, and defining tasks. You can create a Toolbox for each task or set of tasks that you want to perform. Each Toolbox contains two lists:

- A list of activities related to the task
- A list of actions related to each activity

You can define both the activities, which appear on the left-hand side of the Toolbox, and the valid actions for each activity, which appear on the right-hand side.

You will find more details in “The Contents of a Toolbox” on page 19-63.

What is an Object Browser?

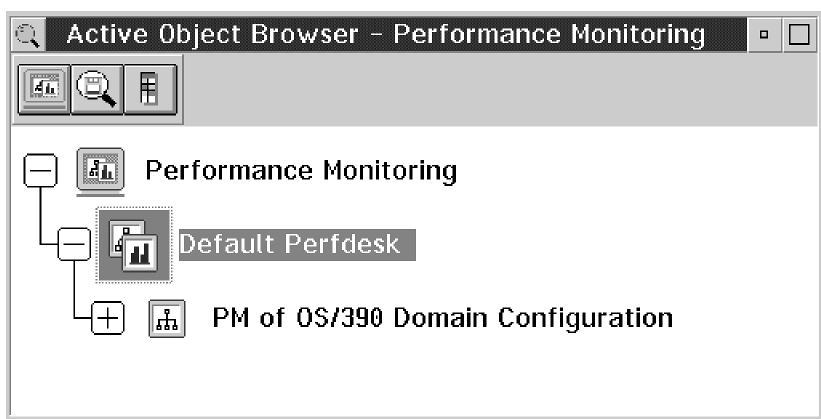

Figure 19-5. Example of Active Object Browser

An Object Browser contains icons or text descriptions of PerfDesks and the objects they contain. You can perform actions on the objects from the Object Browser.

The Object Browser also shows how the objects are related to one another. For example, you can expand a PerfDesk icon to see the DataViews it contains. You can further expand the DataView to see the series it contains.

The objects in the Object Browser have a menu from which you can select an action you want to perform, such as closing a DataView. You can customize these menus by adding global actions as menu items, and you can define local actions that are available in only one menu.

You can also open separate windows, with an object from any level as the root icon of the new window. For example, if you want one PerfDesk to be the root in a separate window, double-click on its PerfDesk icon to open a separate window with this icon as the root of the new window.

PM Common Functions provides two Object Browsers to show objects of two types:

- Active objects
- Saved PerfDesks

The Performance Data Model

This description of the data model supported by PM of OS/390 includes:

- The configuration hierarchy with resource types supported by PM of OS/390, which can be shown in PM configuration views
- Counter types supported by PM of OS/390, which can be shown in PM data views
- Work scopes supported by PM of OS/390

Resources

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.

In PM of OS/390, resources are named according to the naming conventions of PM Common Functions. A resource is identified by

```
<upper level qualifier>,<resource name>,<resource type>.
```

Here, <x> means an instance of x. For example, the MVS image named SYS1 is identified by

```
,SYS1,MVS IMAGE.
```

Note that, in this case, the upper level qualifier is empty, because the resource name SYS1 is unique within the sysplex corresponding to the PM of OS/390 domain. For example, a volume named DATA01 mounted in MVS image SYS2 is identified by

```
SYS2,DATA01,VOLUME.
```

For some resources, only one instance exists within its containing resource. The naming convention for these resources is that the resource name is an '*'. For example, there is only one SQA in an MVS image, so:

```
SYS1,*,SQA
```

denotes the SQA of MVS image SYS1.

The `upper level qualifier` will be referred to in the application as `ulq`.

Resources are shown in configuration views. The resources are ordered hierarchically in a configuration tree. The nodes of this tree can be expanded interactively, starting from the root node. The root node represents the sysplex or stand-alone MVS image belonging to the domain.

The links between the nodes of the configuration hierarchy represent relationships, but the semantics of these relationships are not uniform for the whole tree. In most cases, the relationship is "contains," which means that the resource represented by the higher level node is an aggregate of the resources represented by its lower level nodes.

For some of the listed resource types, multiple instances may exist. For example, there may be multiple MVS images in a sysplex. These instances have names. The names need not necessarily be unique within the PM of OS/390 domain. For example, a volume named DATA01 need only be unique within its containing MVS image.

There are several types of resource in the configuration tree of PM of OS/390:

- Resources of which multiple instances may exist.

  For example, the MVS image named SYS1, identified by:

  `,SYS1,MVS IMAGE.`

- Resources of which exactly one instance exists.

  Their name is an '*'. For example, the logical resource of the processor in an MVS image named SYS1, identified by:

  `SYS1,*,PROCESSOR`

  In this example, SYS1 is the `ulq`.

A single resource can be available several times in the configuration view. For example, an LCU can be defined for several channel paths.

Configuration Hierarchy

PM of OS/390 can visualize information for the following resource types:

- SYSPLEX
  - MVS IMAGE
    - IO SUBSYSTEM
      - CHANNEL PATH
        - LOGICAL CONTROL UNIT
          - VOLUME
    - PROCESSOR
    - STORAGE
      - AUXILIARY STORAGE
      - CENTRAL STORAGE
        - CSA
        - SQA
        - ECSA
        - ESQA
      - EXPANDED STORAGE
    - ENQUEUE
    - OPERATOR
    - SW SUBSYSTEMS

- JES
- HSM
- XCF

In an actual configuration:

> Some resource types can have one and only one instance within one containing resource, and these have an '*' for `<resource name>`
> Some resource types can have multiple instances within one containing resource, and must have a genuine `<resource name>` for each instance
> A `<resource name>` must be unique only within its containing resource, not necessarily within the PM of OS/390 *agent*.

The following list shows the above information within the structure of a configuration tree of PM of OS/390:

- ,<sysplex name>,SYSPLEX
  - <mvs image name>,MVS IMAGE
    - <mvs image name>,*,IO SUBSYSTEM
      - <mvs image name>,<channel path name>,CHANNEL PATH
        - <mvs image name>,<logical control unit name>,LOGICAL CONTROL UNIT
          - <mvs image name>,<volume name>,VOLUME
    - <mvs image name>,*,PROCESSOR
    - <mvs image name>,*,STORAGE
      - <mvs image name>,*,AUXILIARY STORAGE
      - <mvs image name>,*,CENTRAL STORAGE
        - <mvs image name>,*,CSA
        - <mvs image name>,*,SQA
      - <mvs image name>,*,EXPANDED STORAGE
        - <mvs image name>,*,ECSA
        - <mvs image name>,*.,ESQA
    - <mvs image name>,*,ENQUEUE
    - <mvs image name>,*,OPERATOR
    - <mvs image name>,*,SW SUBSYSTEMS
      - <mvs image name>,*,JES
      - <mvs image name>,*,HSM
      - <mvs image name>,*,XCF

The PM configuration hierarchy supports a tree-structured, which means that each node (except the root node) has exactly one parent node. If a resource can be accessed via alternate paths, it will appear several times in the tree. If, for example, in MVS image SYS1, the logical control unit 0009 is attached to channel paths 5F and 69, then the resource:

```
SYS1,0009,LOGICAL CONTROL UNIT
```

will be contained in both:

`SYS1,5F,CHANNEL PATH`

and:

`SYS1,69,CHANNEL PATH.`

Counters

PM of OS/390has two formats for presenting performance data: single-valued counters and multi-valued counters.

- Single-valued counters for a resource

  Examples

  `% utilization` (of a processor, of a channel, ...)
  `i/o activity rate` (of a logical control unit, ...)

  Figure 19-2 on page 19-5 is an example of a DataView with single-valued counters.

  If the counter is for a work scope instance, you must specify not only the resource, but also the work scope in the dialog for the collection specification of the counter.

- Multi-valued counters

  These consist of a list of pairs containing the name of an instance of one of the following, with its corresponding data value:

  – MVS image
  – Channel path
  – Logical control unit
  – Channel path and CU
  – Volume
  – Job
  – WLM service class period
  – WLM report class
  – WLM service class
  – WLM workload

  Examples

  ```
  % utilization by job
  # delayed jobs for i/o by mvs image
  ```

  Figure 19-4 on page 19-6 is an example of a DataView with multi-valued counters.

  The unique indicator in the name of a multi-valued counter is the keyword `by`.

Filters

Multi-valued counters are presented in PM of OS/390 in descending order of the values they contain. You can reduce the amount of data contained in multi-valued counters by specifying a filter. Only counters that match the filter criteria are monitored. The supported filter criteria are:

1. A set of name patterns, which can contain wildcards. If multiple name patterns are specified, list elements with a name matching at least one of the patterns pass the filter
2. List elements with value ≥ threshold (or value ≤ threshold)
3. A specified sequence by either name or value
4. n list elements with largest (or smallest) values, n ≥ 1

If more than one filter criterion is specified, the processing takes place in the order specified above. In addition, sort criteria can be specified.

Work Scopes

A work scope is the specification of an entity of work. PM of OS/390 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

Counters with values for work scope instances are available in two ways:

1. As a single-valued counter, where the corresponding work scope name has been specified
2. As a multi-valued counters (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.

Installing PM of OS/390

There are two types of communication you can use when working with PM of OS/390:

- APPC

  An APPC communication requires that you define an APPC TP profile on your host system and that you define an APPC setup on your workstation.

- TCP/IP

  You can implement the TCP/IP communication either by starting a host server as a started task address space (with operator START command).

The configuration of both types enables you to have one implementation as backup solution for the other one. You can decide at run time to choose either an APPC or TCP/IP connection.

The installation of PM of OS/390 comprises the following steps. As described above, one of the first two steps is optional, but recommended.

1. Preparation of an APPC connection
   a. Host - Definition of an APPC TP profile
   b. Workstation - Definition of the OS/2 APPC setup
2. Preparation of an TCP/IP connection
   a. Host - Setup of the server start procedure
   b. Workstation - Server bring-up by the client
3. Installation of Performance Monitoring Common Functions

   The Performance Monitoring Common Functions (PM Common Functions) is a prerequisite for PM of OS/390. Install this product before installing PM of OS/390. If the Performance Monitoring Common Functions is already installed on your workstation, you can omit this step.
4. Installation of PM of OS/390
5. Tailoring PM of OS/390 to your needs

These steps are described in detail in the following sections. When you have completed them, you are ready to start and use PM of OS/390.

Hardware Requirements

- Pentium PC ≥ 90Mhz
- Super Video Graphics Adapter (SVGA) video card
- Harddisk free space > 40 MB
- RAM 32 MB

Software Requirements

- Minimum OS/2 3.00 Warp (family) with Fix Pack 22 (CSD214): XR_W022 (where _ is a language code, if a non US-version of OS/2).
- For communication with OS/390, one of the following:

  Within an APPC environment:
  – Communications Manager/2 Version 1.11 or higher, with 3270 emulator and APPC support
  – Personal Communications/3270 Version 4.1 for OS/2 CSD #1 and APPC support from Communications Manager/2 Version 1.11 or higher
  – Personal Communications/3270 Version 4.1 for OS/2 CSD #1 and APPC support from Communication Server Access Feature

  Within a TCP/IP environment:
  – TCP/IP Version 3.0 or higher

Preparation of an APPC Connection

Host - Definition of an APPC TP Profile

The host part of PM of OS/390 communicates with the DDS OS/390 server via local sockets. In order to build up the connection between the host part of PM of OS/390 and the DDS OS/390 server, the pathname /tmp/gpmserve/pmos390 will be used. This implies, that

- OS/390 UNIX System Services must be configured and active
- A /tmp directory must exist in the file system used by OS/390 UNIX. If it does not exist, it can be created with the MKDIR command.
- The userid associated with GPMSERVE must have write access to /tmp .
- The file /tmp/gpmserve/pmos390 and the directory /tmp/gpmserve are reserved for use by GPMSERVE, no user file with that name may exist in /tmp.
- To install and run PM of OS/390, you need to change or create the APPC TP Profile on the host according to your needs.

A sample job to create an APPC TP Profile for PM of OS/390, is shipped with OS/390 in SYS1.SAMPLIB(GPMTPSAM). This sample job will store the created TP Profile in the data set SYS1.APPCTP. Every PM of OS/390 user connecting via APPC must have a userid on the OS/390 host.

APPC TP Profile for PM of OS/390

```
//GPMTPSAM JOB (<ACCT>),'<PROGRAMMERNAME>',CLASS=A,REGION=32M,...
//STEP1    EXEC PGM=ATBSDFMU
//SYSPRINT DD SYSOUT=*
//SYSSDLIB DD DSN=SYS1.APPCTP,DISP=SHR
//SYSSDOUT DD SYSOUT=*
//SYSIN    DD DATA,DLM='@@'
     TPDELETE
     TPNAME(GPMMVS)

     TPADD
     TPNAME(GPMMVS)
     ACTIVE(YES)
     TPSCHED_DELIMITER(END_OF_TPSCHED)
     KEEP_MESSAGE_LOG(ERROR)
     MESSAGE_DATA_SET(&SYSUID..&TPDATE..&TPTIME..JOBLOG)
     DATASET_STATUS(MOD)
     CLASS(FAST)
     JCL_DELIMITER(END_OF_JCL)
//<JOBNAM> JOB (<ACCT>),<NAME>,REGION=32M,MSGLEVEL=(1,1)
//*
//STEP1    EXEC PGM=GPMCSRV
//*TEPLIB  DD   DISP=SHR,DSN=SYS1.SCEERUN
//*
//GPMINI   DD   DISP=SHR,DSN=SYS1.SERBPWS(GPMINI)
//CEEDUMP  DD   DUMMY
//SYSPRINT DD   DUMMY
//SYSOUT   DD   DUMMY
//*YSPRINT DD   DISP=SHR,DSN=&SYSUID..GPM.SYSPRINT
//*YSOUT   DD   DISP=SHR,DSN=&SYSUID..GPM.SYSOUT
END_OF_JCL
END_OF_TPSCHED
@@
```

You have to perform the following steps to customize this JCL for your environment:

1. Change the job record
2. Verify the SYSSDLIB data set name - this is the VSAM data set that contains the APPC profiles in your installation.
3. Verify the transaction profile definition, for example
   - TPNAME
   - MESSAGE_DATA_SET
   - KEEP_MESSAGE_LOG

   For details, see *OS/390 MVS Planning: APPC/MVS Management*.
4. Change the job record of the transaction JCL
5. You may have to add a STEPLIB DD-Card to the transaction JCL if the OS/390 Language Environment runtime library SCEERUN is not included in the linklist. In this case, SYS1.SCEERUN has to be APF-authorized.
6. To create the host trace and printout data sets:

- Host trace: Replace the 'SYSOUT DD DUMMY' definition with a data set definition and create the data set with the following attributes: DCB=(DSORG=PS,RECFM=FBA,LRECL=121,BLKSIZE=0)
- If required, you may replace the 'SYSPRINT DD DUMMY' definition with a data set definition and create the data set with the following attributes: DCB=(DSORG=PS,RECFM=VB,LRECL=1028,BLKSIZE=0)

Workstation - Definition of the OS/2 APPC Setup

In this section, we assume that VTAM and APPC have already been set up on the OS/390 host. In the VTAM Setup, an additional VTAM application major node is required for APPC communications.

The following information is provided primarily for verification purposes, that is, you need to check the information in your Communications Manager configuration file (a file with extension `.NDF` on your workstation). Some of this information must be consistent with other information provided with the setup of PM of OS/390 (see also “Tailoring PM of OS/390 for your Needs” on page 19-33) or with information provided on the host. Information that needs to be consistent in different locations is described in this section.

If you need to update your current Communications Manager configuration, you should first make a backup copy of your current configuration file.

You can check (and change) the information provided with your current Communications Manager configuration file on your workstation, or you can create a new configuration file with the *Communications Manager Setup* functions. This section describes the alternative of creating the minimum configuration required for APPC support.

1. After a double-click on the *Communications Manager Setup* icon from the *Communications Manager/2* folder, the system displays the following window:

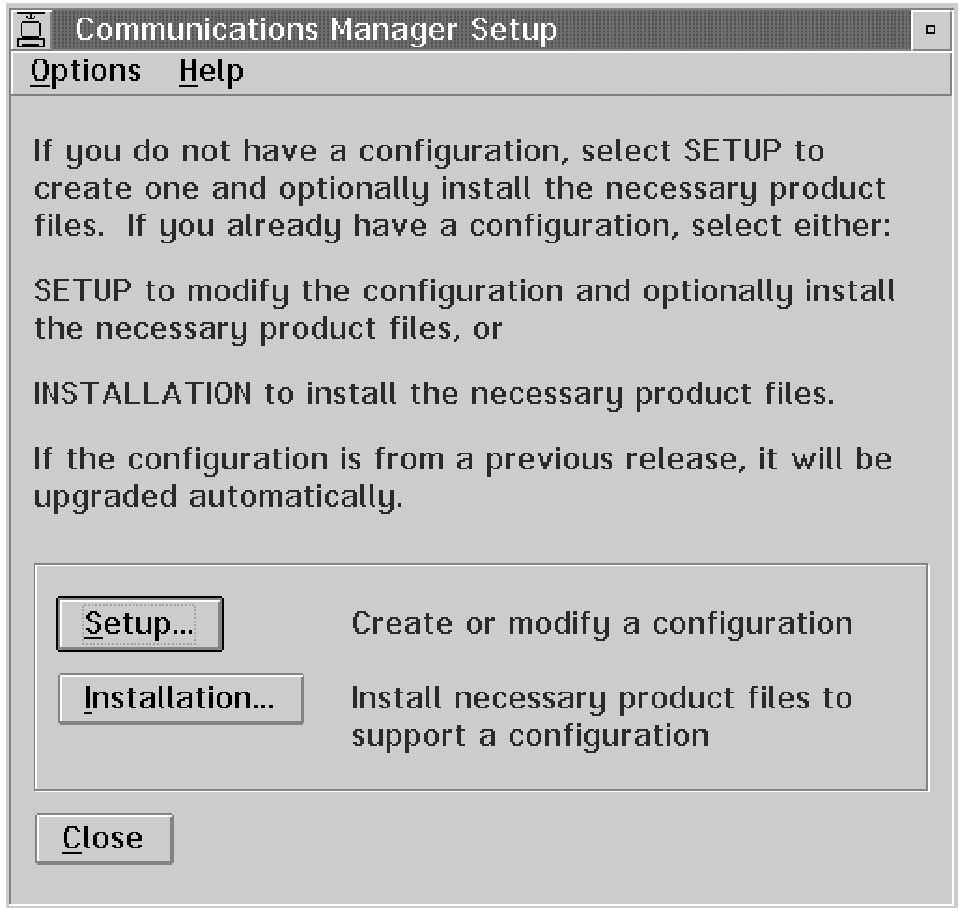

Figure 19-6. Communications Manager Setup Entry Panel

In this window, click on **Setup...** .

2. This will lead you to the following window:

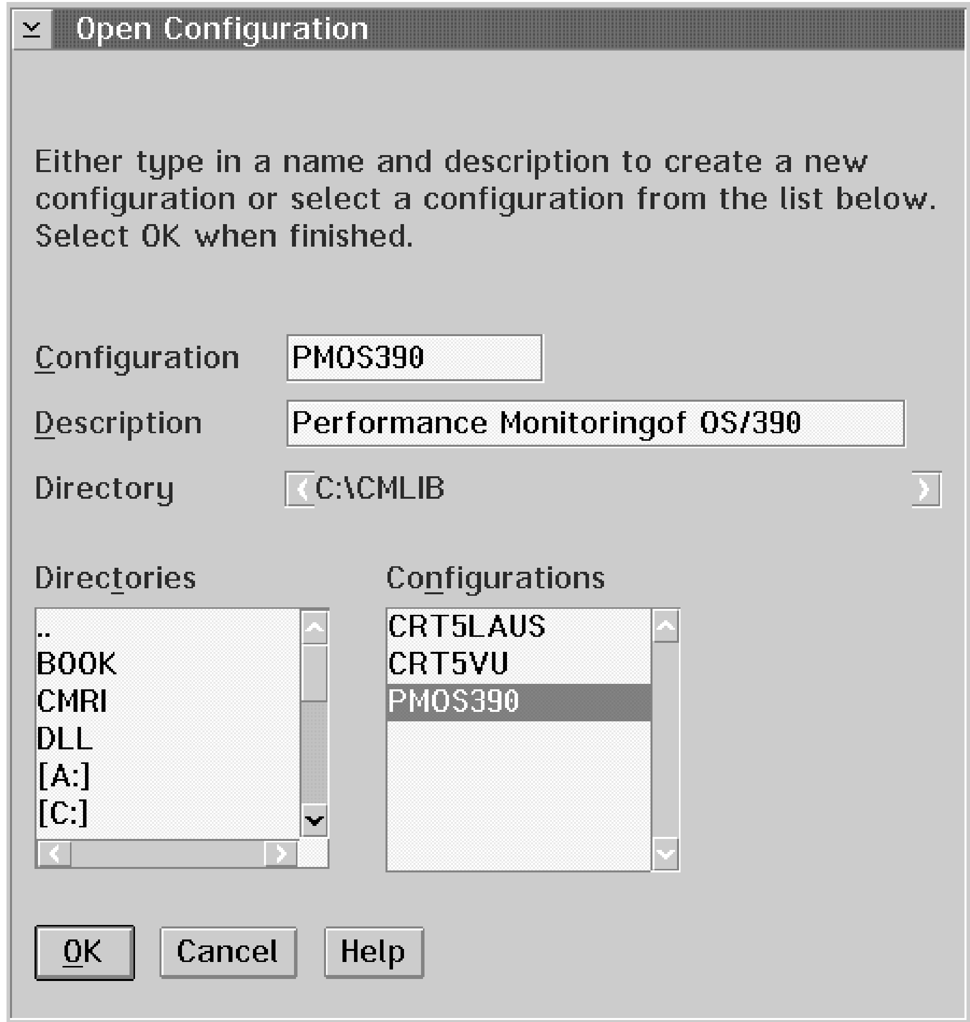

Figure 19-7. Communications Manager Setup: Open Configuration

In this window, if you want to change your existing configuration, select the name of your configuration file from the **Configurations** list. Or, if you want to create a new configuration (as demonstrated in Figure 19-7), enter the name of your new configuration file in the **Configuration** entry field. In the example, the configuration file is called `PMOS390`, so the created file will be called `PMOS390.NDF` on your workstation, `.NDF` being the extension for configuration files. You can also enter an optional description of the configuration.

In this panel, click on **OK** after providing all information. If the configuration is a new one, you must confirm that you want to create it. You may also need to confirm, on a separate confirmation window, that this configuration will be used on your workstation.

3. The next window to be shown is:

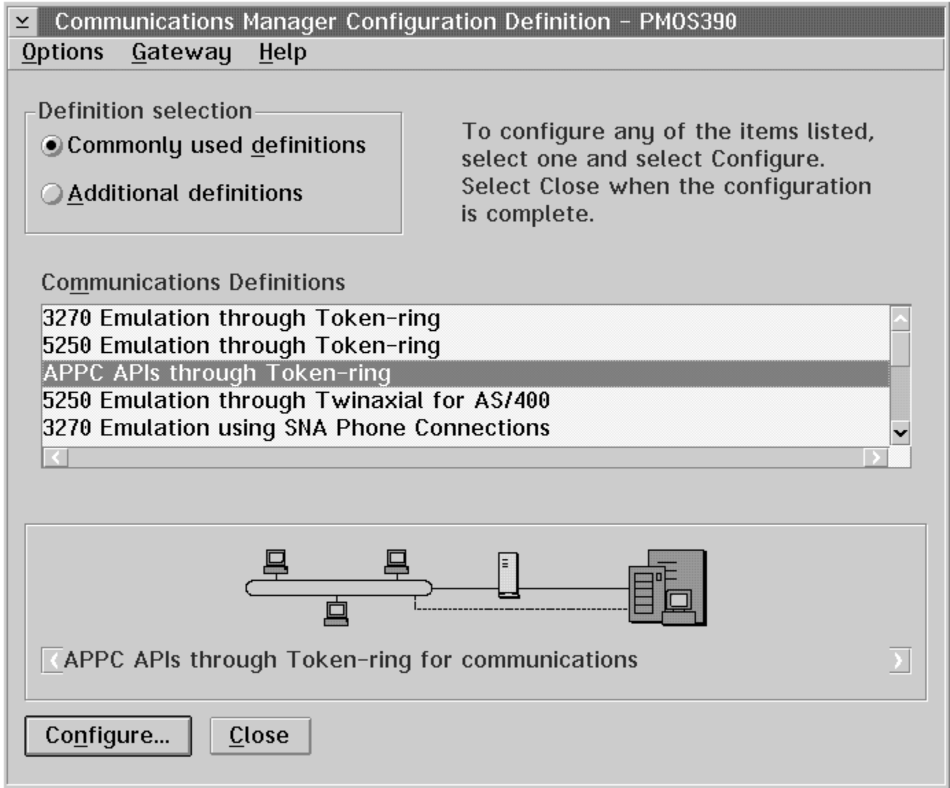

Figure 19-8. Communications Manager Configuration Definition

In this window, select radio button **Commonly used definitions**, and also select **APPC APIs through Token-ring** from the list of **Communications Definitions**. Then click on **Configure...** .

Note: If you have an Ethernet communication, you might select this as well.

4. This will lead you to the following window:

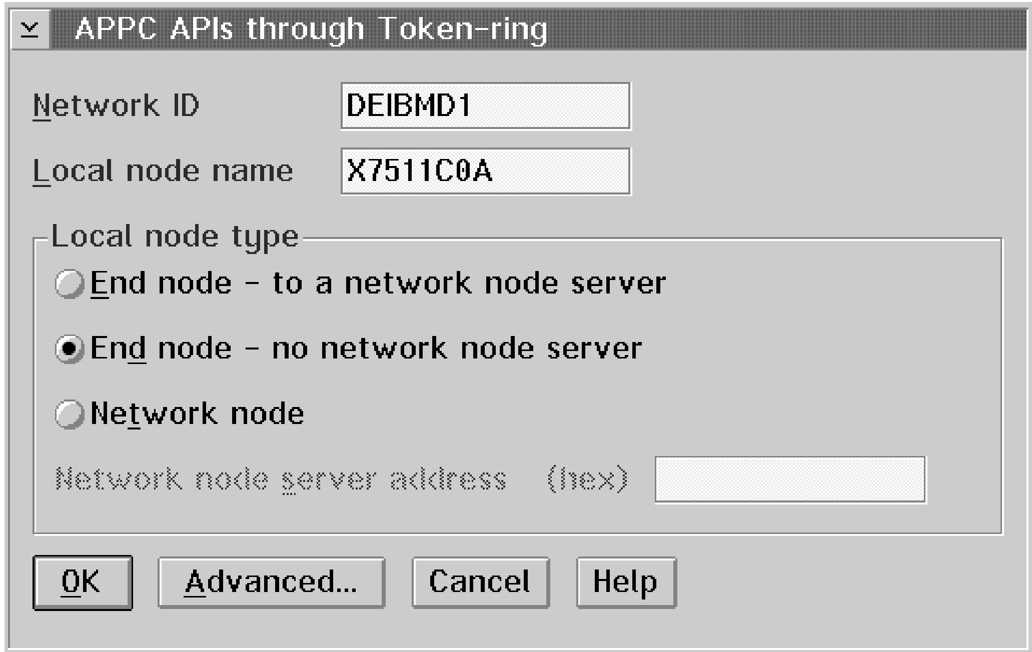

Figure 19-9. Communications Manager: APPC APIs through Token-ring

If you are modifying an existing configuration, this panel will show you the current values defined for your host communication. The **Network ID** and the **Local node name** are the address of the independent logical unit (ILU(, this is the value of the parameter FQ_CP_NAME in the `.NDF` file (see page 19-28). Verify, change or add these values accordingly. For all communications, these values are the same, so you may retrieve them from existing configurations when creating a new one, or ask your network administrator for the correct values.

Note: Make sure that your local node is enabled for APPC (ask your network administrator).

Also select the radio button **End node - no network node server**

Then click on **Advanced...** .

5. This will lead you to the following window:

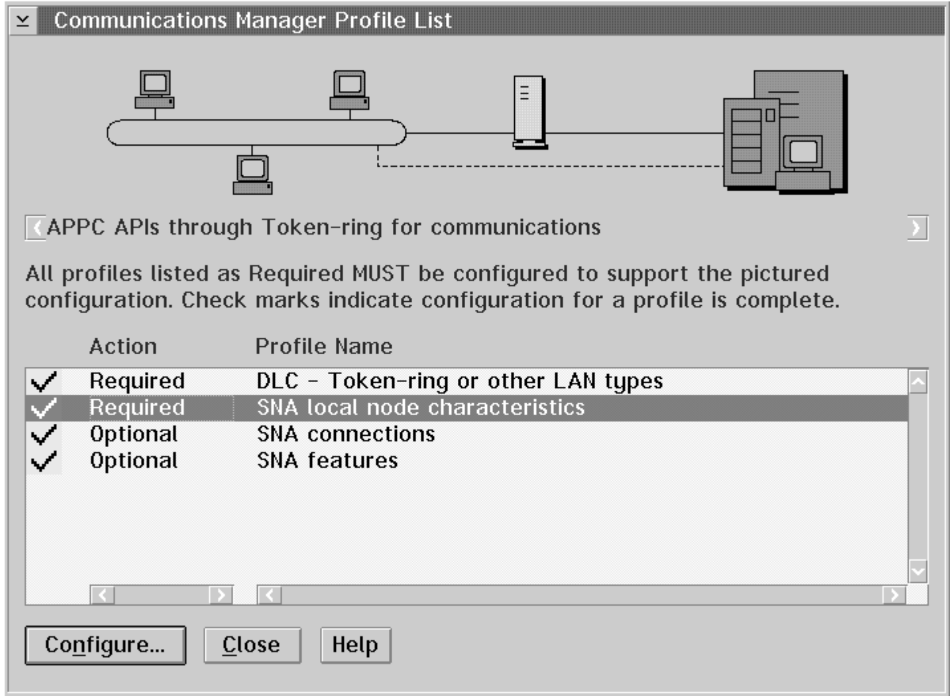

Figure 19-10. Communications Manager Profile List

In this window, select the **SNA local node characteristics** and click on **Configure...** .

6. This will lead you to the following window:

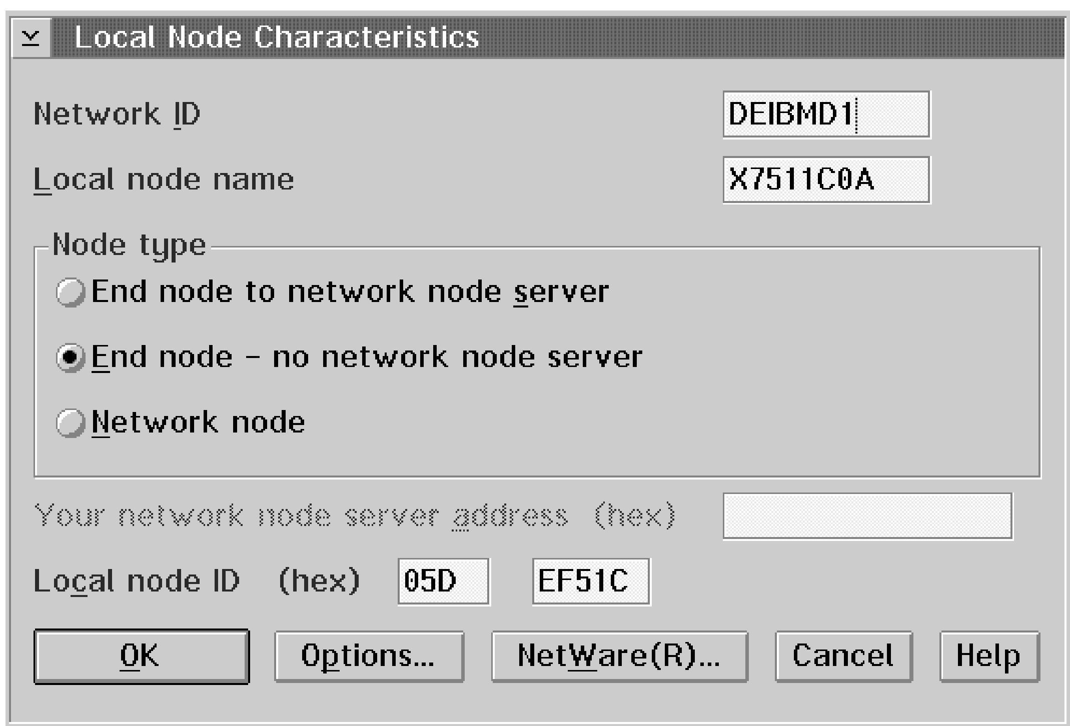

Figure 19-11. Local Node Characteristics

In this window, enter your **Local node name** within the network in hexadecimal characters. You may need to ask your network administrator for this information, or you can retrieve it from existing configurations, as it is the same for any configuration on your workstation. See parameter NODE_ID in existing configuration files, like for example, `NODE_ID(X'05DEF51C')`.

Now you have defined the required communication information for your local node. Pressing **OK** will bring you back to the window shown in Figure 19-10 on page 19-22.

7. As a next step, you need to define the communication information of your partner, in this case, your OS/390 host. Now, in the window of Figure 19-10 on page 19-22, select the **SNA connections** profile from the list and click on again **Configure...** .
8. This will lead you to the following window:

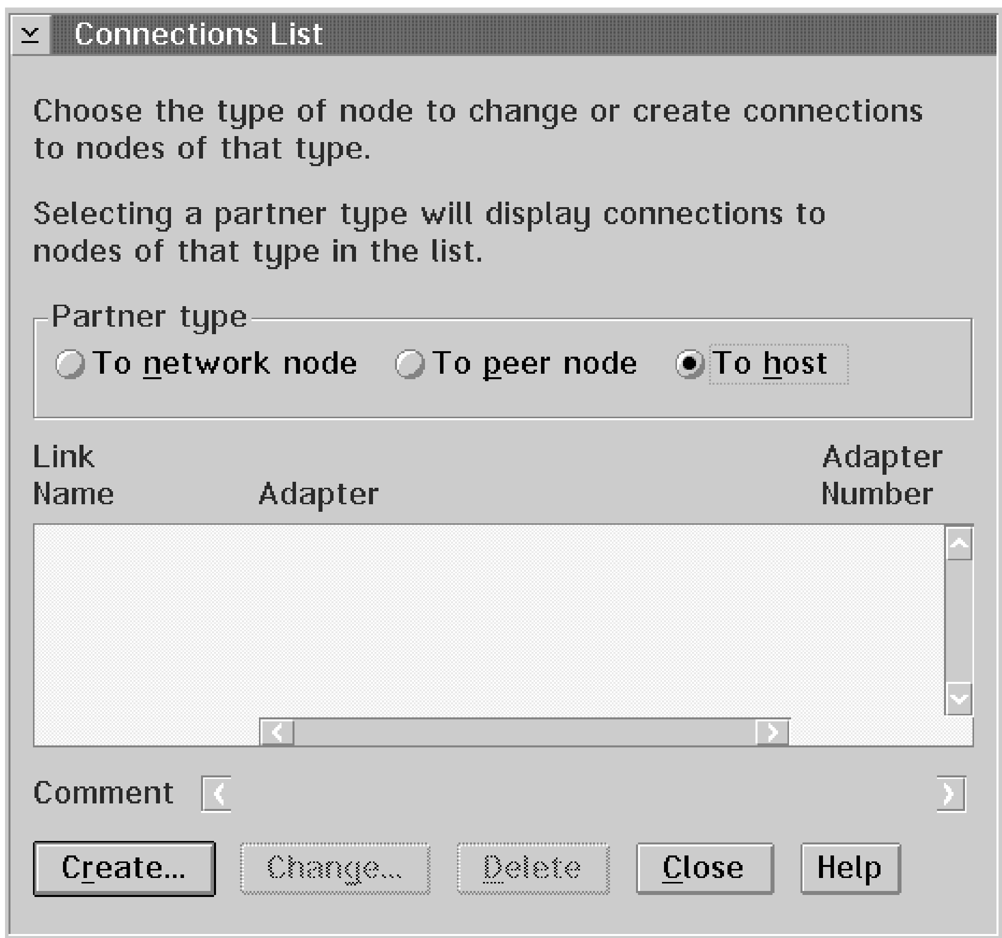

Figure 19-12. Connections List

In this window, select radio button **To host** and click on **Create...** .

9. This will lead you to the following window:

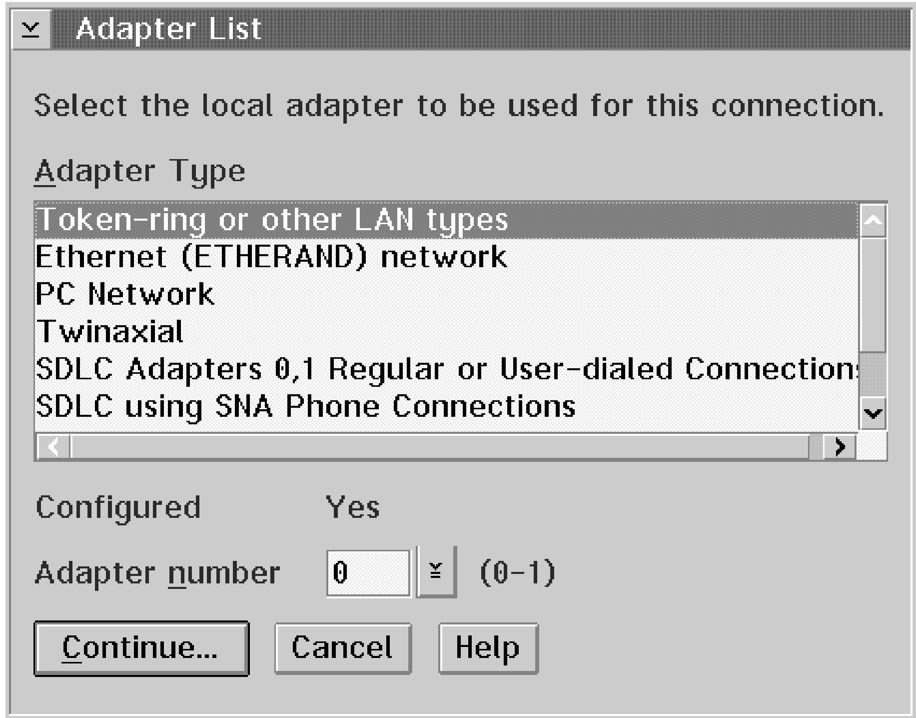

Figure 19-13. Adapter List

In this window, select **Token-ring or other LAN types** from the list of provided adapter types and click on `Continue...` .

10. This will lead you to the following window:

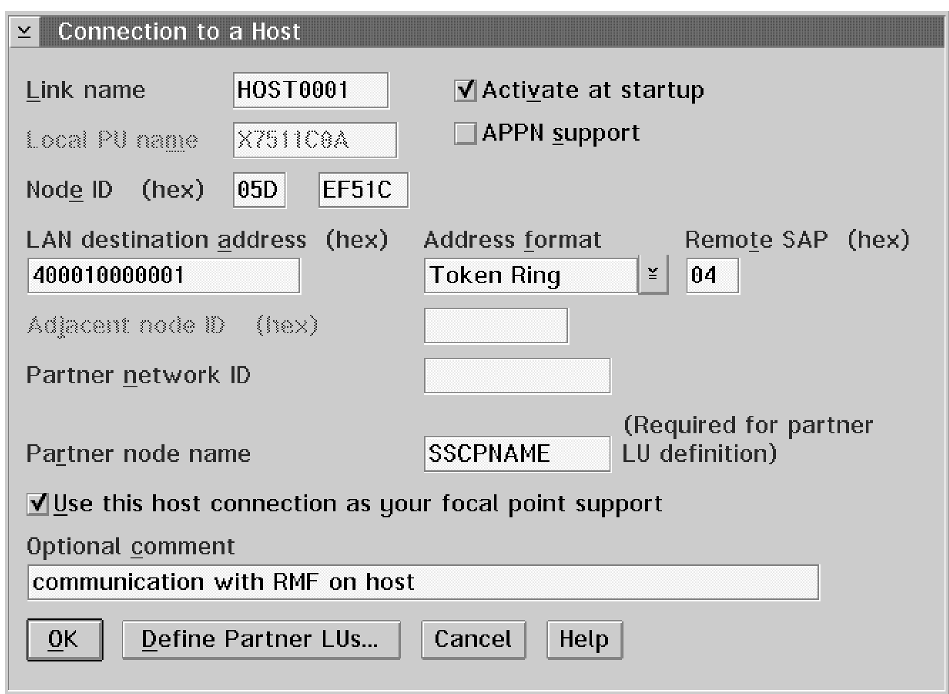

Figure 19-14. Connection to a Host

Now you need to specify a **Partner node name**. This is required to enable the `Define Partner LUs` push button. However, this information is ignored if the check box **APPN support** is not selected. In this case, enter dummy information.

After filling in all required information, click on `Define Partner LUs...`.

11. This will lead you to the following window:

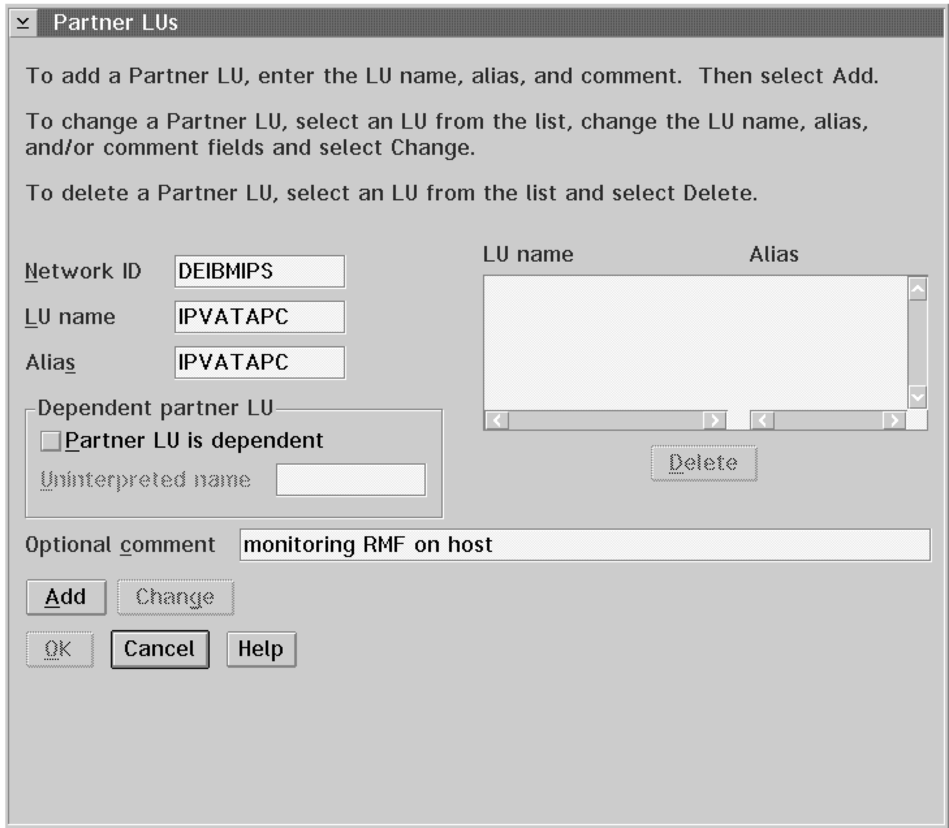

Figure 19-15. Partner LUs

In this window, you define information about the logical unit of the host with which you want to communicate. Specify this information in the entry fields **Network ID** and **LU name**. For ease of use, you can choose an **Alias** name for the fully qualified name of your host in the network. You may need to ask your network administrator for the appropriate names and IDs.

Make sure that the definition of the APPC LU name of your OS/390 system here (and therefore of the `.NDF` file, too) matches the `ACBNAME` entry in the Parmlib member APPCPMxx, like for example:

- Entry in Parmlib member APPCPMxx:

  `ACBNAME(IPVATAPC)`

- Definition in the `.NDF` file (see page 19-28):

  `FQ_PARTNER_LU_NAME(DEIBMIPS.IPVATAPC)`

Now click on `Add` to add this newly defined partner to your configuration. If this is the only partner you need to monitor, you can now click on `OK`.

12. This will bring you back to the window shown in Figure 19-14 on page 19-24, where you now click on `OK`, too.

13. Now you are back in Figure 19-12 on page 19-23. Click on **Close** here. The Communications Manager will lead you stepwise back to the windows previously opened. Close each window with **Close**, until you reach the window called **Communications Manager Profile List** (see Figure 19-10 on page 19-22). In this window, select **SNA features** and click on **Configure...** to define your local CPI communications side information.

14. This will lead you to the following window:

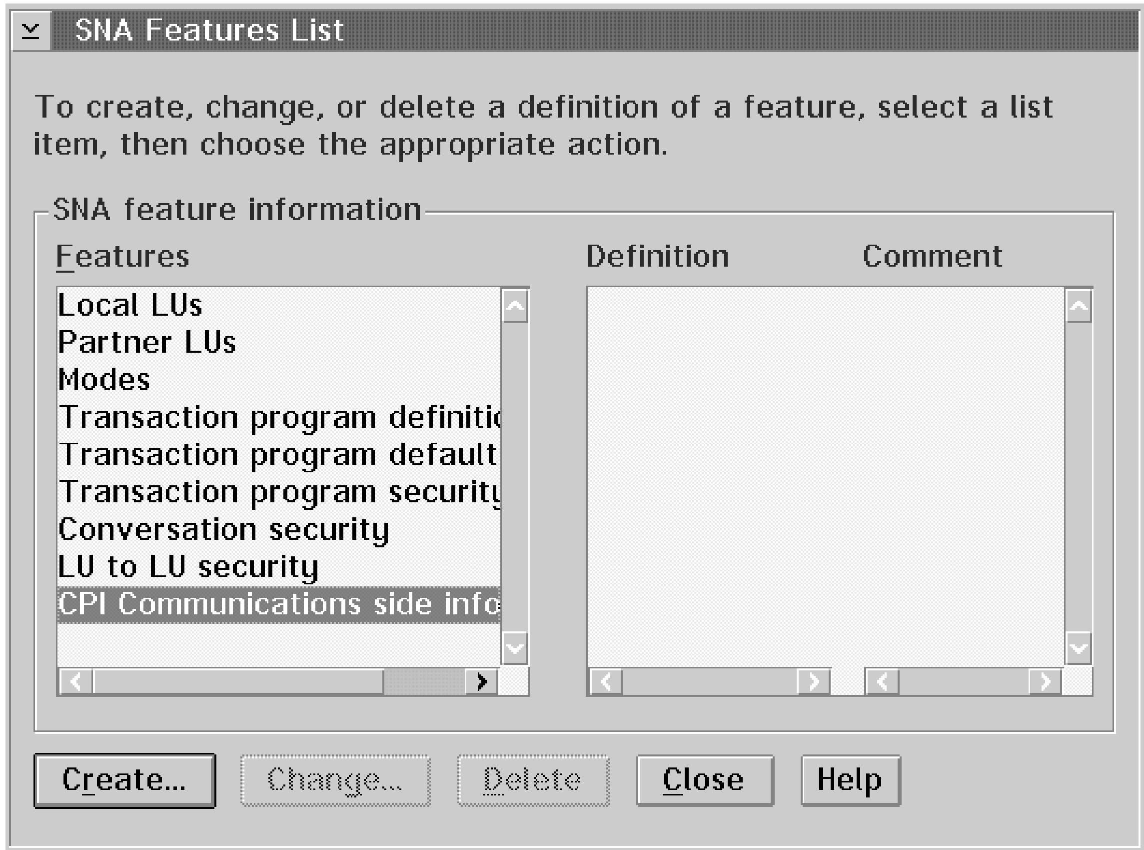

Figure 19-16. SNA Features List

In this window, select **CPI Communications side information** from the list of provided features, and click on **Create...** .

15. This will lead you to the following window:

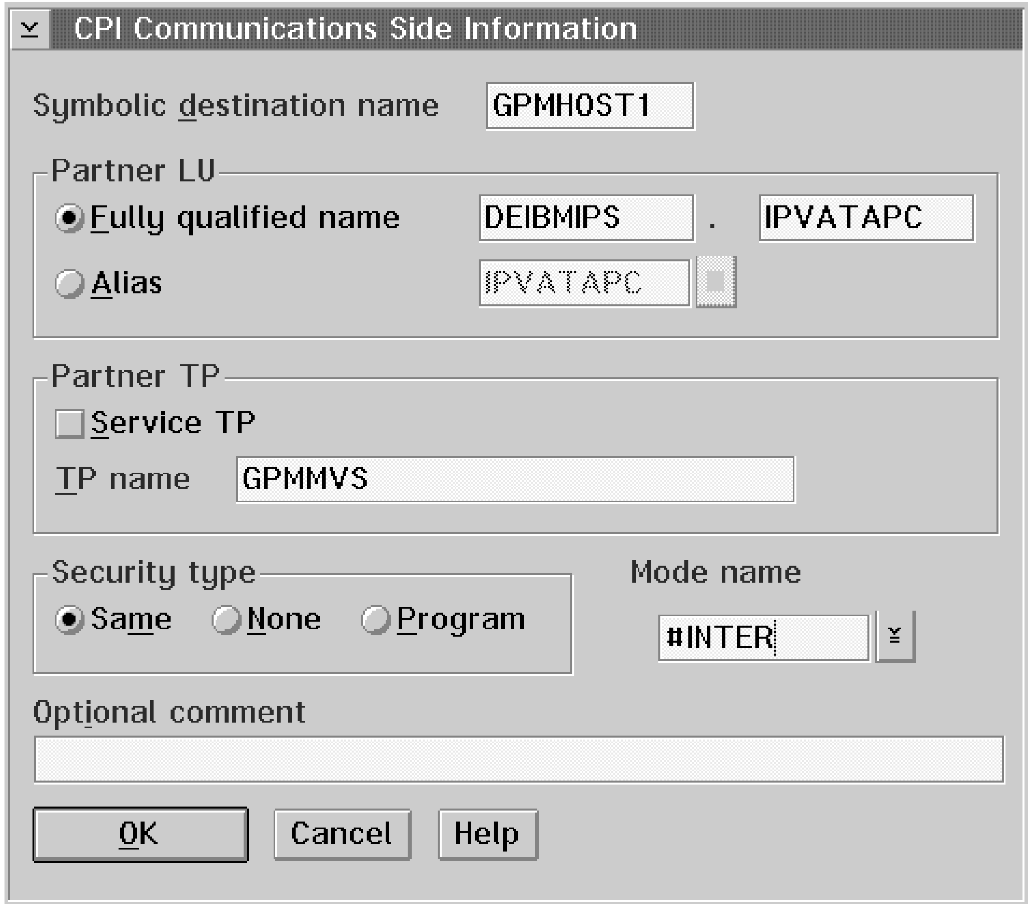

Figure 19-17. CPI Communications Side Information

In this window, specify the following values for your partner LU:

Symbolic destination name
Here you can specify a symbolic destination name of your choice for your APPC communication with the host. This name must match with the `HOST` parameter in the `GPMSETUP.INI` file (see page 19-33), and will reappear as the parameter `SYMBOLIC_DESTINATION_NAME` in the configuration file. See page 19-28.

Fully qualified name
Click the radio button in front of these entry fields and enter the values that you specified in the fields **Network ID** and **LU name** in Figure 19-15 on page 19-25.

Alias
There is an easier way to specify the fully qualified name, if you have specified alias names for the values in **Network ID** and **LU name** in Figure 19-15 on page 19-25. In this case, simply click on the radio button in front of this entry field, click on the arrow to the right of it, and select the appropriate alias name from the pull-down list.

TP name
Here you specify the name of the OS/390 TP-Profile. This name must match the parameter TPNAME in the `GPMSETUP.INI` file (see page 19-33). What you specify here will reappear as the parameter TP_NAME in the configuration file.

Mode name Click on the arrow to the right of this field and select **#INTER** from the pull down list.

When you have specified all the required values, click on **OK**.

16. Now end the *Communications Manager Setup* function by closing all open windows with **Close**. The Communications Manager will finally verify the resulting configuration file.

As an example, of what you will have produced, here is an extract of a configuration file (with extension .NDF on your workstation):

Extract of an Communications Manager/2 .NDF file - Sample

```
/* This is the address of the independent logical unit (ilu) */

DEFINE_LOCAL_CP  FQ_CP_NAME(DEIBMD1.X7511C0A )

DEFINE_CPIC_SIDE_INFO  SYMBOLIC_DESTINATION_NAME(GPMHOST1)
                        DESCRIPTION(side info for GPM TSOLAB)

/* Optional, corresponds to entry field ' ' from CM/2 - Setup facility */

                        FQ_PARTNER_LU_NAME(DEIBMIPS.IPVATAPC)
                        MODE_NAME(#INTER  )
                        TP_NAME(GPMMVS);
```

SYMBOLIC_DESTINATION_NAME
This is a name of your choice; however, it must correspond to the HOST value in the GPMSETUP.INI file (see page 19-33).

FQ_PARTNER_LU_NAME
Ask your system administrator for this name; it can also be found in Parmlib member APPCPMxx.

TP_NAME
The name of the APPC TP profile; it corresponds to the TPNAME value in the GPMSETUP.INI file.

Note: Adapt the parameters to your needs. If you use a text editor to adapt this file instead of the *Communications Manager Setup* facility, make sure that the changes made for the CPI-C Side Info are dynamically activated with the command:

```
CMVERIFY xxxxxxxx.NDF /e
```

where xxxxxxxx.NDF is the name of the .NDF file that you are modifying.

Preparation of the TCP/IP Connection

Host - Setup of the Server Start Procedure

- Make sure that TCP/IP is started on your OS/390 system.
- Modify the GPMSRVxx Parmlib member to your needs. SESSION_PORT(nnnn) is the port number of the TCPIP server on the OS/390 host (nnnn between 5000 and 9999). The port number must correspond to the TCP_PORT = nnnn definition in the workstation file GPMSETUP.INI

You find details in “Setting Up the Distributed Data Server” on page 2-10.

- To start the DDS Server, use the system command

  `START GPMSERVE,MEMBER=xx`

  xx is the suffix for the GPMSRVxx Parmlib member. Once the TCP/IP Server is started on the OS/390 system, the PM of OS/390 client is able to connect to this port.

Starting the TRACE of the DDS address space

- Start the TRACE with the following command:

  `MODIFY GPMSERVE,TRACEON`

- Stop the TRACE with the following command:

  `MODIFY GPMSERVE,TRACEOFF`

The TRACE is directed to DDNAME SYSOUT as it is defined in the GPMSERVE Proclib member.

Installation of Performance Monitoring Common Functions

Performance Monitoring Common Functions (PM Common Functions) is a prerequisite for installing PM of OS/390. If the PM Common Functions is already installed on your workstation you can switch to the information contained in “Installation of PM of OS/390” on page 19-32. This chapter describes the required steps to install PM Common Functions from the OS/390 system to your workstation.

1. Logon to TSO on your OS/390 system.
2. Make sure that your TSO session is ready for file transfer (native TSO READY mode).
3. Open an OS/2 command window and create a temporary subdirectory into which you will download the installation command file for PM Common Functions, for example, `D:\TEMP`.
4. Now you are ready to download the PM Common Functions installation command file into the directory created in step 3.

   Switch to this directory and type the following OS/2 command to download the file:

   `RECEIVE CPMINST.CMD h:'hlq.SCPMPWS(CPMINST)'`

   where **h:** is your 3270 emulation session name.

   This command will download CPMINST.CMD on your workstation.
5. Type the following OS/2 command to install PM Common Functions:

   `CPMINST`

   You will receive the following prompts for the required values. Reply to them as shown:

   - Enter the name for the 3270 emulator session or click on ENTER to accept the default (A)

     ⇒ If the session name of your emulator session is not *A*, enter the letter name here (without a colon).

- Enter the fully qualified host data set name (without quotes) or click on ENTER to accept the default (hlq.SERBPWSV)

  ⇒ The product is delivered in data set hlq.SERBPWSV. If this name was changed during installation on the host, ask your system administrator for the new name.

- Enter the action to perform: I(nstall) / U(pdate) / D(elete) or click on ENTER to accept the default (I)

  ⇒ Here you would enter *i* or just click on ENTER to install the product.

6. Now the Installation and Maintenance utility displays a window as shown in Figure 19-18. Follow the instructions of the installation utility:

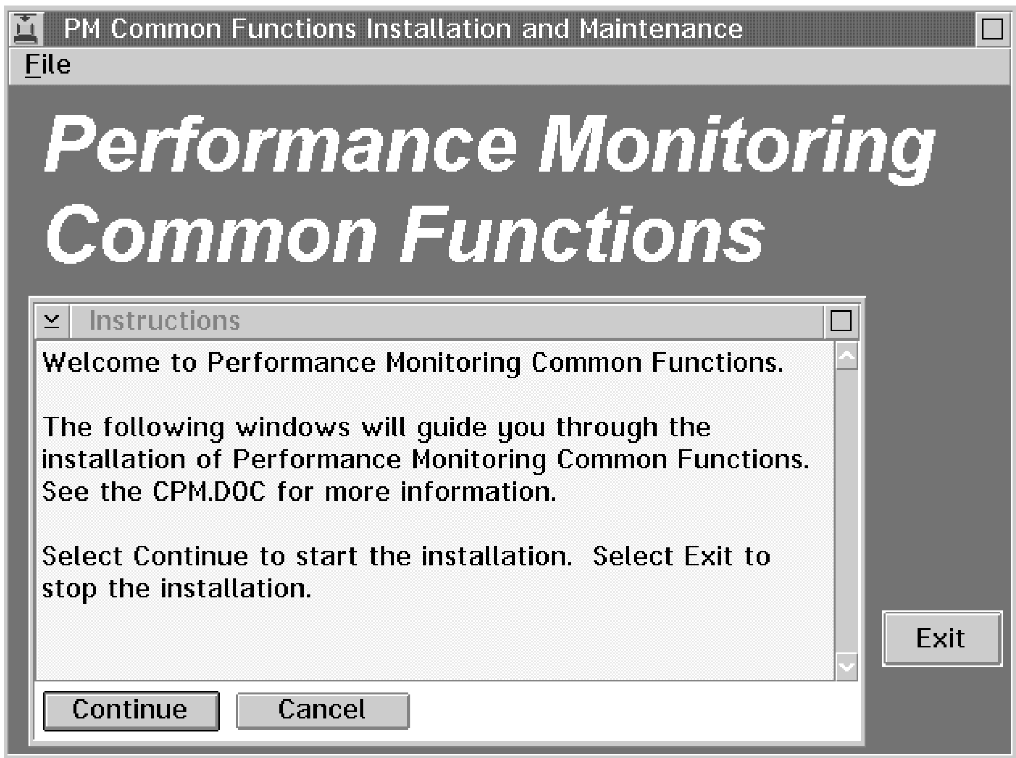

Figure 19-18. PM Common Functions - Installation and Maintenance

Click on **Continue** .

7. The next window shown is:

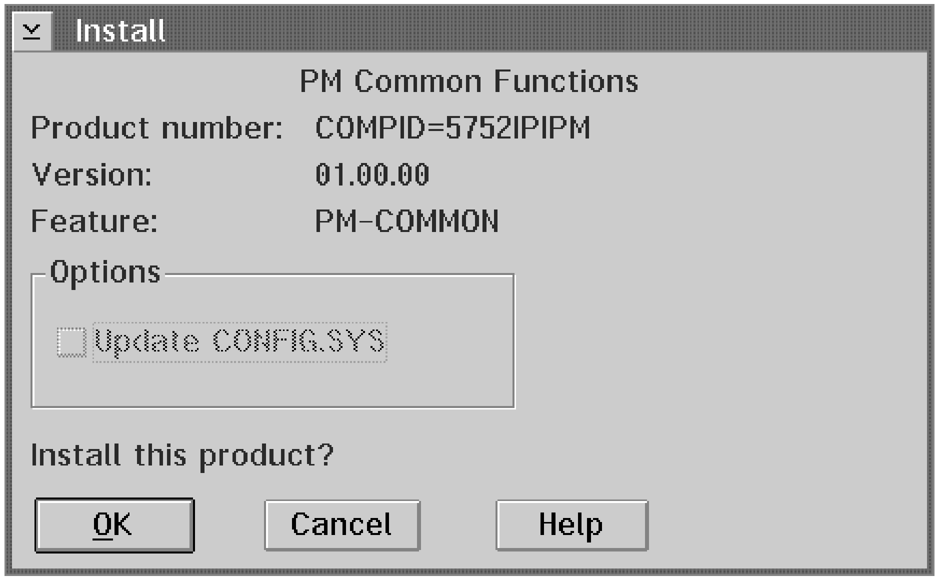

Figure 19-19. PM Common Functions - Install

Now click on **OK**.

8. The next window shown is:

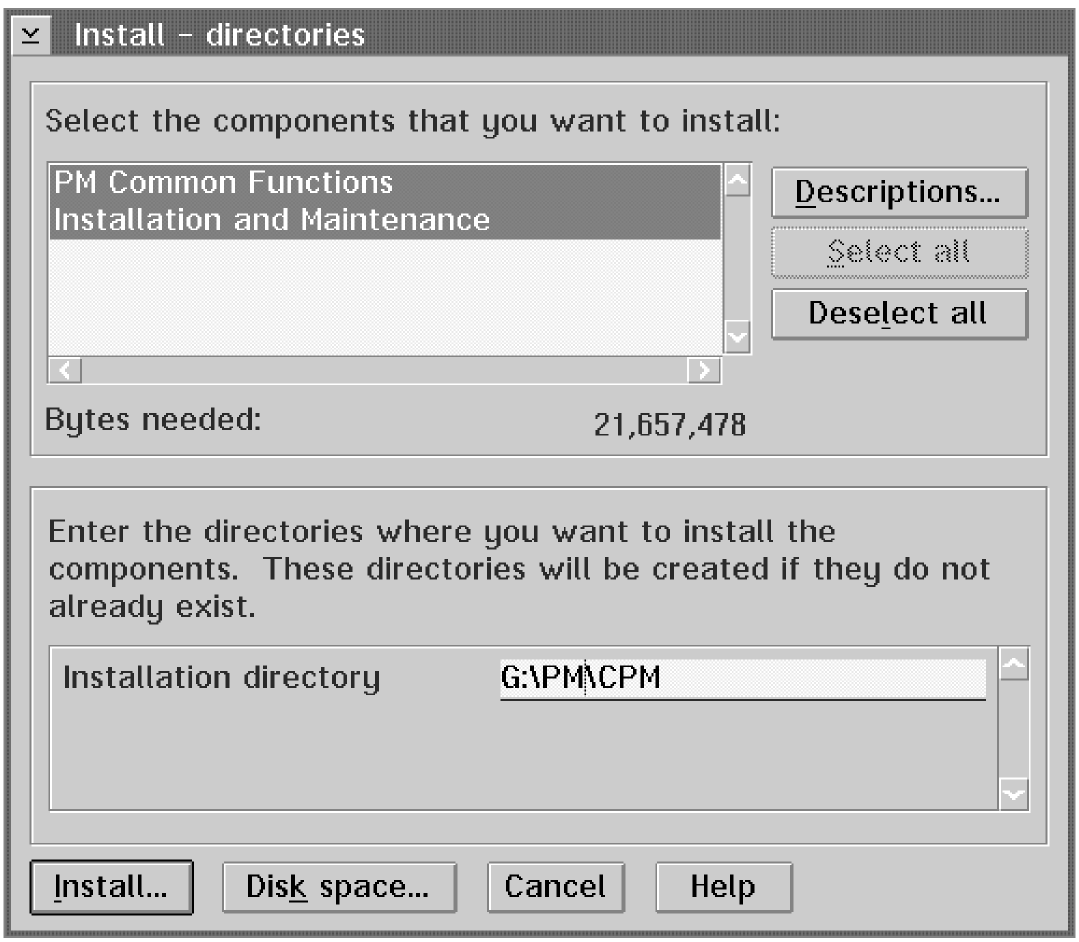

Figure 19-20. PM Common Functions - Install-directories

In this panel, click on Select all to select from the list of products offered for installation both of the components:

- Performance Monitoring Common Functions
- Installation and Maintenance

If you install the Installation and Maintenance component, you will have a handy tool for updating PM Common Functions and PM of OS/390 and applying service.

In the entry field **Installation directory** enter the path in which you want to have the product installed. Now click on Install... .

When installation is complete, the installation utility will display an appropriate message.

Installation of PM of OS/390

As PM Common Functions is a prerequisite for installing PM of OS/390, you will get an error message (see Figure 19-21) if you try to install PM of OS/390 before PM Common Functions.

Figure 19-21. Error Message if PM Common Functions is not installed

Follow these steps to install PM of OS/390 from the OS/390 system to your workstation,

1. Logon to TSO on your OS/390 system.
2. Make sure that your TSO session is ready for file transfer (native TSO READY mode).
3. Switch to a temporary subdirectory to which you want to download the PM of OS/390 installation command file, for example, `D:\TEMP`. In this subdirectory, type the following OS/2 `RECEIVE` command to download the file:

   `RECEIVE GPMINST.CMD h:'hlq.SERBPWSV(GPMINST)'`

   where **h:** is your 3270 emulation session name

   This command will download GPMINST.CMD on your workstation.
4. Now type the following OS/2 command to install PM of OS/390:

   `GPMINST`

   You will receive the following prompts for the required values. Reply to them as shown:

- Enter the name for the 3270 emulator session or click on ENTER to accept the default (A)

  ⇒ If the session name of your emulator session is not *A*, enter the correct letter here (without colon).

- Enter the fully qualified host data set name (without quotes) or click on ENTER to accept the default (hlq.SERBPWSV)

  ⇒ The product is delivered in data set hlq.SERBPWSV. If this name changed during installation on the host, ask your system administrator for the new name.

- Enter the action to perform: I(nstall) / U(pdate) / D(elete) or click on ENTER to accept the default (I)

  ⇒ Here you would enter *i* or just click on ENTER to install the product.

5. Now follow the instructions of the installation utility. Perform the same steps as described in "Installation of Performance Monitoring Common Functions" on page 19-29 for PM of OS/390.
6. When installation is complete, the installation utility will display an appropriate message. After completion, you will have a new folder named PM of OS/390 inside the **Performance Monitoring** folder on your desktop.

Tailoring PM of OS/390 for your Needs

Now, a file called `GPMSETUP.TXT` is available in your PM of OS/390 subdirectory.

1. Edit the file `GPMSETUP.TXT` and change the entries as required. They must match your APPC or TCP/IP definitions as described.
2. When you have finished your modifications, rename this file to `GPMSETUP.INI`. A sample of an `.INI` file is shown below:

GPMSETUP.INI File

```
[GPM_SETUP]
     # This is an example of GPMSETUP.INI for PM of OS/390

  COMMTIMEOUT= 60
  {
   [DOMAIN]

     NAME   = MYDOMAIN
     DESC   = "This is a description for my domain"
     USERID = MYUSERID
     HOST   = GPMHOST1
     TPNAME = GPMMVS

     # The following values are for TCP/IP only

     TCP_PORT = n

}
```

The parameters in the `GPMSETUP.INI` file have the following meaning:

COMMTIMEOUT
: Number of seconds, the workstation will wait for a response from the OS/390 host. If no response is received after this

time, then, depending on the state of the communication, either a warning message will be posted periodically to the PM message browser, or the connection to the host will be terminated, with an error message. The default timeout is 60 seconds.

This parameter is optional.

NAME The domain name shown in the 'PM of OS/390 Configuration Window' representing a SYSPLEX on a OS/390 host.

DESC Description of the domain

USERID Userid on the domain (on OS/390)

HOST Host name

APPC OS/390 symbolic destination name defined in .NDF file

TCP/IP The OS/390 host name defined in etc\host or resolved by a name server, or its IP address

The following value is for APPC only:

TPNAME The name of APPC TP profile for the domain (OS/390)

The following values are for TCP/IP only:

TCP_PORT Port number of the TCP/IP server on the 390 host.

- A value of 5000 < n < 9999 assumes that a TCP/IP server has been started on the OS/390 host.
- A value of 0 causes a TCP/IP server with a port number between 10000 and 65535 to be started on the OS/390 host, at the time the domain is opened.

The following values are required for TCP/IP if TCP_PORT = 0:

TCP_CLASS Class of job to start the TCP/IP server on the OS/390 host

TCP_ACCNT Account information for job to start the TCP/IP server on the OS/390 host

The following value is optional for TCP/IP if TCP_PORT = 0:

MAXIDLE Specifies the number of minutes after which the TCP/IP server on the OS/390 host will terminate, if no requests from the workstation were received. If the value is not specified or 0, the TCP/IP server will never terminate.

APPC or TCP/IP

Depending on the parameters TPNAME or PORT, PM of OS/390 tries to establish either an APPC connection or a TCP/IP connection:

- In case of an APPC connection, TPNAME refers to the transaction profile name in SYS1.APPCTP on the OS/390 system.
- For a TCP/IP connection, PORT identifies the common port number, which is used by the client on the PWS and the server on the OS/390 host.

3. From the OS/2 folder **Performance Monitoring** on your desktop, select the **PM of OS/390** icon. A double-click on this icon will open another folder, the **PM of OS/390 Icon View**.
4. In this folder, double-click on the **PM of OS/390 Register** icon. This will register the new PM of OS/390 application to PM Common Functions:

Figure 19-22. Registering PM of OS/390 to PM Common Functions

5. Now you can select the **PM Common Functions Start Here !** icon from the **Performance Monitoring** folder. Double-click on this icon to start the application.

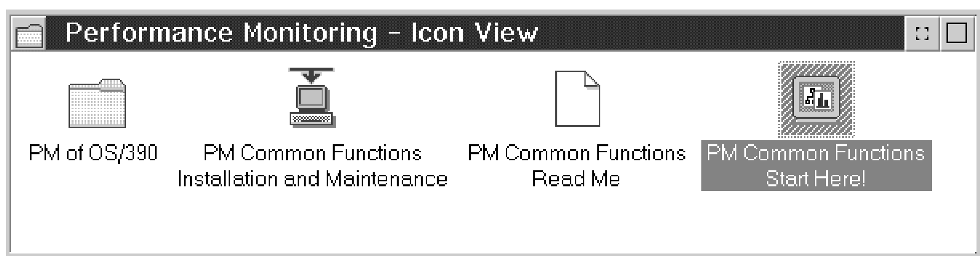

Figure 19-23. Icon View of PM of OS/390

Applying Service to PM of OS/390

With the installation utility, you can also update and apply service to the products PM Common Functions and PM of OS/390. This is the preferred method for doing this.

However, if the **Installation and Maintenance** icon, or the whole PM folder, have for some reason vanished from your desktop, use the command line interface described below.

Note:

Prerequisite for this command line interface is the existence of the data sets **CPMINST.CMD** and **GPMINST.CMD**, which you can receive from the host with:

```
RECEIVE CPMINST.CMD h:'hlq.SCPMPWS(CPMINST)'
RECEIVE GPMINST.CMD h:'hlq.SERBPWSV(GPMINST)'
```

where **h:** is your 3270 emulation session name.

If you delete files manually from the directory, you should afterwards use the Installation and Maintenance utility to ensure a clean deletion of the products.

When deleting the products, you still need connection to your OS/390 host, because certain information is retrieved from it.

Switch to the subdirectory in which you installed the product, and enter the appropriate command:

- To update PM Common Functions, enter:

  `CPMINST h hlq.SCPMPWS; u`

- To delete PM Common Functions, enter:

  `CPMINST h hlq.SCPMPWS; d`

- To update PM of OS/390, enter:

  `GPMINST h hlq.SERBPWSV u`

- To delete PM of OS/390, enter:

  `GPMINST h hlq.SERBPWSV d`

In each of these commands,

- **h** is your 3270 emulation session name (without colon).

If you enter the commands without parameters, you are prompted for the values as described in step 5 in the list in “Installation of Performance Monitoring Common Functions” on page 19-29.

Note:

Whenever you update Performance Monitoring Common Functions, the PM of OS/390 icon will disappear from the OS/2 folder **Performance Monitoring**. To restore this icon, you also need to update PM of OS/390:

`GPMINST h hlq.SERBPWSV u`

Recommendations for Working with PM of OS/390

Data Consistency in the Sysplex

In many cases, performance data of PM of OS/390 will come from multiple MVS images. To get the complete spectrum of data in a sysplex (corresponding to a *domain*), the following recommendations should be implemented:

- All participating MVS images should have the RMF gatherer active, because data can only be reported for systems with an active RMF Monitor III Data Gatherer.
- All participating systems should use the same mintime, since any reported data is based on multiples of the respective mintime. If, for example, one system uses a mintime of 100 seconds and another a mintime of 80 seconds, then data requested with a range of 100 seconds would be provided correctly for the first system, whereas data for a range of 160 seconds would be provided for the second. This could lead to difficulties in interpreting the data.
- All participating systems should use the same cycle time. Otherwise data gathered with varying "precision" will be mixed.

If several sysplexes are to be managed with PM of OS/390, the recommendations given above apply to all MVS images in all sysplexes.

Availability of Performance Data

- PM of OS/390 can provide performance data for the I/O subsystems only if the Monitor III gatherer option IOSUB (default) is active to enable the collection of data about the I/O subsystem configuration. The I/O data is available for all processors, just the 4381 systems are not supported.
- Performance data about common storage, as it is reported in the Monitor III reports STORC and STORCR, is available only if VSM common storage tracking is active. This can be defined by issuing:

  ```
  SET DIAG=01
  ```

  The defaults in the Parmlib member DIAG01 are:

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

- In a VM guest environment, no IOQUEUE data is available.

Running PM of OS/390

You start the application from the PM Common Functions folder:

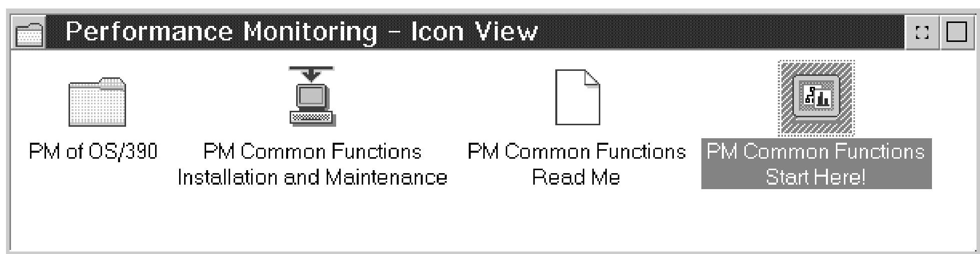

Figure 19-24. PM Common Functions Folder

You get the main window containing all applications that are registered to PM Common Functions.

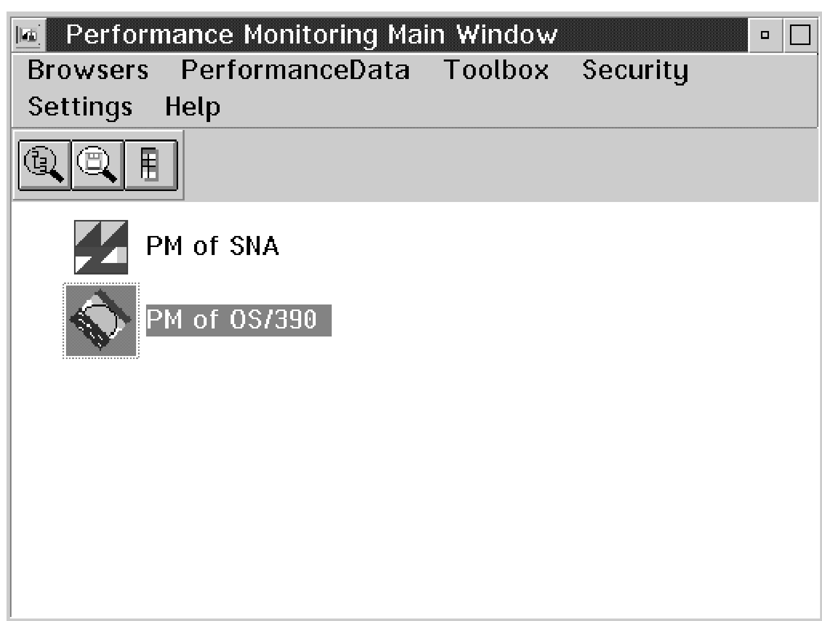

Figure 19-25. Performance Monitoring Main Window

Start-up Problem

As described in "Installing PM of OS/390" on page 19-13, there is a specific requirement for the service level of your OS/2 system:

- OS/2 3.00 Warp (family) with FixPack 22 (CSD214): XR_W022 (where '_' is a language code, if non US-version of OS/2).

If you have not installed at least this service level on your workstation, you will not be able start PM of OS/390. You can check the service level with the command

```
SYSLEVEL
```

Part of the information that you get is the following:

```
SOM Event Management Framework
Current CSD level: SM200xx
```

If xx does not have at least the value of **12**, then you should install the required service package. You might contact your IBM Support Center to get the package.

First Round - Creating a DataView

In the above example, PM of OS/390 is one of two applications. The other one is PM of SNA which is another product to be ordered separately. You can start PM of OS/390 by a double-click, which leads you to the domain configuration. A domain is a sysplex of one or more MVS images, to which there is an APPC connection from the managing workstation.

In general, any performance monitor integrated in PM Common Functions can manage one or more domains, and each domain can contain one or more agents. An agent is a provider of system management information for a certain set of resources. In the case of PM of OS/390, there is one agent per domain, which means that the **PM of OS/390 Domain Configuration** window will have one domain icon for each sysplex.

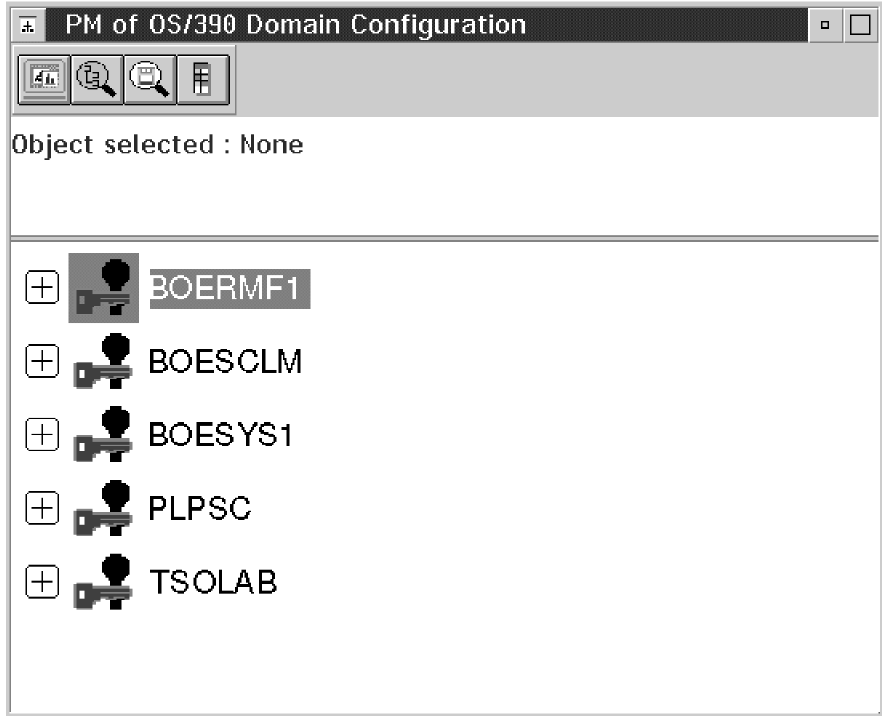

Figure 19-26. Domain Configuration

This domain configuration is the graphic representation of the definitions in your `GPMSETUP.INI` file:

GPMSETUP.INI

```
[GPM_SETUP]
  COMMTIMEOUT = 60
  {
     [DOMAIN]
     NAME   = TSOLAB
     DESC   = TSOLAB
     HOST   = GPMSYS01
     TPNAME = GPMBWSC
     USERID = BWSC

     [DOMAIN]
     NAME   = PLPSC
     DESC   = PLPSC
     HOST   = GPMSYS02
     TPNAME = GPMBWSC
     USERID = D09RMF1

     [DOMAIN]
     NAME   = BOESYS1
     DESC   = BOESYS1
     HOST   = GPMSYS03
     TPNAME = GPMBHBE
     USERID = BHBE

     [DOMAIN]
     NAME   = BOERMF1
     DESC   = BOERMF1
     HOST   = GPMSYS04
     TPNAME = GPMBWSC
     USERID = BWSC

     [DOMAIN]
     NAME   =  BOESCLM
     DESC   =  SCLM-System
     HOST   = GPMSYS05
     TPNAME = GPMBWSC
     USERID = BTEU
  }
```

When you have decided which sysplex you are going to work with, you can expand the corresponding domain by clicking the + box in front of the domain icon. This causes a logon to the sysplex under the user ID that you have defined in the .INI file.

Note: If you double-click on the icon or the name of the domain, this **does not** open the domain, but creates a new *default domain* that overlays the PM of OS/390 domain configuration showing only the selected domain. To get back, you have to move (or close) the window with the default domain.

Please ensure that you provide a valid password during logon. If your password is expired, you must specify a new password during standard TSO logon before you can proceed with logging on to the sysplex in the current PM of OS/390 environment. The successful logon is indicated by the updated domain configuration — now you can expand the display of the configuration as you wish, by clicking on the + and – boxes that are familiar from other application on your workstation.

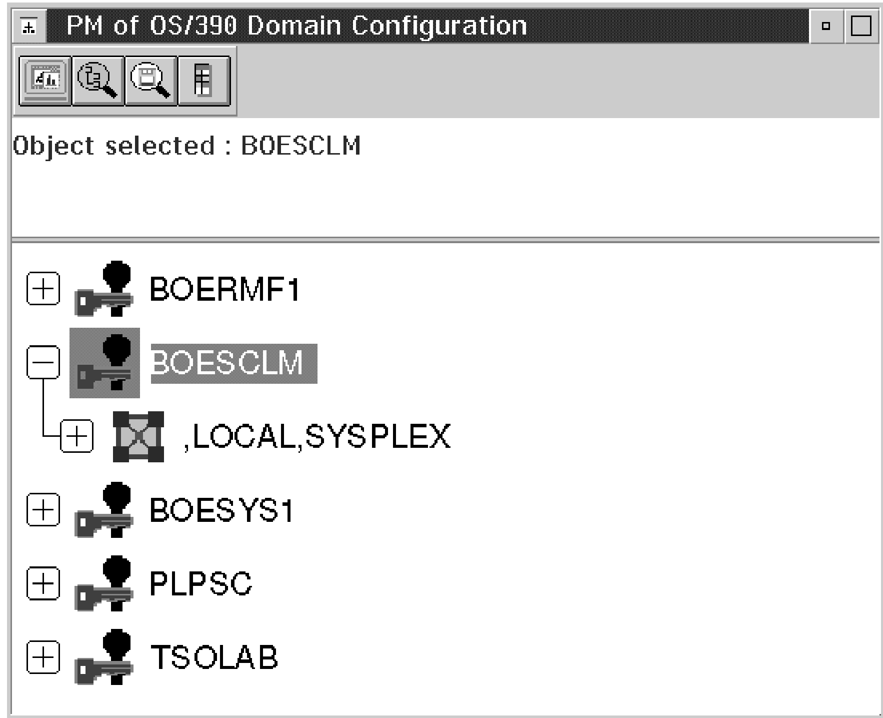

Figure 19-27. Domain Configuration - Successful Logon to a Sysplex

What is this good for?

There are two answers to this question:

1. You can see the current configuration of the sysplex - you get the information about the MVS systems belonging to the sysplex, and their I/O configuration.

   Normally, this is static information that does not change too often, but each change is reflected directly because the actual values are read from the sysplex.

2. You are interested in getting performance data for a specific resource, for example,

   - the processor
   - the storage
   - the workload

   You can expand the configuration so as to display the resource you are interested in, and then request the data, which will be provided as a DataView.

You might have expanded the configuration as in this example:

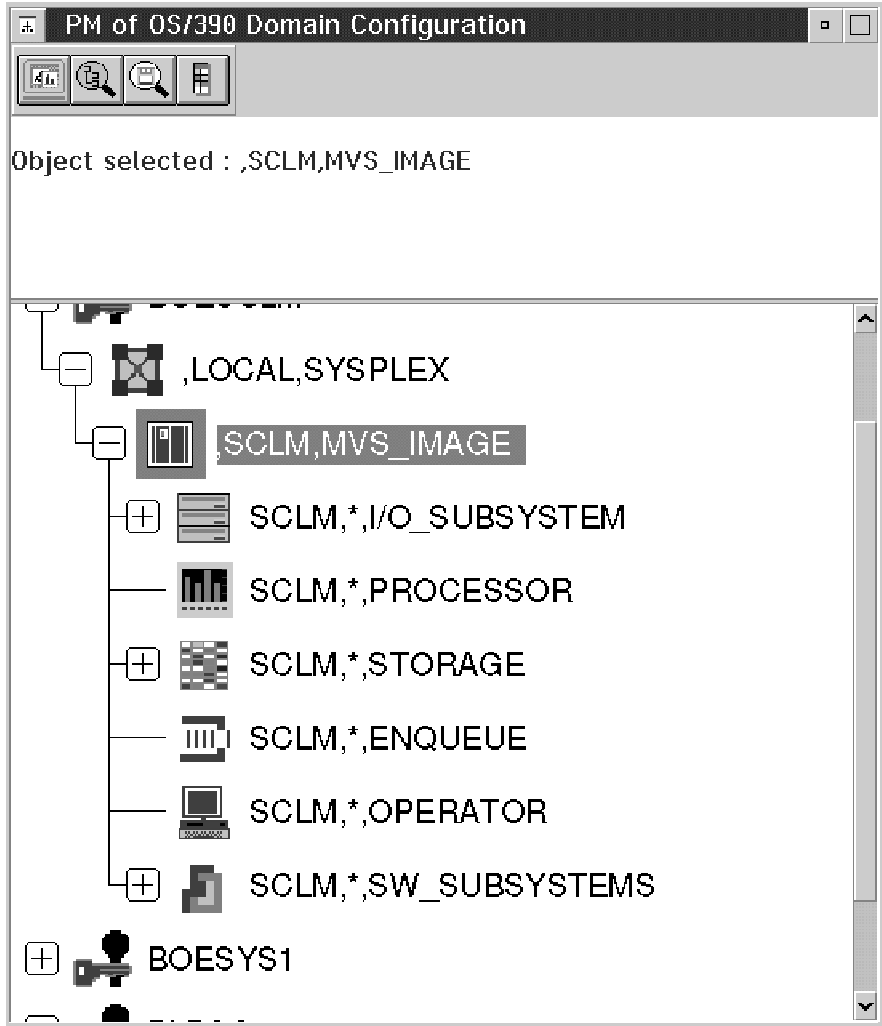

Figure 19-28. Domain Configuration - Expanded Sysplex Configuration

Assuming you want to get data on the system level, you would select the resource `MVS_IMAGE`. Clicking on the right mouse button leads to the context menu offering you **Add series...**. If you select this, you get the panel for defining an object DATAVIEW. The name that you specify here will become the title of the new DataView. If you leave the field blank, the DataView is named **Default DataView**.

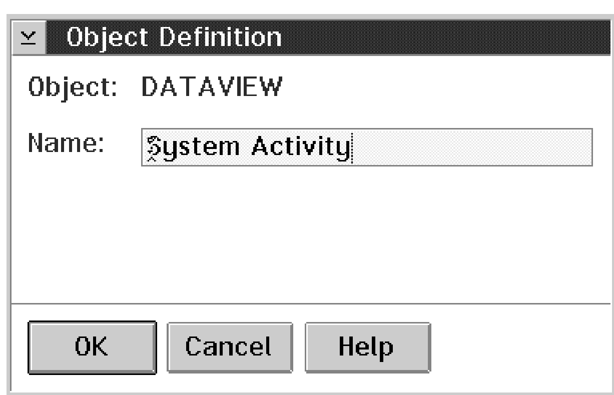

Figure 19-29. Object Definition - DataView

The next panel lets you define one or more series to build up the DataView with the performance data you are interested in.

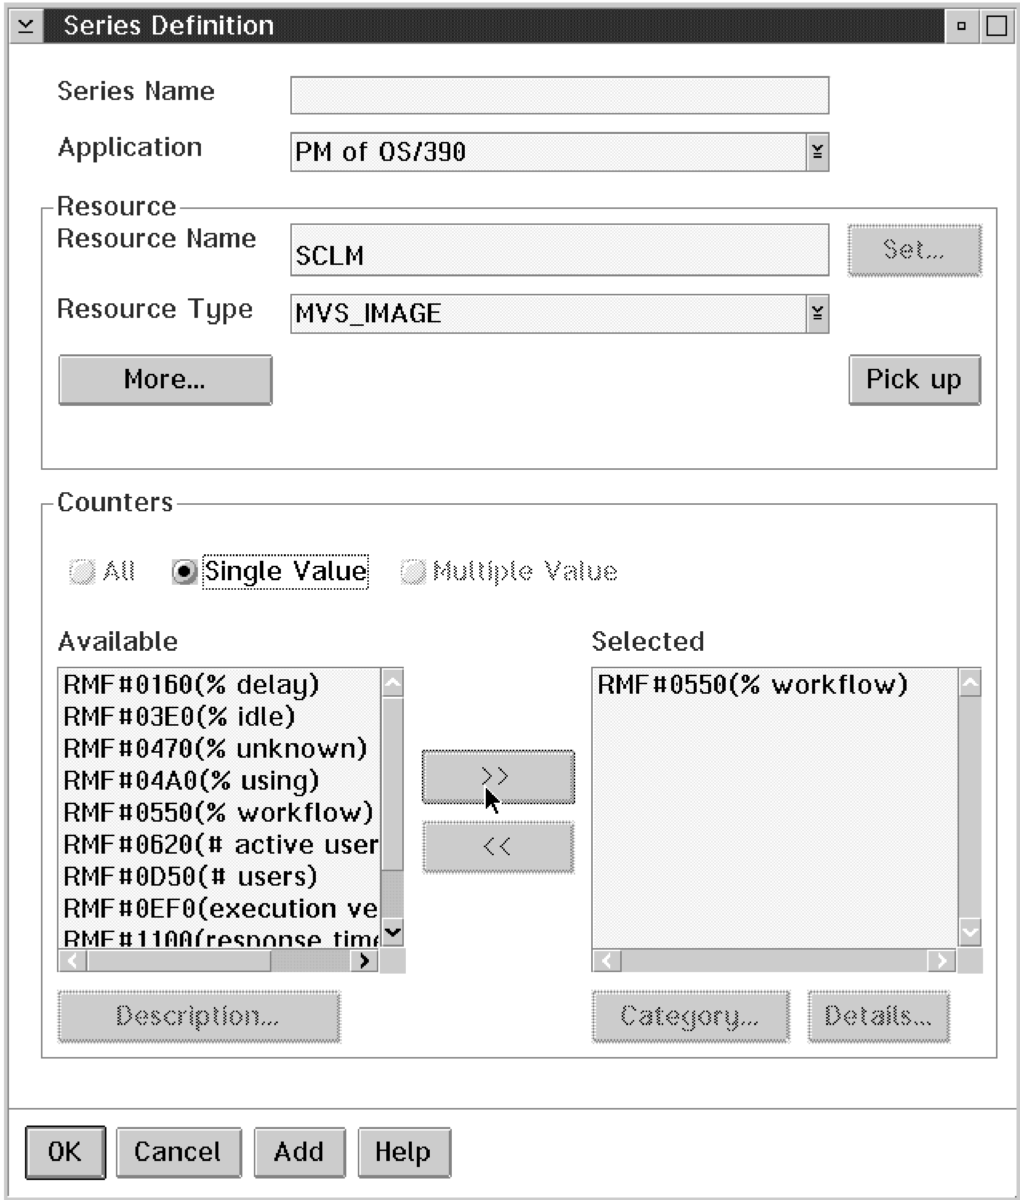

Figure 19-30. Object Definition - DataView

Automatically, some fields in this panel have default values based on your current selections (for example, resource name and resource type), but you can overwrite them if you want to. If you still want to proceed with getting data for the resource type MVS_IMAGE, you can now concentrate on the counters.

As described in Counters, there are two types of counter:

- Single-valued counters (for example, % delay, % using)
- Multi-valued counters (for example, % delay by job, % using by job)

Initially, all the counters are displayed, and you can make your selection. If you know which type you want, you can specify that only those counters should be shown.

Please make your selection ...

Do this by clicking one or several (use the CTRL-key) counters in the **Available** box.

Note: If you have selected only one counter, you will see that the **Description...** key is active, and you can get help information for this counter.

Now, there are two possibilities:

1. The >> key is active (black) - your selection is valid. When you press the key, the counters of your choice appear in the **Selected** box.
2. The >> key is not active (gray) - your selection is not valid. You have combined single-valued counters and multi-valued counters, and this is not allowed.

In the above example, you find that one counter (% workflow) has been selected. You can now press **OK**, if this is to be the only counter in the DataView, or press **Add** to add the counter and continue the series definition.

When you have selected a single-valued counter, the left box now shows single-valued counters only, because only those are valid for the DataView you are defining.

If you have decided that the DataView should have a second counter (% using), you can define it in the same way. If this concludes the definition, leave the panel with **OK**, and you will get the DataView **System Activity**. The legend tells you about the counters that you have selected. The question marks are the place-holders for the measurement data you are looking for.

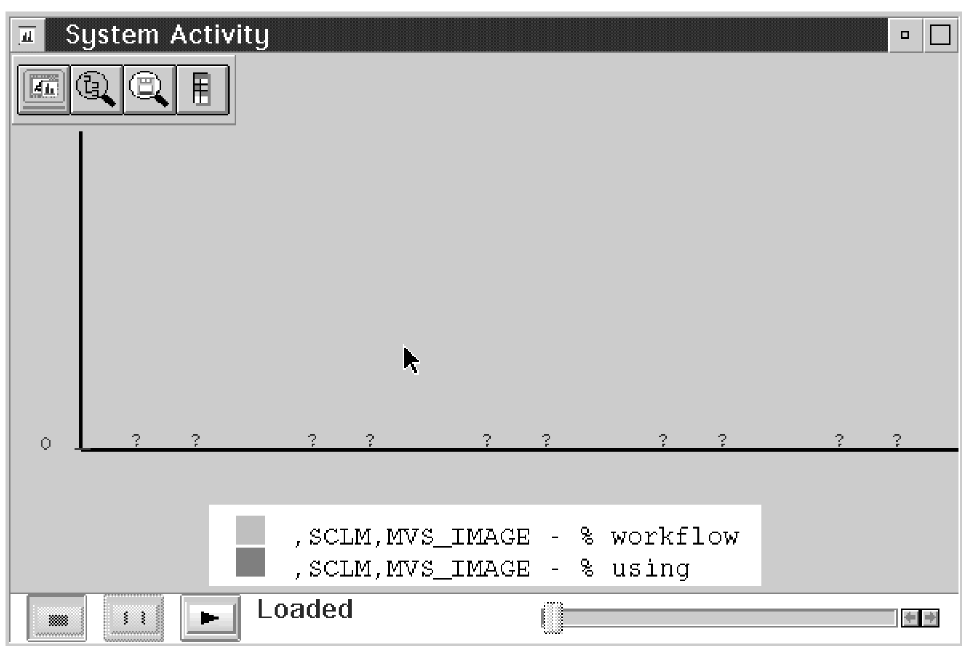

Figure 19-31. DataView - Loaded

By pressing the Start key, you start data reporting with PM of OS/390. The example shows the data received from the Monitor III data gatherer after five measurement intervals.

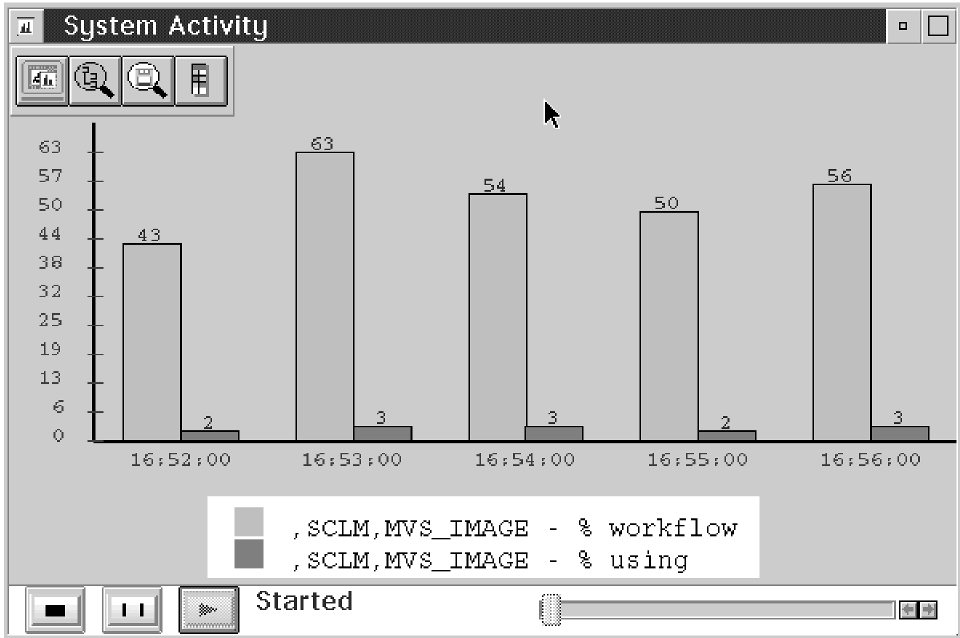

Figure 19-32. DataView - Started

If you would like to save this DataView for further use, you can do so via the context menu, which leads you to this panel:

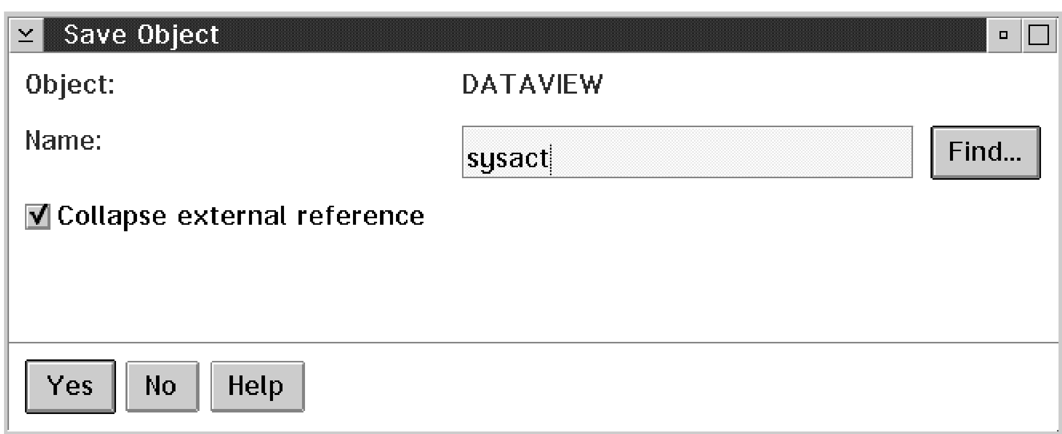

Figure 19-33. DataView - Save DataView

When you specify the name for the DataView, you will be informed that the name of a DataView needs an extension of `.do`. This will be appended automatically.

Figure 19-34. DataView - Save DataView

You should select a data set name that gives a clue to the contents of the DataView.

Working with PerfDesks

In First Round - Creating a DataView you saw how to logon to a domain and how to define, start, and save a DataView. This method can be used either to define a permanent DataView for further processing, and to create a temporary DataView for displaying specific performance data. In this section, we will concentrate on understanding how to use PerfDesks together with DataViews.

In the Performance Monitoring Main Window, you can select **Browsers** from the action bar with several choices in the pull-down menu. We will start by selecting the **Active Objects**.

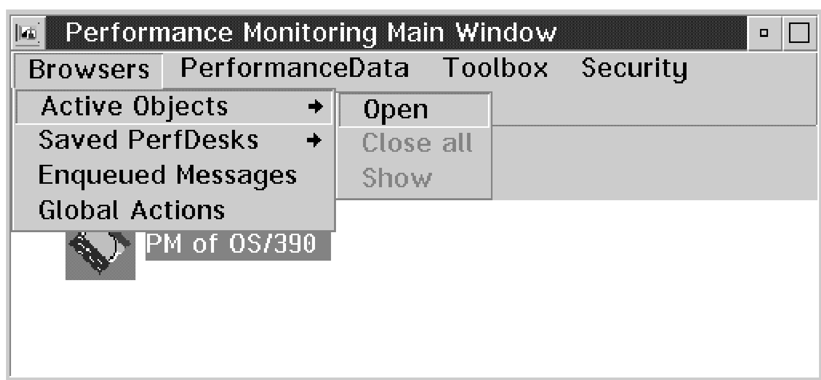

Figure 19-35. Performance Monitoring Main Window - Browsers

This brings up the window for the Active Object Browser.

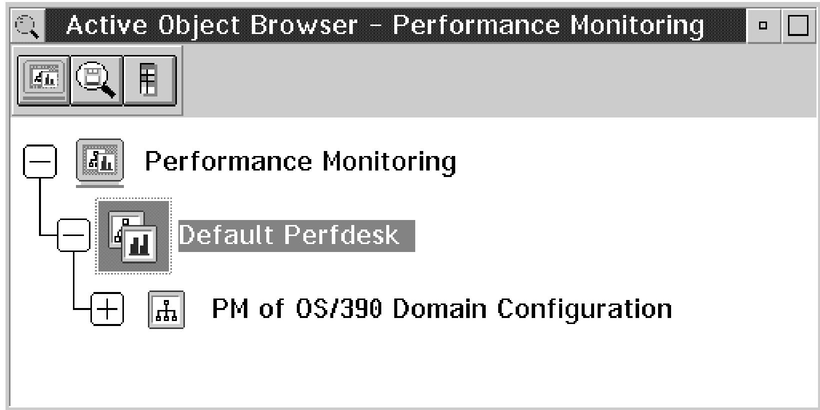

Figure 19-36. Active Object Browser

Currently, there is one PerfDesk available: the Default PerfDesk, containing the PM of OS/390 Domain Configuration. You can use this PerfDesk to load an existing DataView. As usual, all available functions for an object can be seen in the context menu:

Figure 19-37. Load DataView

This functions offers a selection window showing all files in the current directory with the extension `.DO`.

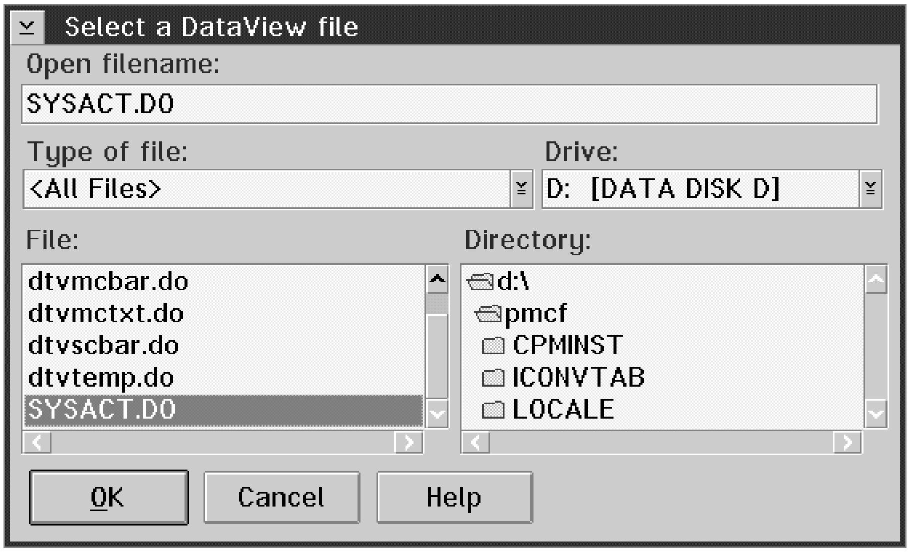

Figure 19-38. Select DataView File

When the appropriate DataView has been selected, it is loaded and you can start it as described earlier.

You can see that the window of the Active Object Browser contains the updated information, including the recently loaded DataView called 'System Activity'.

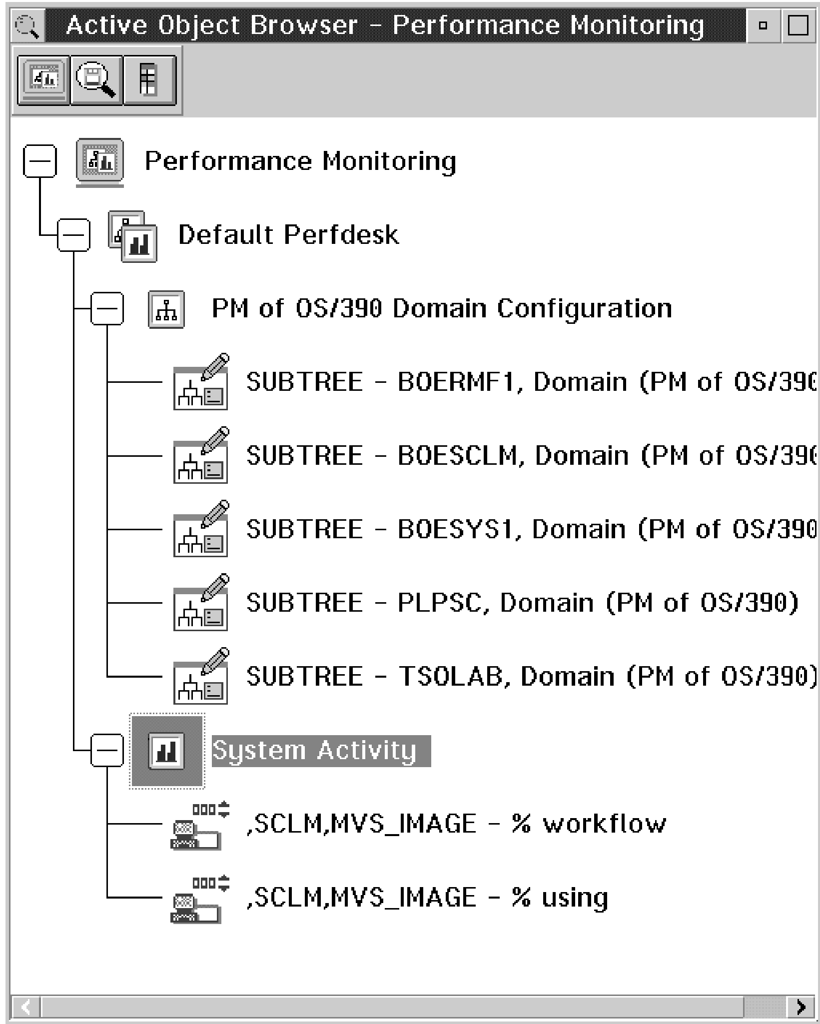

Figure 19-39. Active Object Browser

As with any other object, you can save the Default PerfDesk, but this creates a second Default PerfDesk, with just a different file name. Therefore, we recommended you to create your own PerfDesk with all the DataViews that you want to use on a more or less regular basis.

Creating a PerfDesk

You can create a new PerfDesk from the context menu of the Performance Monitoring application:

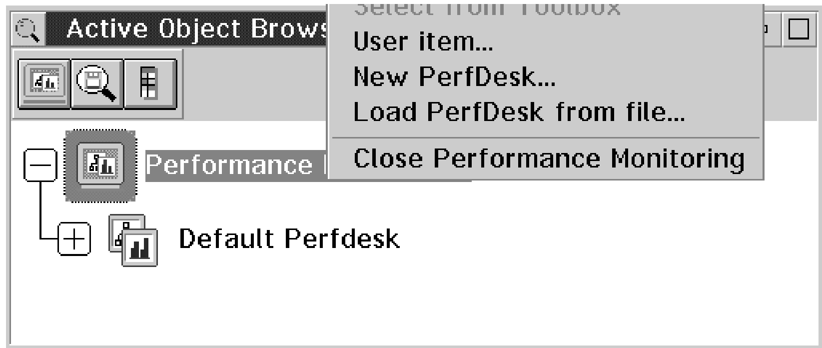

Figure 19-40. Create New PerfDesk

You will be prompted for a name of the new PerfDesk:

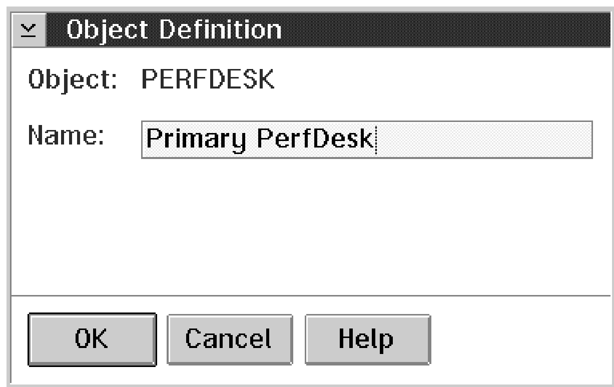

Figure 19-41. Define PerfDesk Name

Then you get the refreshed Active Object Browser window:

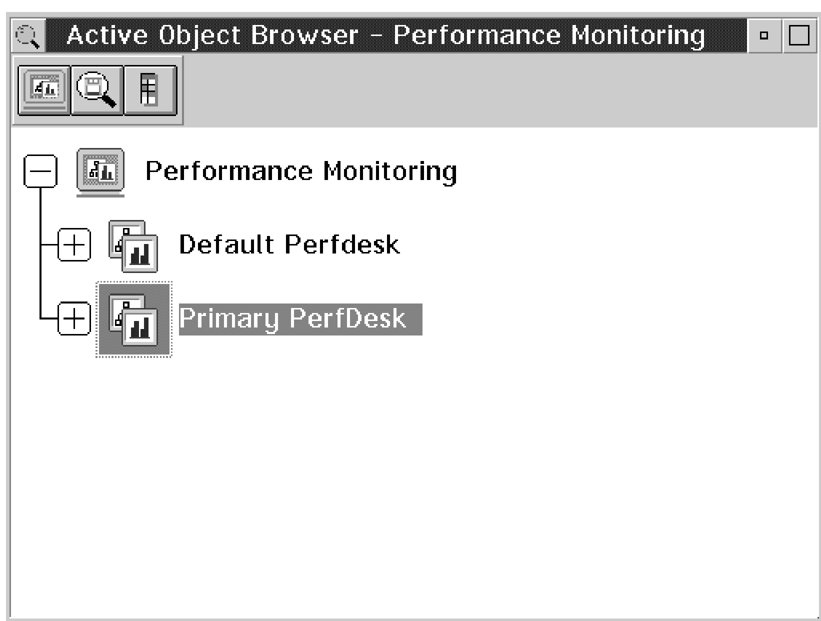

Figure 19-42. Active Object Browser with new PerfDesk

Now you can tailor the PerfDesk as required by loading DataViews that have been defined and saved earlier. If you stay with the one DataView 'System Activity', you get these active objects:

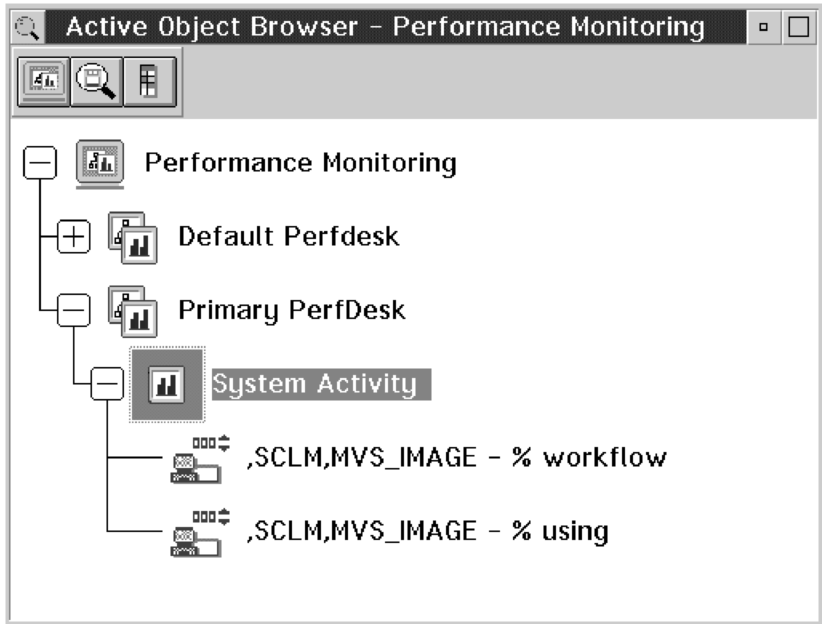

Figure 19-43. Customized PerfDesk

Assuming that this is exactly the PerfDesk that you want, now is the time to save it:

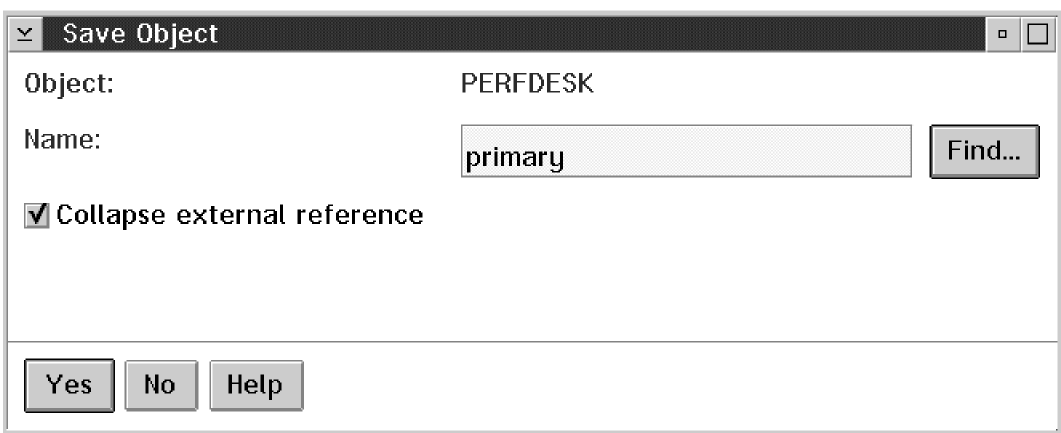

Figure 19-44. Save PerfDesk

The prompting message explains that `.PO` is the valid extension for PerfDesk objects.

Figure 19-45. Save PerfDesk

Of course you will accept this, and this concludes the creation and saving of our first PerfDesk.

Working with Saved PerfDesks

The second function that is available through **Browsers** in the action bar of the Performance Monitoring Main Window is **Saved PerfDesks**:

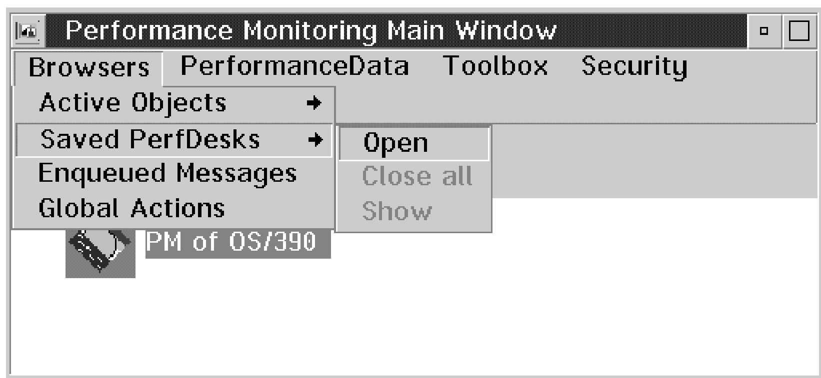

Figure 19-46. Browsers - Saved PerfDesks

Here, you get an overview of all PerfDesks, including the information about their file names. This is important to know, in case you want to erase a DataView or PerfDesk, because the only function available to do this is the OS/2 Delete command — and for that you need to know the file name.

Through the context menu of the selected PerfDesk you can load and (if you want) start all DataViews belonging to the PerfDesk. The Load & Start function ensures that all DataViews are started together, if there are more than one.

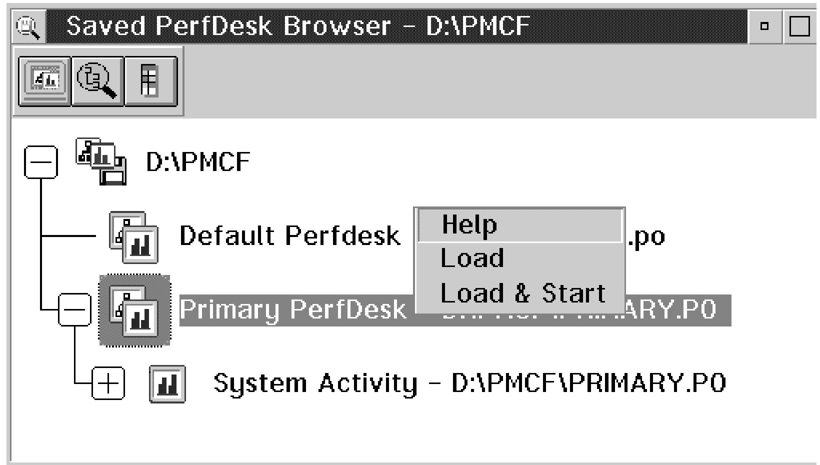

Figure 19-47. Load and Start PerfDesk

If you try to start a DataView to display performance data for a domain before you have logged on to that domain, you will be prompted for the logon parameters as described in First Round - Creating a DataView. This means that you do not have to select the domain you want to work with explicitly through the PM of OS/390 Domain Configuration. You can start implicitly by selecting one of your saved PerfDesks.

PerfDesks

- Create one or more PerfDesks and tailor them to your requirements
- You have immediately access to the reports that you need for your performance monitoring tasks

Some more Details about Working with PM of OS/390

DataView

Start - Freeze - Stop

After either loading or creating a DataView, you see the status information **Loaded**. In this status, you can modify the DataView, either by adding another series, or by changing the collection details. Both of these functions are available through the context menu.

In the lower left-hand corner, you see the three buttons **Stop**, **Pause**, and **Start** as you know them from your video recorder at home. You can start data collection, you can hold it (it will be shown as **Frozen**), and you can stop it.

Note: Modifications of the DataView are possible only before starting data collection. If you really want to make any changes, you have to stop and **reset** data collection.

DataView Collection Details

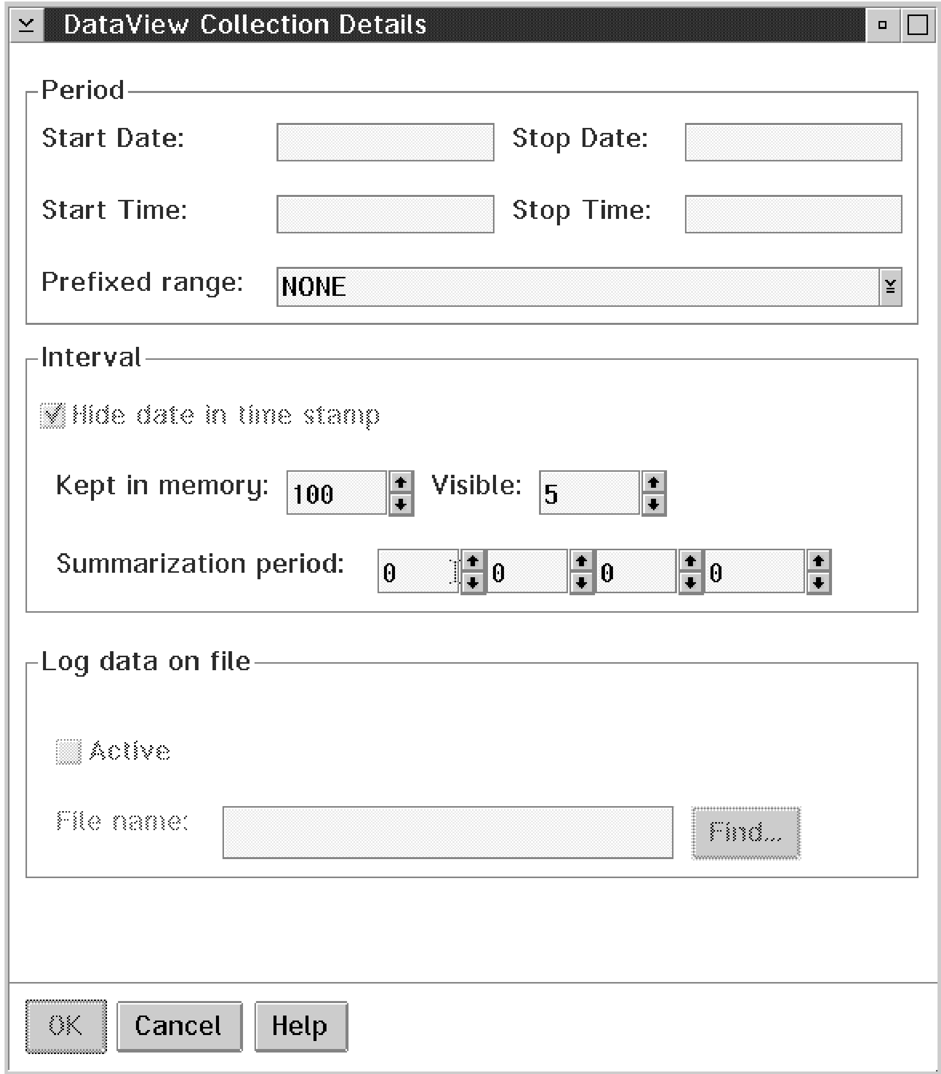

Figure 19-48. DataView Collection Details

Here, you can select the reporting period to access historical data that are available in the buffer or the VSAM data sets allocated to the Monitor III gatherer session.

Furthermore, you can define the amount of data to be stored (default value is 100 samples), the number of series to be displayed (default value is 5 samples), and a summarization period (in hours, minutes, seconds, and milliseconds) to get the data in a more compact form.

You can also specify data logging, which requires that you define a summarization period.

Series List

This function provides a list of all series belonging to the DataView, with their details. By clicking the right button on one of the series, you get the context menu that allows to modify or remove that series.

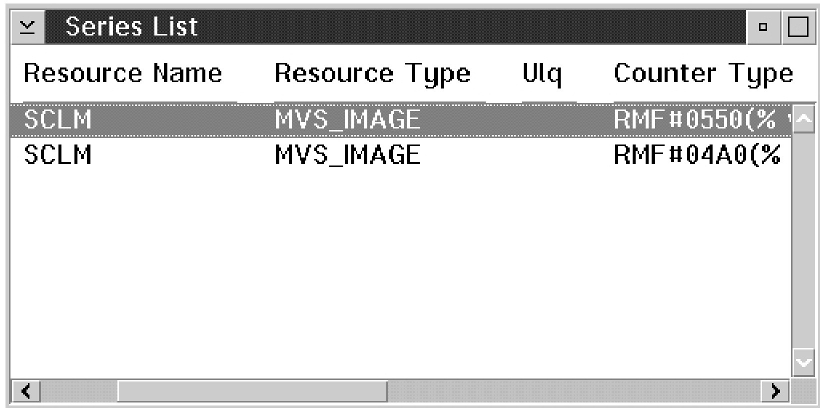

Figure 19-49. DataView - Series List

Series Definition

The panel to define a new series offers several entry fields. In the first round, we defined a series for the resource that we had selected in the domain configuration. But this panel offers more capabilities.

Series Name
Name of the series to be displayed in the DataView. This should be used for DataViews with one series only.

Application
If you have several PM applications installed, here you can select which one you want to get the data from.

Resource
Here you specify the resource that you want, by name and type. **Pick up** displays the selected resource from the CfgView (if one has been selected), **More...** gives some more details.

Counters
Depending on the resource you have chosen, you get a list of all available counters. If you add a series to a DataView, you will see only single- or multi-valued counters, because they cannot be mixed in one DataView

With **>>** (copy to **Selected**) and **<<** (remove from **Selected**) you can define the counter you want to add. If you click one selected counter, you might see the active **Details...** button that can lead you to the work scope dialog, or to the filter dialog, depending on the type of counter.

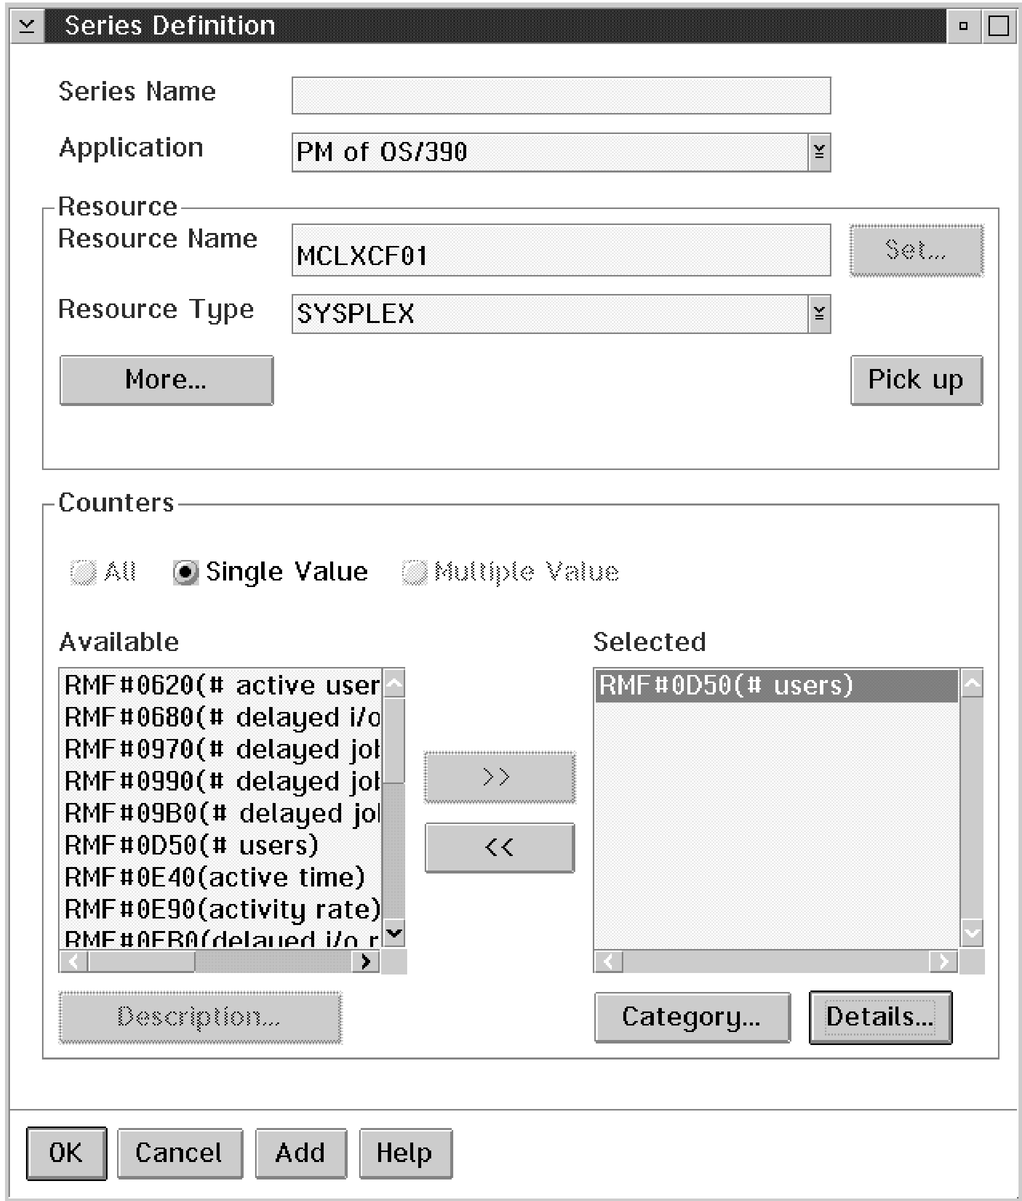

Figure 19-50. Series Definition

Work Scope Dialog

Note: This dialog is available only for systems running in goal mode.

In this example, we want to get the number of TSO users depending on the service class period.

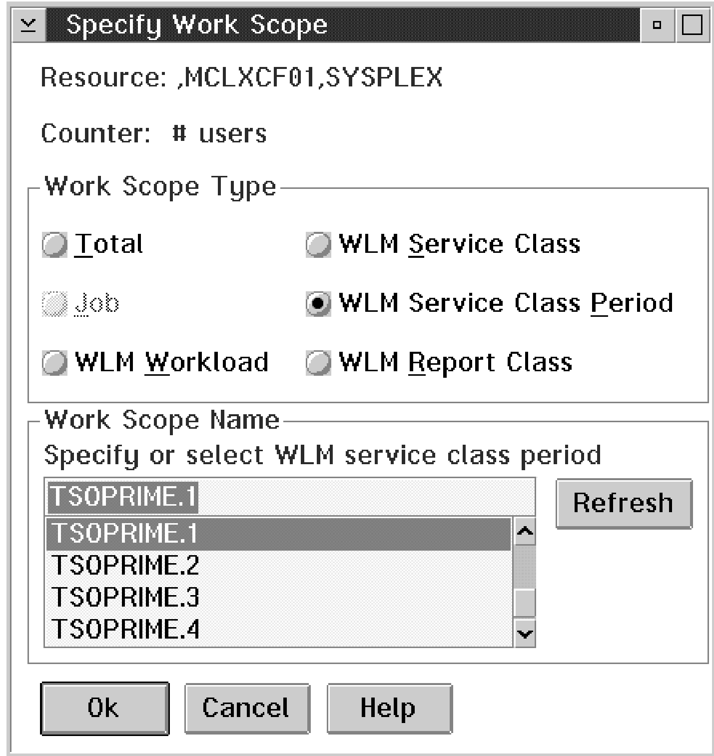

Figure 19-51. Work Scope Dialog

The resulting DataView is shown in Figure 19-52.

Resource and Counter Information

This area informs you about the resources and counters for which a work scope name can be specified in this panel.

Work Scope Type

You must select a work scope type. Only the radio buttons of the choices that are available for the combination of Resource and Counter are enabled.

Work Scope Name

For `Work Scope Type = Total` nothing has to be entered here.

For other work scope types, the Work Scope Name can be one of the following:

- Jobname,
- WLM workload name
- WLM service class name
- WLM service class period name
- WLM report class name

A list of valid Work Scope Names 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 has been filled, the entry field may be filled in by selecting an item from the list.

PM of OS/390 will remember the selection list and show it again the next time. **Refresh** can be used to get a refreshed selection list.

Push Buttons Ok, Cancel, Help

Ok causes PM of OS/390 to show the counter value for the given resource and the specified work scope. In the case of a syntax error, a pop-up message is shown. Errors can be avoided by selecting an entry of the selection list, rather than typing the name in.

Cancel causes PM of OS/390 to ignore all information entered during this invocation of the dialog.

Help causes PM of OS/390 to show a Help panel.

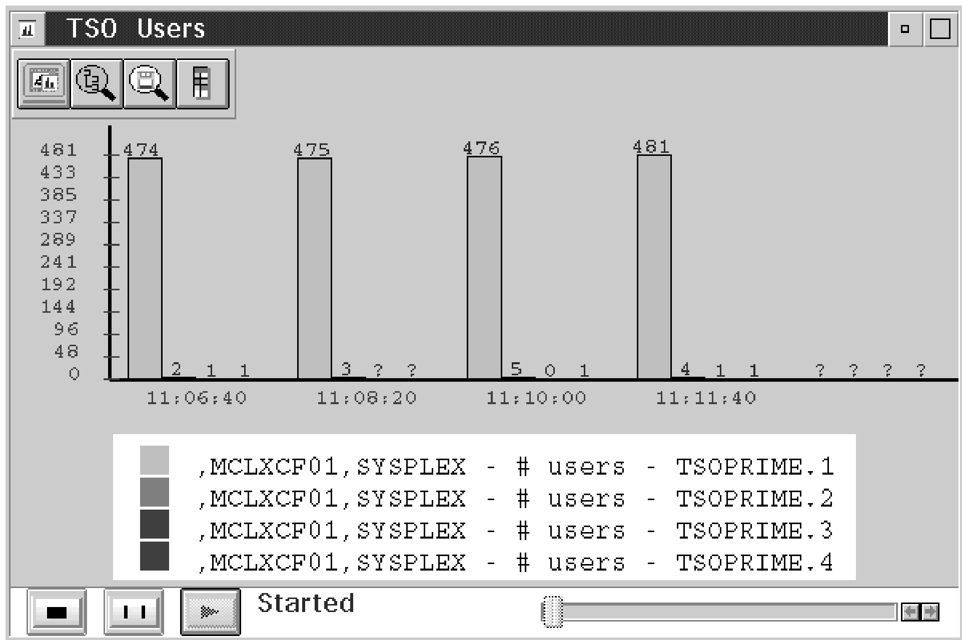

Figure 19-52. Performance Data for WLM Work Scopes

Filter Dialog

A multi-valued counter 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 counter, and to sort them by name or value.

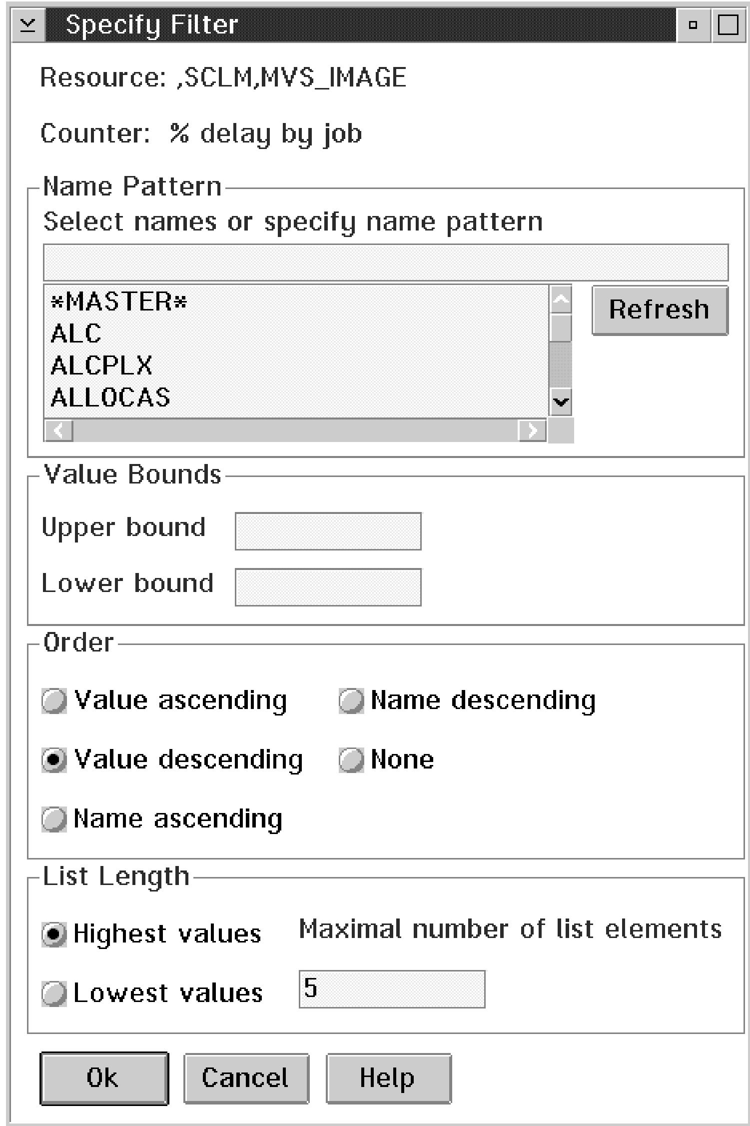

Figure 19-53. Filter Dialog

Filtering is performed stepwise in the following sequence:

1. Name-pattern matching
2. Value bound comparison
3. Ordering
4. List length reduction

Resource and Counter Information

This area informs you about the resources and counters for which a filter can be specified in this panel.

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 counter 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*), it has to be enclosed in back slashes

Example

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*
```

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.

PM of OS/390 will remember the selection list and will show it again the next time. **Refresh** can be used to get a refreshed selection list.

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 counter 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.

Order

You must select one of the choices:

- Value ascending
- Value descending
- Name ascending
- Name descending

- None

for ordering the values in the list of name-value pairs of the multi-valued counter.

List Length

You must select one of the choices:

- Highest values
- Lowest values

for restricting the length of the list of name-value pairs of the multi-valued counter. Additionally, the maximum number of list elements must be specified. It must be a positive whole number.

Push Buttons Ok, Cancel, Help

Ok causes PM of OS/390 to show the multi-valued counter for the given Resource and the specified Filter. In case of an input error, a pop-up message is shown. Filtering is performed in steps with the following sequence:

1. Name pattern matching
2. Value bound comparison
3. Ordering
4. List length reduction

Cancel causes PM of OS/390 to ignore all information entered via this invocation of the dialog.

Help causes PM of OS/390 to show a Help panel.

Batch Collection

PM of OS/390 offers the capability to store data that are being gathered into a log file. You can initiate this process by clicking on **Performance Data** in the action bar and by selecting **Batch Collection**.

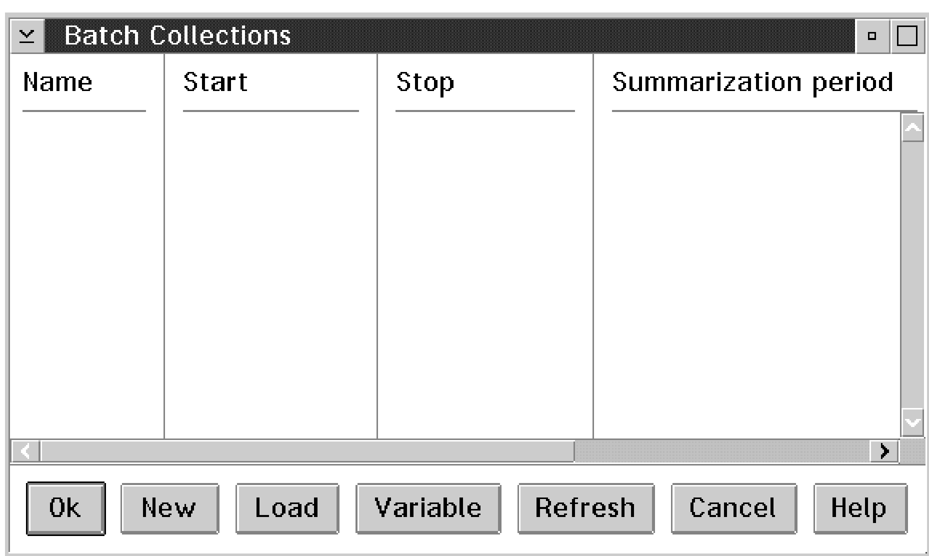

Figure 19-54. List of Batch Collections

You start with a panel showing the list of batch collections already defined, in the above example the list is empty. If you want to create a new collection, press **New** , and you get this panel:

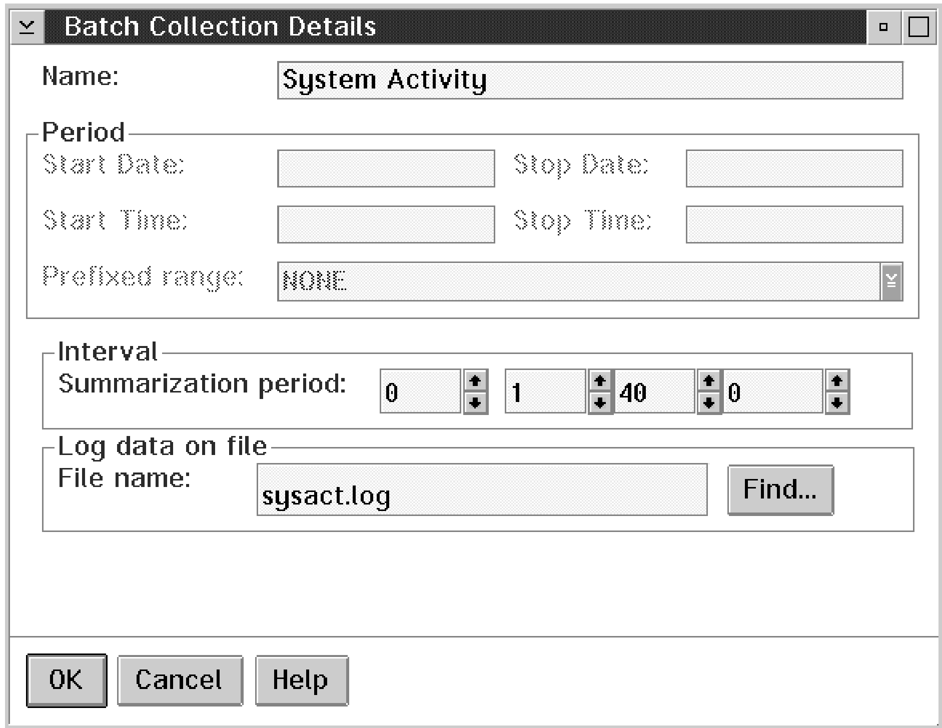

Figure 19-55. Define new Batch Collection

There are two fields that are important:

Summarization period
Specifies how often performance data is summarized. The interval is in the form hh:mm:ss:ms. A new record is added for each interval in the log file specified in the File name field.

Note: It is mandatory to define the summarization period, otherwise no data storing will be performed.

File name
The name of the file where the data should be stored.

The format of data created as batch collection is the following:

```
"PM performance data batch file"

"1996-12-16-15.07.06";",SCLM,MVS_IMAGE";"GpmAppl.RMF#0160";"% delay";3;"12/16/96";"15:05:20";"12/16/96";"15:07:00"
"1996-12-16-15.07.06";",SCLM,MVS_IMAGE";"GpmAppl.RMF#04A0";"% using";2;"12/16/96";"15:05:20";"12/16/96";"15:07:00"
"1996-12-16-15.07.06";",SCLM,MVS_IMAGE";"GpmAppl.RMF#0550";"% workflow";45;"12/16/96";"15:05:20";"12/16/96";"15:07:00"
"1996-12-16-15.08.46";",SCLM,MVS_IMAGE";"GpmAppl.RMF#0160";"% delay";1;"12/16/96";"15:07:00";"12/16/96";"15:08:40"
"1996-12-16-15.08.46";",SCLM,MVS_IMAGE";"GpmAppl.RMF#04A0";"% using";1;"12/16/96";"15:07:00";"12/16/96";"15:08:40"
"1996-12-16-15.08.46";",SCLM,MVS_IMAGE";"GpmAppl.RMF#0550";"% workflow";63;"12/16/96";"15:07:00";"12/16/96";"15:08:40"
"1996-12-16-15.10.28";",SCLM,MVS_IMAGE";"GpmAppl.RMF#0160";"% delay";0;"12/16/96";"15:08:40";"12/16/96";"15:10:20"
"1996-12-16-15.10.28";",SCLM,MVS_IMAGE";"GpmAppl.RMF#04A0";"% using";0;"12/16/96";"15:08:40";"12/16/96";"15:10:20"
"1996-12-16-15.10.28";",SCLM,MVS_IMAGE";"GpmAppl.RMF#0550";"% workflow";68;"12/16/96";"15:08:40";"12/16/96";"15:10:20"
```

Figure 19-56. Log File containing Data for 3 Counters

The Contents of a Toolbox

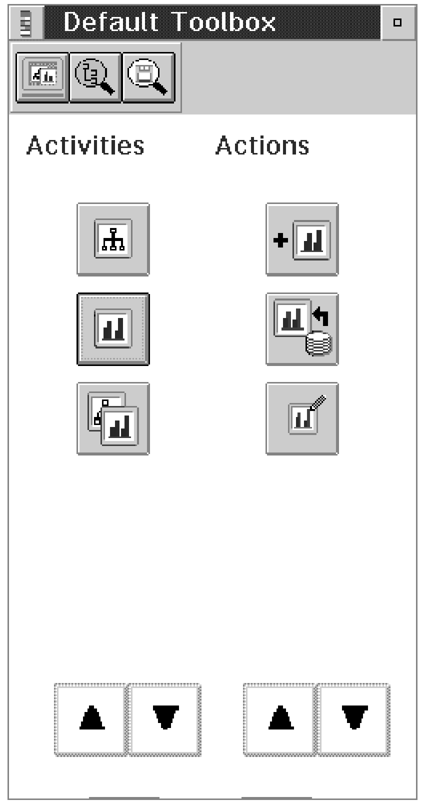

Figure 19-57. Example of a Toolbox

A Toolbox provides a graphical interface for managing, configuring, and defining tasks. You can create a Toolbox for each task or set of tasks that you want to perform. Each Toolbox contains two lists: a list of activities related to the task, and a list of actions related to each activity. You can define both the activities, which appear on the left-hand side of the Toolbox, and the valid actions for each activity, which appear on the right-hand side.

For example, assume you want a Toolbox to create and change DataViews. You create a Toolbox that has the following activities on the left-hand side:

- Create a DataView
- Display data
- Control a DataView

For the activity **Create a DataView** you could have these actions on the right-hand side of the Toolbox:

- Add a new DataView to a PerfDesk
- Add series to a DataView
- Save DataView as external

The action **Add a new DataView to a PerfDesk** could also be an action in another Toolbox that you create to build PerfDesks.

Each Toolbox can be stored in a file.

Performance Analysis

PM of OS/390 offers the capability to analyze performance data being displayed in a DataView.

If you have defined a DataView to show delayed jobs for an MVS image, you might get the following information:

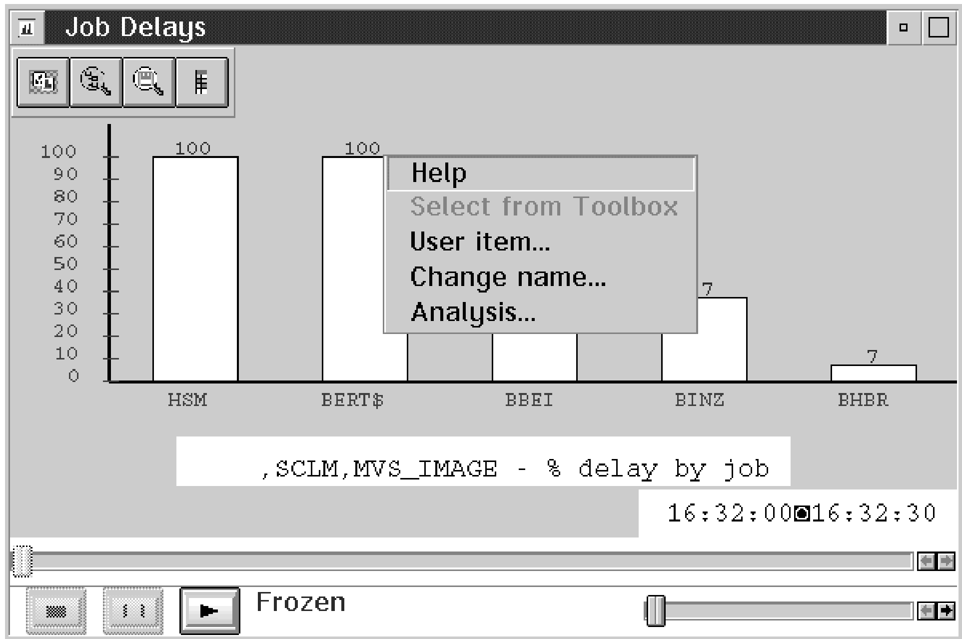

Figure 19-58. DataView - Delay by Job

In this example, you find two address spaces (HSM and BERT$) being delayed for 100%. The **analysis** function of PM of OS/390 is available through the context menu for each series on a DataView, you get it by right-clicking one of the series. If you want to start analysis for job BERT$, you select **Analysis...** , and you get the analysis dialog. The title shows the domain for which analysis is intended.

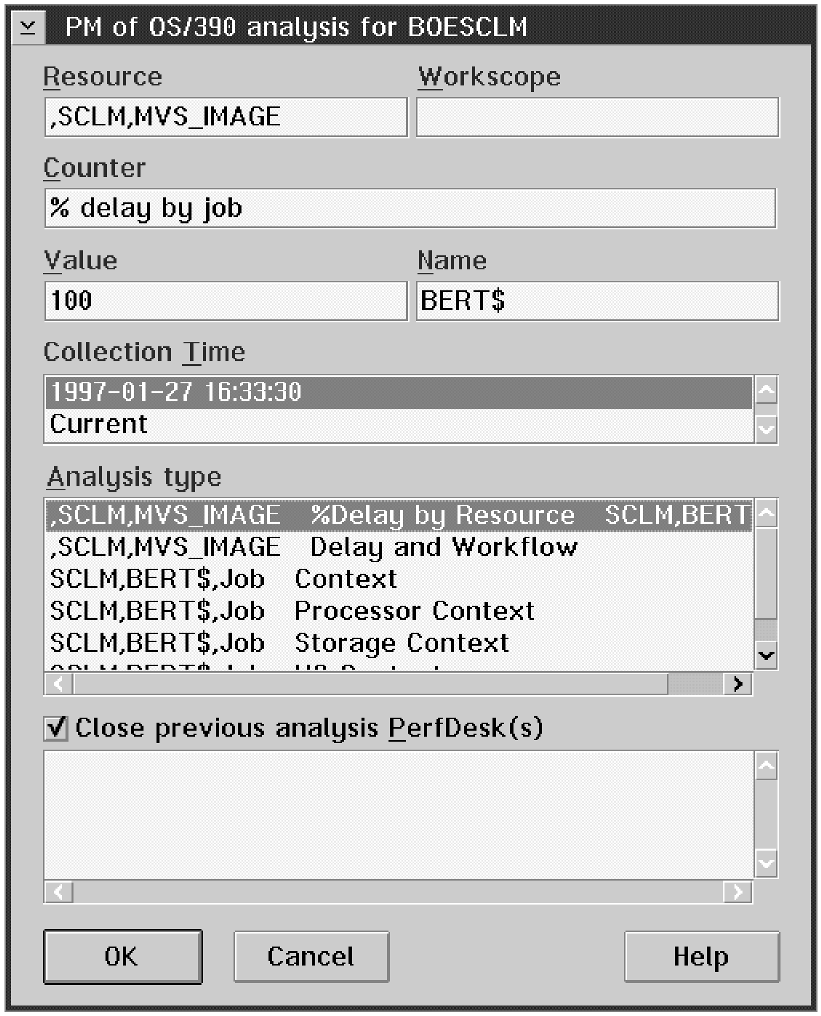

Figure 19-59. PM of OS/390 - Analysis Dialog

The fields

Resource, Work Scope, Counter, Value, Name, and Collection Time

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
- the **current** data which analyses the next counter sample.

If the work scope field is empty, the counter is global rather than for a specific work scope.

Analysis Type

This listbox allows you to select the alternative of the next analysis step. Each alternative is shown with the resource and its work scope (if required) as the target for the analysis.

Close Previous Analysis PerfDesk(s)

The PerfDesks that were previously shown during the analysis process can be closed to avoid a "messy" screen with too many DataViews. The selected PerfDesks will be closed, when the check box is selected.

In this example, you see that several types of analysis are offered, **%Delay by Resource** is currently the default. If you feel that this one is the best offering, press **OK**, and you get the corresponding DataView.

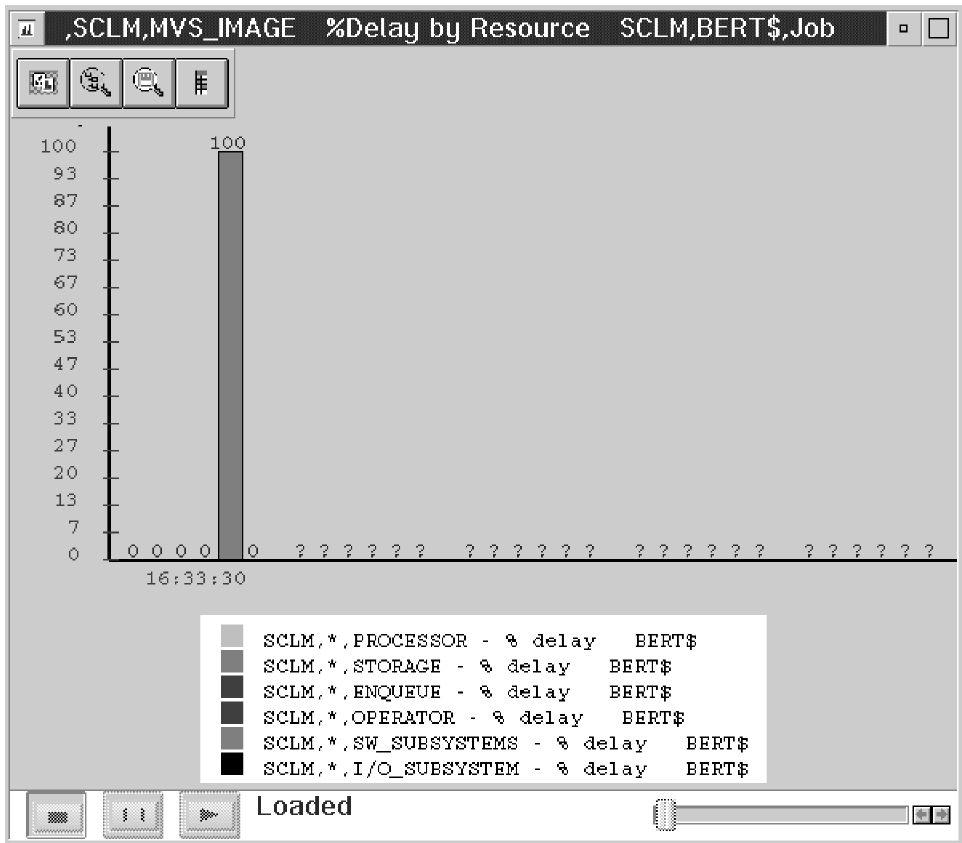

Figure 19-60. DataView - Delay by Resource

This DataView will show all delay reasons for one job, in this example, the delay for a S/W subsystem is the only reason - 100%. If you are familiar with Monitor III, you know that there are three types of S/W subsystems that can cause the delay: JES, XCF, and HSM - the latter is obvious the right one in this example. The current implementation of the analysis function cannot show this explicitly.

As seen in Figure 19-58 on page 19-64, HSM is also delayed by 100%, therefore the analysis for HSM could be the next step. The analysis dialog looks like this:

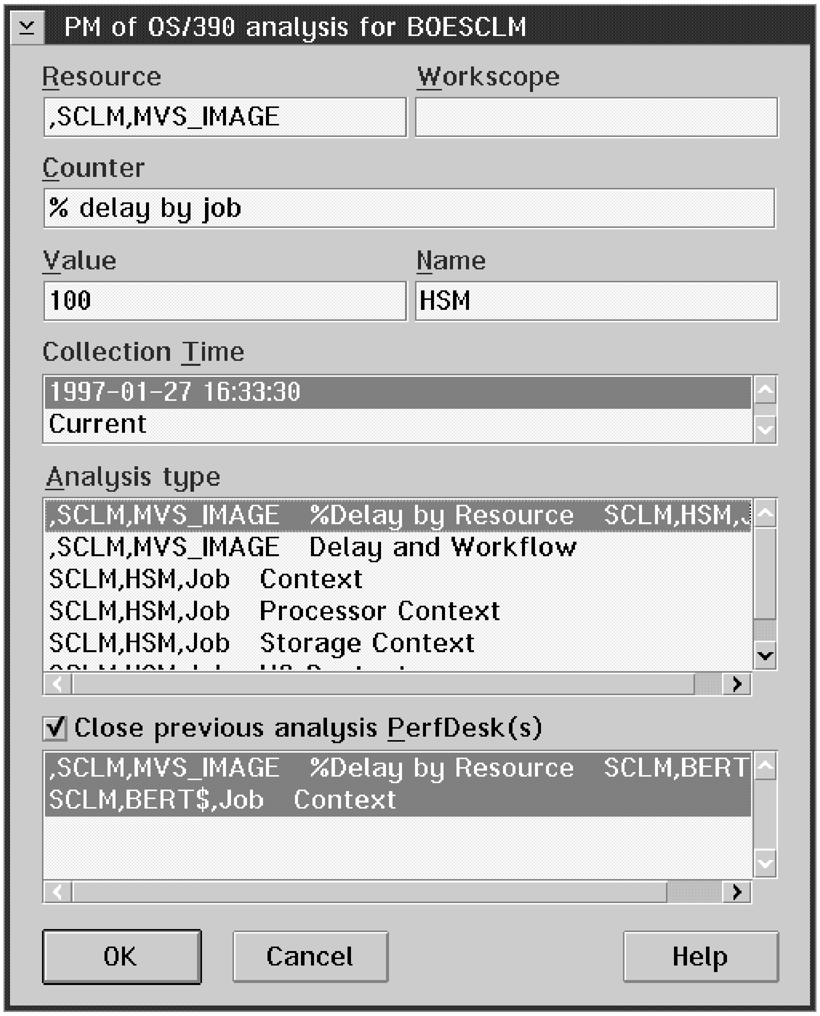

Figure 19-61. PM of OS/390 - Analysis Dialog

Continuing with the default selection leads to a DataView that shows all delay reasons for HSM.

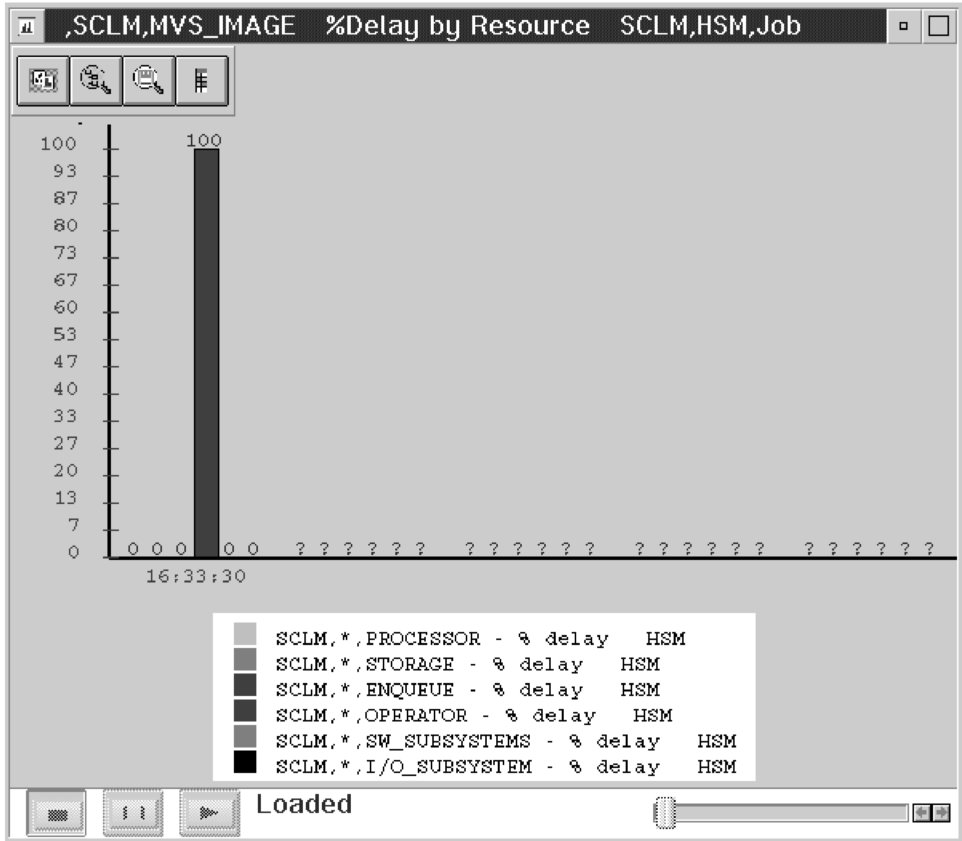

Figure 19-62. DataView - Delay by Resource

HSM is delayed by an operator request - probably an outstanding mount request for a migration volume.

Saved PerfDesks

The analysis support of PM of OS/390 requires several PerfDesks consisting of one or more DataViews. They are available on your workstation and you will see them if you click on icon **Saved PerfDesks Browser** in the Performance Monitoring Main Window:

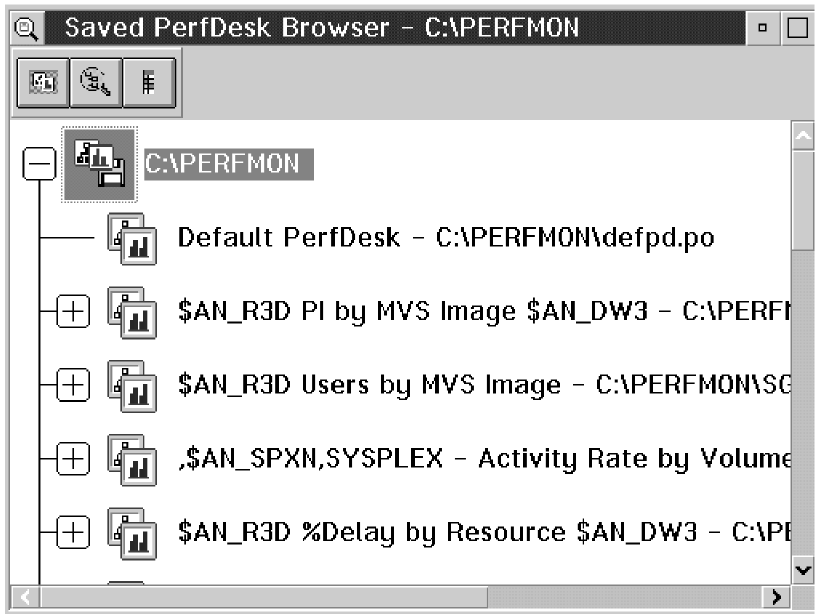

Figure 19-63. DataView - Delay by Resource

It is **not recommended** to load and start these PerfDesks directly because they need some initialization which will be performed during the analysis dialog. From this point, you should load and start only those PerfDesks which you have defined and stored by yourself.

Available Performance Data

In this chapter, we have talked a lot about *series*, which are *resource-counter pairs*. Here, *resource* is the name of a specific instance of a resource type. Not every counter is available for every resource type, and some resource types have a lot of counters associated with them, so here is an overview. In the lists that follow, the headers designate the resource types, the list items the counters that are associated with them:

AUX - Auxiliary Storage

slots by job
slots

CENTRAL - Central Storage

% frames active
% frames available
% frames CSA
% frames idle
% frames LPA
% frames NUC

% frames SQA
# frames online
unreferenced interval count
working set by job
working set

CHANNEL - Channel Subsystem

% channel path partition utilization
% channel path total utilization

CSC - Common Storage Area

% available
% not released
% utilization
# frames defined
# frames not released
% utilization by job
# frames not released by job
# frames used by job
% utilization
# frames not released
# frames used

ECSC - Extended Common Storage Area

% available
% not released
% utilization
# frames defined
# frames not released
% utilization by job
# frames not released by job
# frames used by job
% utilization
# frames not released
# frames used

ENQUEUE - Enqueue Delays

% delay
# delayed jobs
% delay by job
# delayed jobs by WLM service class period
# delayed jobs by WLM report class
# delayed jobs by WLM service class
# delayed jobs by WLM workload
% delay
# delayed jobs

ESQA - Extended System Queue Area

% available
% not released
% utilization
# frames defined
# frames not released
% utilization by job
# frames not released by job
# frames used by job
% utilization
# frames not released
# frames used

EXPANDED - Expanded Storage

% frames active
% frames available
% frames CSA
% frames idle
% frames LPA
% frames SQA
# frames online
migration age
working set by job
working set

HSM - HSM Subsystem

% delay
# delayed jobs
% delay by job
# delayed jobs by WLM service class period
# delayed jobs by WLM report class
# delayed jobs by WLM service class
# delayed jobs by WLM workload
% delay
# delayed jobs

IOSUB - I/O Subsystem

% delay
# delayed i/o requests
# delayed jobs
# using jobs
i/o activity rate
delayed i/o request rate
% channel path partition utilization by channel path
% channel path total utilization by channel path
% connect time by job
% delay by job
% using by job
% all channel paths busy by LCU
# delayed i/o requests by LCU
delayed i/o request rate by LCU
delayed jobs by WLM service class period

using jobs by WLM service class period
delayed jobs by WLM report class
using jobs by WLM report class
delayed jobs by WLM service class
using jobs by WLM service class
% CU busy by channel path and CU
% director port busy by channel path and CU
% active time by volume
% connect time by volume
% delay CU busy by volume
% delay device busy by volume
% delay director port busy by volume
% disconnect time by volume
% pending time by volume
i/o activity rate by volume
i/o intensity by volume
IOS queue time by volume
response time by volume
delayed jobs by WLM workload
using jobs by WLM workload
% connect time
% delay
% using
% delay by volume
# delayed jobs
# using jobs

JES - JES Subsystem

% delay
# delayed jobs
% delay by job
# delayed jobs by WLM service class period
# delayed jobs by WLM report class
# delayed jobs by WLM service class
# delayed jobs by WLM workload
% delay
# delayed jobs

LCU - Logical Control Units

% all channel paths busy
# delayed i/o requests
i/o activity rate
delayed i/o request rate
% CHPID taken by channel path and CU
% CU busy by channel path and CU
% director port busy by channel path and CU
% active time by volume
% connect time by volume
% delay CU busy by volume
% delay device busy by volume
% delay director port busy by volume
% disconnect time by volume
% pending time by volume

i/o activity rate by volume
IOS queue time by volume
response time by volume

MVSIMAGE - MVS System

% delay
% idle
% unknown
% using
% workflow
# active users
# users
execution velocity
transaction ended rate
% delay by job
% idle by job
% unknown by job
% using by job
% workflow by job
% workflow by WLM service class period
# active users by WLM service class period
# users by WLM service class period
execution velocity by WLM service class period
response time by WLM service class period
transaction ended rate by WLM service class period
% workflow by WLM report class
# active users by WLM report class
# users by WLM report class
execution velocity by WLM report class
response time by WLM report class
transaction ended rate by WLM report class
% workflow by WLM service class
# active users by WLM service class
# users by WLM service class
execution velocity by WLM service class
response time by WLM service class
transaction ended rate by WLM service class
% workflow by WLM workload
# active users by WLM workload
# users by WLM workload
execution velocity by WLM workload
response time by WLM workload
transaction ended rate by WLM workload
% delay
% idle
% unknown
% using
% workflow
% workflow
# active users
# users
execution velocity

response time
transaction ended rate

OPERATOR - Operator Delays

% delay
# delayed jobs
% delay by job
# delayed jobs by WLM service class period
# delayed jobs by WLM report class
# delayed jobs by WLM service class
# delayed jobs by WLM workload
% delay
# delayed jobs

PROC - Processor Delays

% delay
% partition utilization
% SRB
% TCB
% TCB + SRB
% total utilization
# delayed jobs
# processors online
# using jobs
% delay by job
% TCB + SRB by job
% using by job
# delayed jobs by WLM service class period
# using jobs by WLM service class period
# delayed jobs by WLM report class
# using jobs by WLM report class
# delayed jobs by WLM service class
# using jobs by WLM service class
# delayed jobs by WLM workload
# using jobs by WLM workload
% delay
% TCB + SRB
% using
# delayed jobs
# using jobs
% available
% not released
% utilization
# frames defined
# frames not released
% utilization by job
# frames not released by job
# frames used by job
% utilization
# frames not released
# frames used

STORAGE - Storage Delays

% delay
# delayed jobs
% delay by job
% delay for COMM by job
% delay for LOCL by job
% delay for OTHR by job
% delay for OUTR by job
% delay for SWAP by job
# frames active by job
# frames DIV by job
# frames fixed by job
# frames idle by job
# frames total by job
es rate per residency time by job
pgin rate per residency time by job
working set by job
delayed jobs by WLM service class period
delayed jobs for COMM by WLM service class period
delayed jobs for LOCL by WLM service class period
delayed jobs for OTHR by WLM service class period
delayed jobs for OUTR by WLM service class period
delayed jobs for SWAP by WLM service class period
frames active by WLM service class period
frames fixed by WLM service class period
frames idle by WLM service class period
pgin rate by WLM service class period
delayed jobs by WLM report class
delayed jobs for COMM by WLM report class
delayed jobs for LOCL by WLM report class
delayed jobs for OTHR by WLM report class
delayed jobs for OUTR by WLM report class
delayed jobs for SWAP by WLM report class
frames active by WLM report class
frames fixed by WLM report class
frames idle by WLM report class
pgin rate by WLM report class
delayed jobs by WLM service class
delayed jobs for COMM by WLM service class
delayed jobs for LOCL by WLM service class
delayed jobs for OTHR by WLM service class
delayed jobs for OUTR by WLM service class
delayed jobs for SWAP by WLM service class
frames active by WLM service class
frames fixed by WLM service class
frames idle by WLM service class
pgin rate by WLM service class
delayed jobs by WLM workload
delayed jobs for COMM by WLM workload
delayed jobs for LOCL by WLM workload
delayed jobs for OTHR by WLM workload
delayed jobs for OUTR by WLM workload
delayed jobs for SWAP by WLM workload

frames active by WLM workload
frames fixed by WLM workload
frames idle by WLM workload
pgin rate by WLM workload
% delay
% delay for COMM
% delay for LOCL
% delay for OTHR
% delay for OUTR
% delay for SWAP
# frames active
# frames DIV
# frames fixed
# frames idle
# frames total
es rate per residency time
pgin rate per residency time
working set
delayed jobs
delayed jobs for COMM
delayed jobs for LOCL
delayed jobs for OTHR
delayed jobs for OUTR
delayed jobs for SWAP
frames active
frames fixed
frames idle
pgin rate

SWSUB - Subsystem Delays

% delay
# delayed jobs
% delay by job
# delayed jobs by WLM service class period
# delayed jobs by WLM report class
# delayed jobs by WLM service class
# delayed jobs by WLM workload
% delay
# delayed jobs

SYSPLEX - Sysplex

active users
delayed i/o requests
delayed jobs for i/o
delayed jobs for processor
delayed jobs for storage
users
i/o activity rate
delayed i/o request rate
service units (capacity) / transaction
transaction ended rate
% channel path partition utilization by channel path
% channel path total utilization by channel path

% all channel paths busy by LCU
# delayed i/o requests by LCU
delayed i/o request rate by LCU
% delay by mvs image
% idle by mvs image
% processor utilization by mvs image
% unknown by mvs image
% using by mvs image
% workflow by mvs image
# active users by mvs image
# delayed jobs for i/o by mvs image
# delayed jobs for processor by mvs image
# delayed jobs for storage by mvs image
# users by mvs image
execution velocity by mvs image
transaction ended rate by mvs image
active time by WLM service class period
execution velocity by WLM service class period
execution velocity goal by WLM service class period
percentile achieving response time goal by WLM service class period
performance index by important WLM service class period
performance index by WLM service class period
queue time by WLM service class period
response time by WLM service class period
response time goal by WLM service class period
response time goal percentile by WLM service class period
transaction ended rate by WLM service class period
active time by WLM report class
execution velocity by WLM report class
queue time by WLM report class
response time by WLM report class
transaction ended rate by WLM report class
active time by WLM service class
execution velocity by WLM service class
queue time by WLM service class
response time by WLM service class
transaction ended rate by WLM service class
% CU busy by channel path and CU
% director port busy by channel path and CU
% active time by volume
% connect time by volume
% delay CU busy by volume
% delay device busy by volume
% delay director port busy by volume
% disconnect time by volume
% pending time by volume
i/o activity rate by volume
i/o intensity by volume
IOS queue time by volume
response time by volume
active time by WLM workload
execution velocity by WLM workload
queue time by WLM workload
response time by WLM workload

transaction ended rate by WLM workload
execution velocity goal
percentile achieving response time goal
performance index
response time goal
response time goal percentile
active users
users
active time
execution velocity
queue time
response time
transaction ended rate

VOLUME - DASD Information

% active time
% connect time
% delay CU busy
% delay device busy
% delay director port busy
% disconnect time
% pending time
i/o activity rate
i/o intensity
IOS queue time
response time
% delay by job
% using by job
% delay
% using

XCF - XCF Subsystem

% delay
# delayed jobs
% delay by job
# delayed jobs by WLM service class period
# delayed jobs by WLM report class
# delayed jobs by WLM service class
# delayed jobs by WLM workload
% delay
# delayed jobs

Chapter 20. RMF Client/Server Enabling (RMFCS)

About RMF on a Client

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 chapter covers the following topics:

- Overview of RMFCS
- RMFCS scenarios
- Installation and setup 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 OS/390 systems without an active TSO/TCAS subsystem on the host.

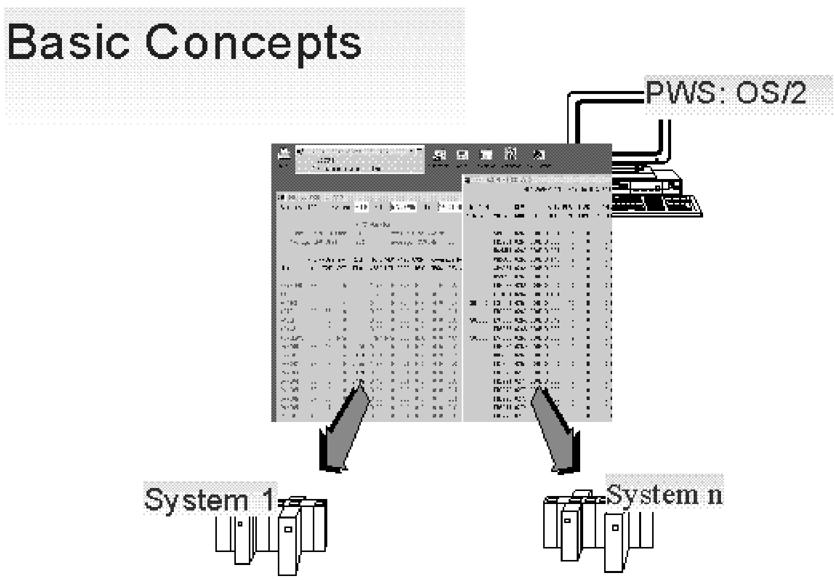

Figure 20-1. RMFCS Performance Data View - Example

The example shows an RMF-PWS client with

- A Monitor III SYSINFO report session connected to *MVS System_1*
- A Monitor II DEVICE report session connected to *MVS System_n*

running concurrently.

With RMFCS, you can establish as many sessions as you want with any MVS 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 OS/2 windows by using the ISPF/SPLIT function, which allows 32 logical screens. Each SPLIT creates a new OS/2 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 OS/390 performance management with a state-of-the-art user front end.

Hitherto, one or more 3270 TSO sessions were used for online monitoring of MVS performance data. The new concept of RMFCS uses an OS/2-based workstation as the single point of control for multiple MVS 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 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.

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

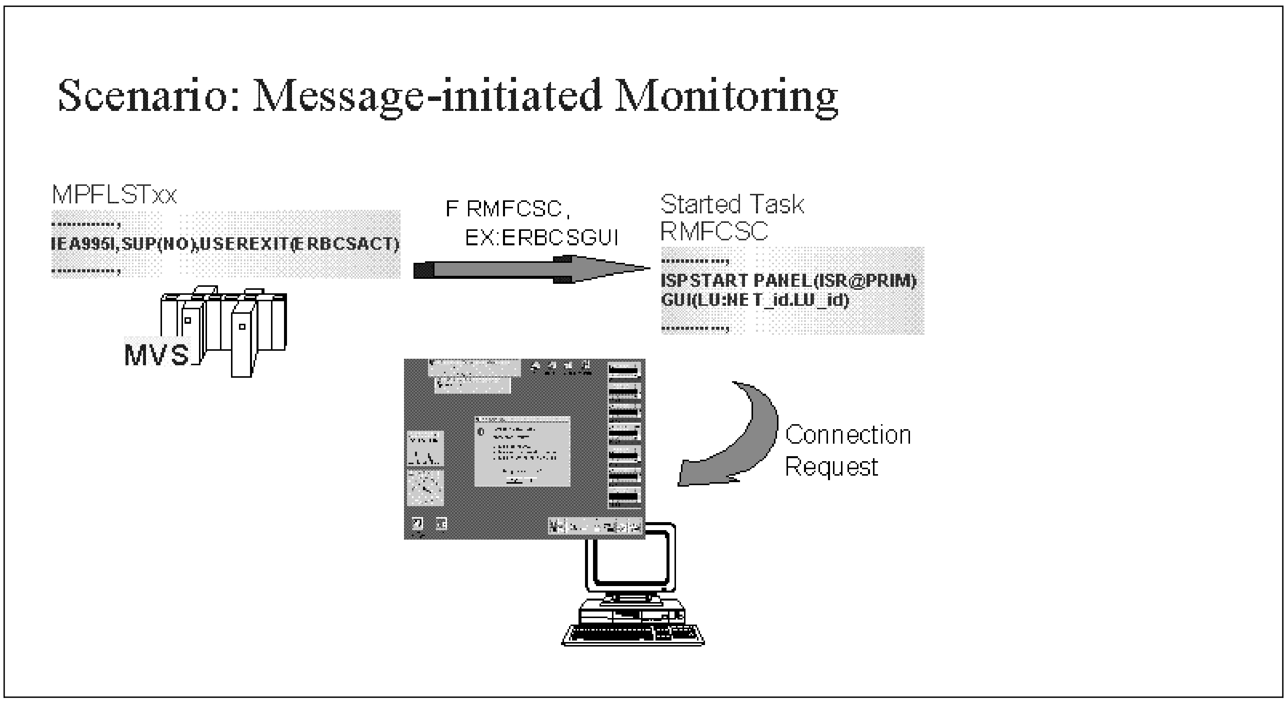

Figure 20-2. RMFCS Scenario I: Message-Initiated Monitoring

In this scenario, the MPF parmlib member is used for event handling and further processing of system alerts. It is assumed here 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 OS/2 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 OS/2 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 MVS hosts are supported.

Scenario II: Exception-Initiated Monitoring

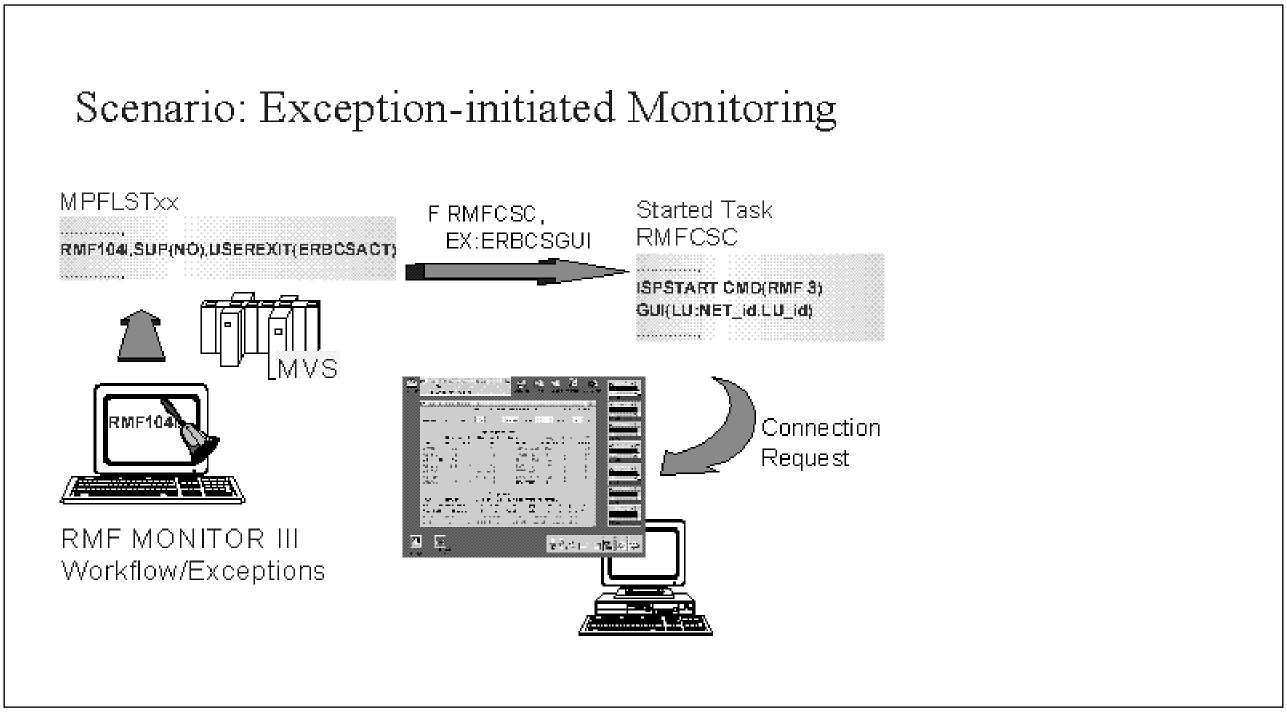

Figure 20-3. RMFCS Scenario II: Exception-Initiated Monitoring

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 MVS 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 MVS 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 fulfils 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 TCP/IP (Version 2.2) or APPC (ACF/VTAM Version 3.4) network connection from the workstation to the host. If APPC is used, the connection must be capable of supporting parallel LU 6.2 sessions.
  - ISPF 4.2.0
- Workstation Software
  - OS/2 2.x or higher
  - Communication Manager 1.x or TCP/IP 2.x
- Workstation Hardware
  - A 386-processor-based workstation with 8MB of memory. For optimum performance, we recommend 16MB of memory and a 486 processor.
  - Display: XGA graphic card is recommended (or compatible graphics with 1024 * 768 resolution)

For details related to ISPF, please refer to *OS/390 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.LINKLIB:

  ERBCSCTL
  ERBCSACT
  ERBCSWTO

- RMFCS REXX procedures to the RMF CLIST library SYS1.SERBCLS:

  ERBCSGUI
  ERBCSINI
  ERBM3B
  ERBM3BWX
  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 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
   - Install ISPFCS code on your workstation (see description under ISPF 3.7 on your MVS 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 will be 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 ERBM3BWX. This procedure performs similar steps to those in the following example. It assumes that:

   - You are working with TSO userid TSO1
   - You have selected qualifier BAT1 for your RMFCS data sets

   a. Rename your current Monitor III table data set:

   ```
   ren RMFOS270.isptable rmftmp.isptable
   ```

   b. Start an RMF session. This results in the creation of a new table data set:

   ```
   rmf
   ERB0TABL dataset 'TSO1.RMFOS270.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 2.7.0  Workflow/Exceptions          Line 1 of 19
Command ===> _                                                Scroll ===> HALF

Samples: 100     System: L96S  Date: 04/04/99  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         *DMN001        4      1          7
ALL BATCH     30     12          2         *DMN002       55      0        100
ALL ASCH       1      0    No work         *DMN003        0      0         88
ALL OMVS                  Not avail        *DMN004        0      0         92
*PROC         96      2         99         *DMN005        2      1         93

------------------------------ Exceptions -------------------------------------
Name       Reason        Critical val. Possible cause or action
*DMN001    OPER-Message    1.0 users   Awaiting reply to operator request 15.
ALL BATCH  OPER-Message   12.1 users   Awaiting reply to operator request 37.
*DMN006    OPER-Message    1.8 users   Awaiting reply to operator request 90.
*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
  __     DMN      1          WF                        2       3
  __     DMN      2          WF                        2       4
  __     DMN      3          WF                        2       5
  __     DMN      4          WF                        2       6
  __     DMN      5          WF                        2       7
  __     DMN      6          WF                        Not Displayed
  __     DMN      7          WF                        Not Displayed
  __     SYSTEM              EX-ANY
  __     JOB                 EX-UNAVAIL
  __     STOR                EX-AVG
  __     STOR                EX-AVG
  __     STOR                EX-AVG     *STOR
  __     SRVCLS   TSO        EX-AVG
```

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 RMFOS270.isptable 'bat1.rmfm3b.isptable'
ren rmftmp.isptable RMFOS270.isptable
```

This ensures that your Monitor III batch session can run with the definitions in data set BAT1.RMFM3B.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 20-18). By default, it issues the WTO message:

```
RMF101I
```

for MPF processing.

Startup

RMFCS Control Session

Either start procedure RMFCSC 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 RMFCSC,HLQ=USER#1
S RMFCSC,HLQ=USER#2
S RMFCSC,HLQ=USER#3
```

Depending on whether message-initiated or exception-initiated monitoring is in effect, the MODIFY command:

```
F RMFCSC,EX:ERBCSGUI
```

will be 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 doe 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 OS/390 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 RMFCSC 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(ERBCSWTO) 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

RMFCD Enabling consists of a number of:

- JCL procedures
- Modules
- REXX procedures
- CLISTs

These are described in this section.

JCL Procedures

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

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

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 OS/2 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 20-11.)

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

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 O
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 = "GUISCRW(121)"
guiscrd = "GUISCRD(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
/*-------------------------------------------------------------------*/
/* 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       "
reasn = "Reason        "
delay = "Critical val."
process = "Processing WFEX Report..."
wtomsg1 = msgid process
/*--------------------------------------------------------------------*/
/* Header lines for samples 2+3                                       */
/*--------------------------------------------------------------------*/
wtomsg2 = msgid name reasn 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: WFEX 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 WFEX exception lines exist        */
/*       2. if yes, issues WTO message RMF101I and sets             */
/*          the hardcopy request to 1                               */
/*                                                                  */
/********************************************************************/
Wfex_Handler_1: Procedure

tabnam = "ERBWFXT3"
msgid = "RMF101I 3B:"
msgtext = "WFEX 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
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 = "ERBWFXT3"
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
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 = "ERBWFXT3"

proc = "*PROC"
cpus = "CPUS%"                      /* Set the exception type       */
cpus_limit = "90"                   /* Set the threshold            */
cpus_msgid = "RMF103I 3B:"
```

```
ecsa = "*ECSA*"
secs = "SECS%"                       /* Set the exception type       */
secs_limit = "60"                    /* Set the threshold            */
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)                    /* Loops through the table      */
  If (SUBSTR(wfxname,1,5) = proc) &,
     (SUBSTR(wfxreasn,1,5) = cpus) &,
     (SUBSTR(wfxdelay,2,2) >= cpus_limit) Then
  Do                                 /* CPUS%  Threshold exceeded ?  */
    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                                 /* SECS% Threshold exceeded ?   */
    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 (systsvvc) SHARED"           /* Obtain actual value           */

If SUBSTR(systsvvc,1,3) > tcbsrb_limit Then
Do                                 /* Threshold exceeded ?          */

  msgtext = tcbsrb_text systsvvc]]"%" limit tcbsrb_limit]]"%"
  wtomsg = tcbsrb_msgid msgtext
  "SELECT PGM(ERBCSWTO) 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
```

RMF 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.

A

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.

B

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.

C

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

compatibility mode. The implicit state of an MVS system when no workload manager service policies are in effect. Contrast with *goal mode*.

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*.

D

data sample. See *sample*

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

DMN. Domain

domain. In compatibility mode, an optional method for setting bounds for the amount of service to be granted to a particular service class.

DP. Dispatching priority

E

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). (1) On an IBM 3090 processor complex, an extension of processor storage. (2) Optional high-speed storage that transfers 4KB pages to and from central storage.

G

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*.

goal mode. The implicit mode of an MVS system that has active service policies and performance goals defined by the workload manager. Contrast with *compatibility 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

H

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

I

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.

installation performance specification (IPS). In MVS, a set of installation-supplied control information used by the system workload manager. An IPS includes performance group definitions, performance objectives, and coefficients used to establish the service rate. See also service rate.

interactive session. In RMF, a monitor display-session that is controlled from the display terminal. Contrast with *background session*.

J

JES. Job Entry Subsystem

L

LCU. Logical control unit

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)

M

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

O

OMVS. Reference to OS/390 UNIX System Services

P

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. (1) The activity which monitors and allocates data processing resources to applications according to goals defined in a service level agreement or other objectives. (2) The discipline that encompasses collection of performance data and tuning of resources.

performance group. Group of work with the same performance objectives managed by the SRM.

PG. Performance group

PGN. Performance group number

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*.

R

range. The time interval you choose for your report.

Resident time. The time the address space was swapped in, in units of seconds.

S

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.

T

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.

U

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.

V

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

W

workflow. (1) 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. (2) 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

WSM. Working Set Manager

X

XCF. Cross-system coupling facility

XES/CF. See *Cross-system Extended Services/Coupling Facility*.

Index

A

B

C

D

E

F

G

H

I

N

O

P

Q

R

S

T

U

V

W

X

Z

File Number: S370/S390-34
Program Number: 5647-A01

Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

SC28-1949-04