

z/OS
3.2

Data Gatherer User's Guide



Note

Before using this information and the product it supports, read the information in [“Notices” on page 103.](#)

This edition applies to IBM® z/OS® 3.2 (5655-ZOS) and to all subsequent releases and modifications until otherwise indicated in new editions.

Last updated: 2025-09-30

© **Copyright International Business Machines Corporation 1990, 2025.**

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

Contents

Figures.....	vii
Tables.....	ix
About this document.....	xi
Who should use this document.....	xi
z/OS Data Gatherer library.....	xi
z/OS RMF library.....	xi
z/OS information.....	xi
How to read syntax diagrams.....	xii
Symbols.....	xii
Syntax items.....	xii
Syntax examples.....	xiii
How to provide feedback to IBM.....	xv
Summary of changes.....	xvii
Summary of changes for z/OS 3.2.....	xvii
Summary of changes for z/OS 3.1.....	xvii
Part 1. Introduction.....	1
Chapter 1. Performance management in z/OS.....	3
Gathering data.....	4
Short-term data collection with Monitor III.....	4
Snapshot monitoring with Monitor II.....	4
Long-term data gathering with Monitor I and Monitor III.....	4
Data gathering for SMF record producer instances with Monitor P.....	5
Monitored activities.....	5
Reporting data.....	6
Accessing data across the sysplex.....	6
Sysplex Data Server.....	6
Sysplex data services for SMF data.....	7
Sysplex data service for Monitor III data.....	8
Sysplex data gathering service for Monitor II data.....	8
Accessing data using REST services.....	8
Enhancing z/OS observability with z/OS OpenTelemetry Emitter.....	10
Part 2. Administration.....	13
Chapter 2. Setting up z/OS Data Gatherer.....	15
Migrating from previous releases.....	15
Customizing the system environment.....	15
Define z/OS Data Gatherer library authorization.....	15
Check the program properties table (PPT).....	16
Remove GRBMFRES.....	16
Global performance data control authority.....	16
Specifying access definitions.....	17
Define user IDs and ensure access to z/OS UNIX System Services.....	17

Assign started task procedures to user IDs.....	17
Controlling access to data for sysplex data services.....	17
Setting up the RMF control session and monitoring sessions.....	19
Customizing the RMF control session.....	19
Specifying the priority for data gatherer started tasks.....	20
Storing gatherer defaults.....	20
Setting up the Monitor III gatherer session RMFGAT.....	20
Defining VSAM data sets.....	21
Ensuring common storage tracking.....	22
Synchronizing SMF recording intervals.....	22
Defining SMF record writing.....	22
Defining SMF synchronization.....	23
Storing gatherer options.....	24
Naming parmlib option members.....	24
Generalizing parmlib members.....	25
Defining parameters for Monitor I.....	26
Defining parameters for Monitor II.....	28
Defining parameters for Monitor III.....	30
Defining parameters for Monitor P.....	30
Setting up z/OS Data Gatherer REST services.....	30
Application properties.....	32
Configuring z/OS Data Gatherer REST services for high availability.....	33
Controlling the usage of SMF record producer service GRBSMFP.....	33
I Chapter 3. Setting up z/OS OpenTelemetry Emitter.....	35
Part 3. Operation.....	37
Chapter 4. Starting and stopping z/OS Data Gatherer.....	39
Starting z/OS Data Gatherer.....	39
Starting z/OS Data Gatherer in the sysplex.....	40
Starting z/OS Data Gatherer without JES.....	41
Stopping z/OS Data Gatherer.....	41
Controlling the SMF buffer.....	42
Controlling sysplex-wide optimized coupling facility hardware data collection (CFOPT).....	43
Chapter 5. Starting and controlling Monitor sessions.....	45
Specifying session options.....	45
Conflicts between session options.....	45
Starting data gatherer sessions.....	45
Starting a specific monitor session.....	45
Modifying data gatherer session options.....	47
Displaying the status of data gatherer sessions.....	48
Stopping data gatherer sessions.....	49
Stopping a specific session.....	49
Chapter 6. How session options are processed.....	51
Processing session options when you start a session.....	51
Displaying the options in effect for a session.....	52
Examples of processing options when you start a session.....	52
Processing session options when you issue the modify command.....	54
Examples of processing options when you issue the MODIFY command.....	54
Part 4. Data gathering reference.....	57
Chapter 7. Long-term data gathering with Monitor I.....	59
Summary of session options.....	59

Default session options.....	60
Descriptions of Monitor I session options.....	61
CACHE.....	61
CHAN.....	62
CPU.....	62
CRYPTO.....	62
CYCLE.....	62
DEVICE.....	62
ENQ.....	64
ESS.....	65
EXITS.....	65
FCD.....	66
INTERVAL.....	66
IOQ.....	67
MEMBER.....	68
OPTIONS.....	69
PAGESP.....	70
PAGING.....	70
RECORD.....	70
STOP.....	70
SYNC.....	71
SYSOUT.....	71
TRACE.....	72
VSTOR.....	77
WKLD.....	78
VMGUEST.....	78
Special considerations.....	78
INTERVAL and CYCLE options.....	78
Device class selection for the DEVICE option.....	79
 Chapter 8. Snapshot data gathering with Monitor II.....	 81
 Chapter 9. Short-term data gathering with Monitor III.....	 83
Summary of gatherer session options.....	83
Default gatherer session options.....	84
Descriptions of Monitor III data gatherer options.....	85
CACHE.....	85
CFDETAIL.....	85
CRYPTO.....	86
CYCLE.....	86
DATASET.....	87
EADM.....	87
HFSNAME.....	88
IOSUB.....	88
LOCK.....	88
MASTER.....	88
MEMBER.....	89
MINTIME.....	90
OPD.....	90
OPTIONS.....	90
PCIE.....	91
RESOURCE.....	92
SCM.....	92
SGSPACE.....	92
STOP.....	93
SYNC.....	93
SYSOUT.....	93
VSAMRLS.....	94

WSTOR.....	94
ZFS.....	95
ZIIPUSE.....	95
Controlling data set recording.....	96
Starting data set support.....	98
Modifying the data set support options.....	99
Stopping data set support.....	100
Data set support for daylight saving time.....	100
Appendix A. Accessibility.....	101
Notices.....	103
Terms and conditions for product documentation.....	104
IBM Online Privacy Statement.....	105
Policy for unsupported hardware.....	105
Minimum supported hardware.....	105
Programming interface information.....	106
Trademarks.....	106
Glossary.....	107
Index.....	113

Figures

1. Sysplex Data Server Data Flow.....	7
2. Overview of z/OS Data Gatherer SMF REST services.....	8
3. Overview of z/OS Data Gatherer Monitor II REST services.....	9
4. Overview of z/OS Data Gatherer Monitor III REST services.....	10
5. z/OS OpenTelemetry Emitter overview.....	11
6. Console sample output with Monitor I OPTIONS in effect.....	69
7. Console output with OPTIONS in effect.....	91

Tables

1. Syntax examples..... xiii

2. Monitored activities and SMF record types..... 5

3. Available shared application properties..... 32

4. Available SMF REST services application properties..... 32

5. Summary of Monitor I Session Options..... 59

6. Monitor I Default Session Options..... 61

7. Where to specify Monitor I options.....69

8. Monitor III data gatherer session options.....83

9. Monitor III default session options..... 84

10. Monitor III OPTIONS command sources..... 91

About this document

z/OS Data Gatherer is the strategic IBM performance measurement tool in a z/OS host environment. It measures selected areas of system activity and records the collected data in System Management Facilities (SMF) records and, in certain cases, optional VSAM data sets. Applications, such as z/OS Resource Measurement Facility (RMF), depend on the data gathered by z/OS Data Gatherer and can display that data in various formats that you can use to evaluate system performance and identify reasons for performance problems.

This publication introduces z/OS Data Gatherer and describes how to set up, customize, and operate it to gather the data needed by RMF and other performance tools.

Who should use this document

This document is intended for use by:

- System administrators and programmers responsible for installing, customizing, and operating z/OS Data Gatherer.
- Performance analysts responsible for measuring and improving system performance
- System operators

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

z/OS Data Gatherer library

The z/OS Data Gatherer library contains the following information units:

- *z/OS Data Gatherer Programmer's Guide*
- *z/OS Data Gatherer User's Guide*

Messages issued by z/OS Data Gatherer are included in *z/OS Resource Measurement Facility Messages and Codes*.

z/OS RMF library

The z/OS RMF library contains the following information units:

- *z/OS Resource Measurement Facility Programmer's Guide*
- *z/OS Resource Measurement Facility User's Guide*
- *z/OS Resource Measurement Facility Report Analysis*
- *z/OS Resource Measurement Facility Messages and Codes* (includes z/OS Data Gatherer messages)

z/OS information

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

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

How to read syntax diagrams

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

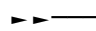
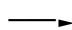
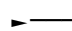
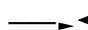
Syntax diagrams pictorially display the order and parts (options and arguments) that comprise a command statement. They are read from left to right and from top to bottom, following the main path of the horizontal line.

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

Symbols

The following symbols may be displayed in syntax diagrams:

Symbol	Definition
--------	------------

	Indicates the beginning of the syntax diagram.
	Indicates that the syntax diagram is continued to the next line.
	Indicates that the syntax is continued from the previous line.
	Indicates the end of the syntax diagram.

Syntax items

Syntax diagrams contain many different items. Syntax items include:

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

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

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

Item type	Definition
-----------	------------

Required	Required items are displayed on the main path of the horizontal line.
Optional	Optional items are displayed below the main path of the horizontal line.
Default	Default items are displayed above the main path of the horizontal line.

Syntax examples

The following table provides syntax examples.

Table 1. Syntax examples

Item	Syntax example
Required item. Required items appear on the main path of the horizontal line. You must specify these items.	➤➤ KEYWORD — required_item ➤➤
Required choice. A required choice (two or more items) appears in a vertical stack on the main path of the horizontal line. You must choose one of the items in the stack.	➤➤ KEYWORD — <div>required_choice1 required_choice2</div> ➤➤
Optional item. Optional items appear below the main path of the horizontal line.	➤➤ KEYWORD — <div>optional_item</div> ➤➤
Optional choice. An optional choice (two or more items) appears in a vertical stack below the main path of the horizontal line. You may choose one of the items in the stack.	➤➤ KEYWORD — <div>optional_choice1 optional_choice2</div> ➤➤
Default. Default items appear above the main path of the horizontal line. The remaining items (required or optional) appear on (required) or below (optional) the main path of the horizontal line. The following example displays a default with optional items.	➤➤ KEYWORD — <div>default_choice1 optional_choice2 optional_choice3</div> ➤➤
Variable. Variables appear in lowercase italics. They represent names or values.	➤➤ KEYWORD — <i>variable</i> ➤➤
Repeatable item. An arrow returning to the left above the main path of the horizontal line indicates an item that can be repeated. A character within the arrow means you must separate repeated items with that character. An arrow returning to the left above a group of repeatable items indicates that one of the items can be selected, or a single item can be repeated.	➤➤ KEYWORD — <div>repeatable_item</div> ➤➤ ➤➤ KEYWORD — <div>,repeatable_item</div> ➤➤ ➤➤ KEYWORD — <div>repeatable_item</div> ➤➤
Fragment. The fragment symbol indicates that a labelled group is described below the main syntax diagram. Syntax is occasionally broken into fragments if the inclusion of the fragment would overly complicate the main syntax diagram.	➤➤ KEYWORD — <div>fragment</div> ➤➤ fragment ➤➤ — <div>,required_choice1 ,required_choice2 <div>default_choice optional_choice</div></div> ➤➤

How to provide feedback to IBM

We welcome any feedback that you have, including comments on the clarity, accuracy, or completeness of the information. For more information, see [How to send feedback to IBM](#).

Summary of changes

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

Note: IBM z/OS policy for the integration of service information into the z/OS product documentation library is documented on the z/OS Internet Library under [IBM z/OS Product Documentation Update Policy](http://www.ibm.com/docs/en/zos/latest?topic=zos-product-documentation-update-policy) (www.ibm.com/docs/en/zos/latest?topic=zos-product-documentation-update-policy).

Summary of changes for z/OS 3.2

The following content is new, changed, or no longer included in z/OS 3.2.

New

The following content is new.

September 2025 release

- Information is added in “[Enhancing z/OS observability with z/OS OpenTelemetry Emitter](#)” on page 10 and Chapter 3, “[Setting up z/OS OpenTelemetry Emitter](#),” on page 35. (APAR OA66345 (www.ibm.com/support/pages/apar/OA66345), which also applies to z/OS 3.1)
- The optional ASCBV31=YES parameter is added to the START command in “[Starting z/OS Data Gatherer](#)” on page 39.

Changed

The following content is changed.

September 2025 release

- None.

Deleted

The following content is deleted.

September 2025 release

- None.

Summary of changes for z/OS 3.1

The following content is new, changed, or no longer included in z/OS 3.1.

New

The following content is new.

October 2024 refresh

- The topic, “[Monitor III REST services](#)” on page 9, is added, and information about setting up Monitor III REST services is added in “[Setting up z/OS Data Gatherer REST services](#)” on page 30 and “[Application properties](#)” on page 32. (APAR OA65072)

September 2023 release

- Information about the SMF record producer service has been added.

Changed

The following content is changed.

June 2024 refresh

- Updates are made in “Setting up z/OS Data Gatherer REST services” on [page 30](#) to identify the 2 applicable chapters in *IBM z/OS Management Facility Configuration Guide*.

September 2023 release

- None.

Deleted

The following content was deleted.

September 2023 release

- None.

Part 1. Introduction

The introduction provides an overview of z/OS Data Gatherer, including:

- Data gathering with Monitor I, Monitor II, Monitor III, and Monitor P
- Cross-sysplex data sharing with Sysplex Data Server
- Low-level Assembler interfaces for SMF and Monitor III data access

Chapter 1. Performance management in z/OS

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 perform these tasks. Applications, such as z/OS Resource Measurement Facility (RMF), use the data gathered by z/OS Data Gatherer to help each of these people do the job effectively.

z/OS Data Gatherer consists of the following components:

- Monitor I
- Monitor II
- Monitor III
- Monitor P
- Sysplex Data Server
- z/OS Data Gatherer REST services
- Low-level Assembler interfaces for SMF as well as for Monitor II and III data access
- Low-level Assembler interface for SMF record producer

These components work together to provide the foundational capabilities that you need for performance management.

z/OS Data Gatherer and RMF

Starting with z/OS 2.5, the data gathering functionality provided by z/OS Data Gatherer has been separated from the information display and reporting functionality provided by RMF. However, because of the historical relationship between z/OS Data Gatherer and RMF, you will still see references to "RMF" even within a data gathering context, such as:

- The RMF control session
- The RMF and RMFGAT started task procedure names, address space names, and user IDs
- The START, STOP, and MODIFY operator commands for controlling sessions and session options
- The names of certain parameters in APIs, such as the sysplex data services
- The names of certain fields in z/OS Data Gatherer data tables
- The names of certain SMF records, record sections, and fields

z/OS Data Gatherer, z/OS Advanced Data Gatherer, and RMF provide the following levels of functionality:

- The **z/OS Data Gatherer** base element provides the following basic functionality:
 - Ability to write SMF type 70 subtype 1 records to the SMF buffer and to SMF data sets or log streams.
 - Access to Monitor II and SMF type 70 subtype 1 record data via the Sysplex Data Server API.
 - Access to any SMF record type and Monitor III data from SMF and Monitor III VSAM data sets via the z/OS Data Gatherer API, GRBSMFR, and the Sysplex Data Server API, ERB3XDRS.
 - Access to the SMF record producer service, which application programs can use to write application-defined SMF records.
- The **z/OS Advanced Data Gather** optional, priced feature, when enabled, adds the following functionality:
 - System administrators can choose to write SMF type 70 - 78 records to the SMF buffer and to SMF data sets or log streams.
 - Application programs can access collected Monitor II, Monitor III, and SMF type 70 - 78 record data via the Sysplex Data Server and z/OS Data Gatherer APIs.

- The **z/OS Resource Measurement Facility (RMF)** optional, priced feature, when enabled, generates performance reports based on the data provided by the z/OS Advanced Data Gatherer feature. Enabling the RMF feature also implicitly enables the Advanced Data Gatherer feature, thus providing the same full functionality for data gathering and reporting as before. Generation of RMF reports with sysplex scope is only supported if the RMF feature is enabled on all systems in a Parallel Sysplex.

For details about these and other z/OS base elements and optional features, see *z/OS Introduction and Release Guide* and *z/OS Planning for Installation*.

Gathering data

z/OS Data Gatherer gathers data using the following monitors:

- Short-term data collection with Monitor III
- Snapshot monitoring with Monitor II
- Long-term data gathering with Monitor I and Monitor III
- Data gathering for SMF record producer instances with Monitor P

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

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

Short-term data collection with Monitor III

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

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

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

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

Snapshot monitoring with Monitor II

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

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

Long-term data gathering with Monitor I and Monitor III

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

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

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

Data gathering for SMF record producer instances with Monitor P

Applications can delegate data gathering and SMF recording to the z/OS Data Gatherer by registering an SMF record producer instance with the z/OS Data Gatherer. When an application registers a new instance, an SMF record producer descriptor must be provided to the z/OS Data Gatherer. The descriptor contains instructions that are used by Monitor P to determine from where measurements shall be gathered and into which SMF record type and SMF data section field these measurements shall be written.

Monitor P starts data gathering for a producer instance at the end of each SMF interval. When the Data Gatherer recognizes the signal for SMF interval end, Monitor P starts as many SMF record producer instances as needed to process all registered SMF record producer descriptors. The gathered data is finally written into an SMF record without any interaction of the application that registered the SMF record producer instance. When the application submits a deregistration request, no SMF record producer instance is started anymore at SMF interval end.

Monitored activities

z/OS Data Gatherer stores data in SMF and VSAM record data:

- Monitor I, II, and III write SMF records (types 70–79) if you define the appropriate SMF recording options.
- Monitor P writes any SMF record type data as requested by applications that have registered SMF record producer instances with the z/OS Data Gatherer.
- In addition, Monitor III writes VSAM records to in-storage buffers or into VSAM data sets.

Table 2 on page 5 provides the following information:

- Displays the SMF record types of all records that are written by data gatherer sessions
- Indicates all Monitor III data stored in VSAM data sets
- Shows all corresponding reporting components

Table 2. Monitored activities and SMF record types								
Gathering				Activity	Reporting			
Short-term Mon III		Snapshot Mon II	Long-term Mon I		Interactive Mon III	Snapshot Mon II	Real-time Mon I	Long-term Post-processor
SMF	VSAM	SMF	SMF					
	*	79.1/2/5		Address space	*	*		*
	*		74.5	Cache	*			*
	*	79.12	73	Channel path	*	*	*	*
74.4	*			Coupling facility	*			*
	*		70.2	Cryptographic hardware	*		*	*
	*	79.9	74.1	Device	*	*	*	*
74.10	*			EADM Activity	*			*
	*			Enclave	*			
	*	79.7	77	Enqueue	*	*	*	*
			74.8	Enterprise Storage Server (ESS)				*
			74.7	FICON director				*
		79.15		IRLM long locks		*		
	*	79.14	78.3	I/O queuing	*	*	*	*
		79.11	75	Page data set		*	*	*
		79.4	71	Paging		*	*	*

Table 2. Monitored activities and SMF record types (continued)								
Gathering				Activity	Reporting			
Short-term Mon III		Snapshot Mon II	Long-term Mon I		Interactive Mon III	Snapshot Mon II	Real-time Mon I	Long-term Post-processor
SMF	VSAM	SMF	SMF					
74.9	*			PCIE Activity	*			*
	*	79.3	70.1	Processor	*	*	*	*
		79.6		Reserve		*		*
72.5				Serialization Delay				*
72.4	*	79.3		Storage	*	*		*
			76	System counters			*	*
74.3/6	*			UNIX	*	*		*
	*		78.2	Virtual storage	*		*	*
	*		72.3	Workload Service classes and report classes	*			*
74.2	*			XCF	*			*
	*			zFS	*			

Reporting data

Reporting capabilities are provided by RMF. For more information about RMF reporting components, see *z/OS Resource Measurement Facility User's Guide*.

Accessing data across the sysplex

The topics that follow describe how to access performance data across the sysplex.

Sysplex Data Server

The Sysplex Data Server is a distributed function of z/OS Data Gatherer. 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. The data gatherer uses this sysplex communication method to provide access to distributed performance measurement data from any point in the sysplex.

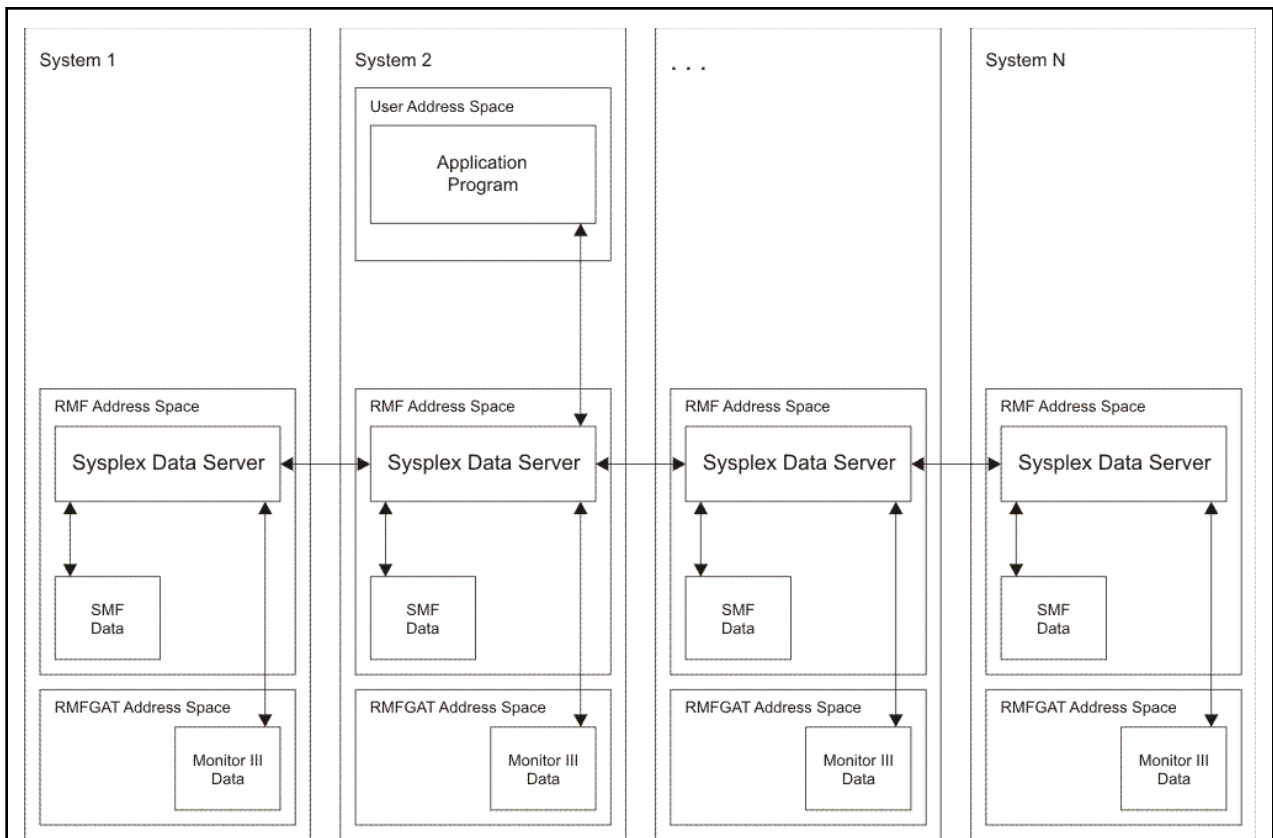


Figure 1. Sysplex Data Server Data Flow

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

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

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

To call the services for SMF data, you need authorization to access the SMF data.¹ For details, see [“Controlling access to data for sysplex data services”](#) on page 17.

Sysplex data services for SMF data

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

¹ Authorization of application programs is provided by the z/OS Security Server RACF®, or products with similar functions, that define the user group authorized to access measurement data returned by the Sysplex Data Server callable services. The services may be invoked by programs running under any PSW key and in problem state.

Sysplex data service for Monitor III data

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

Sysplex data gathering service for Monitor II data

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

Accessing data using REST services

You can access z/OS Data Gatherer and Function Registry for z/OS data by using REST calls via HTTP. To use Data Gatherer REST services, you must set up a separate z/OSMF managed Liberty server instance.

SMF REST services

The SMF REST services provide access to SMF data via REST calls. As shown in [Figure 2 on page 8](#), SMF REST services run inside the separate z/OSMF managed Liberty server instance. SMF REST services provide several REST endpoints which can be called by client applications. SMF REST services invoke the z/OS Data Gatherer interface, GRBSMF, to retrieve SMF data set metadata and SMF record data, depending on the request input. The application then processes the data into a JSON representation of the requested SMF data.

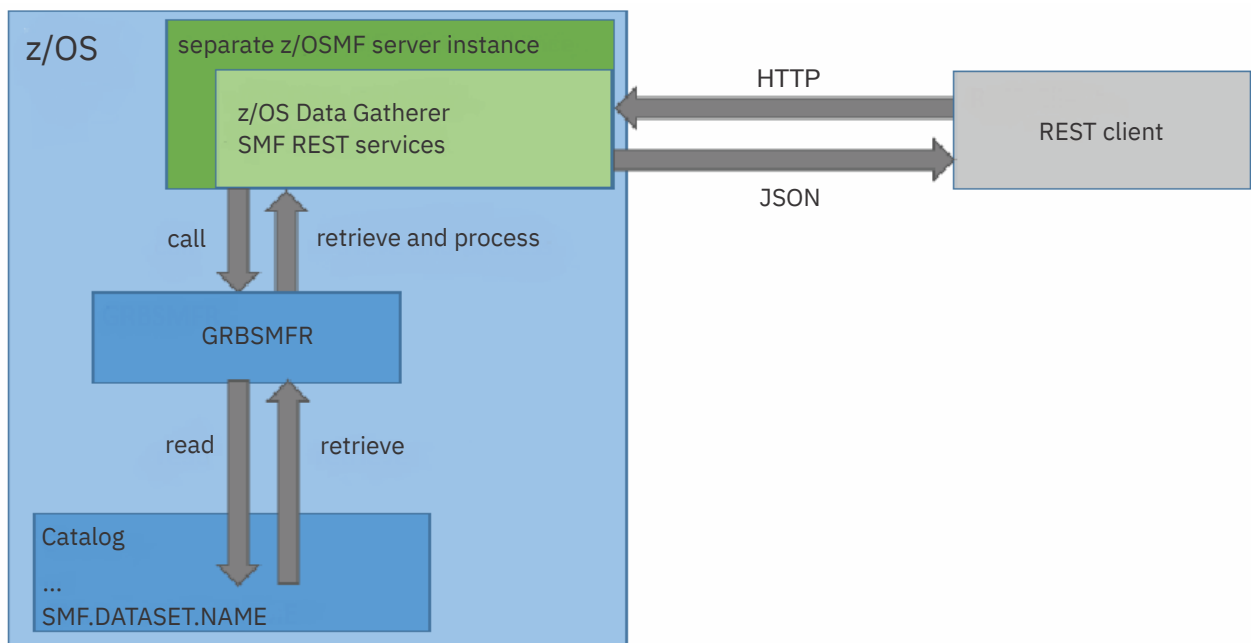


Figure 2. Overview of z/OS Data Gatherer SMF REST services

To retrieve SMF data on a call to SMF REST services, the caller's user ID must have authorization to access the SMF data.

Monitor II REST services

The Monitor II REST services provide access to IBM Function Registry for z/OS data via REST calls. As shown in [Figure 3 on page 9](#), Monitor II REST services runs within the separate z/OSMF managed

Liberty server instance. Monitor II REST services provide several REST endpoints which can be called by client applications. z/OS Data Gatherer Monitor II REST services read and retrieve data from the Function Registry for z/OS from any system in the sysplex that supports this function. The application then processes the data into a JSON representation of the requested function registry data.

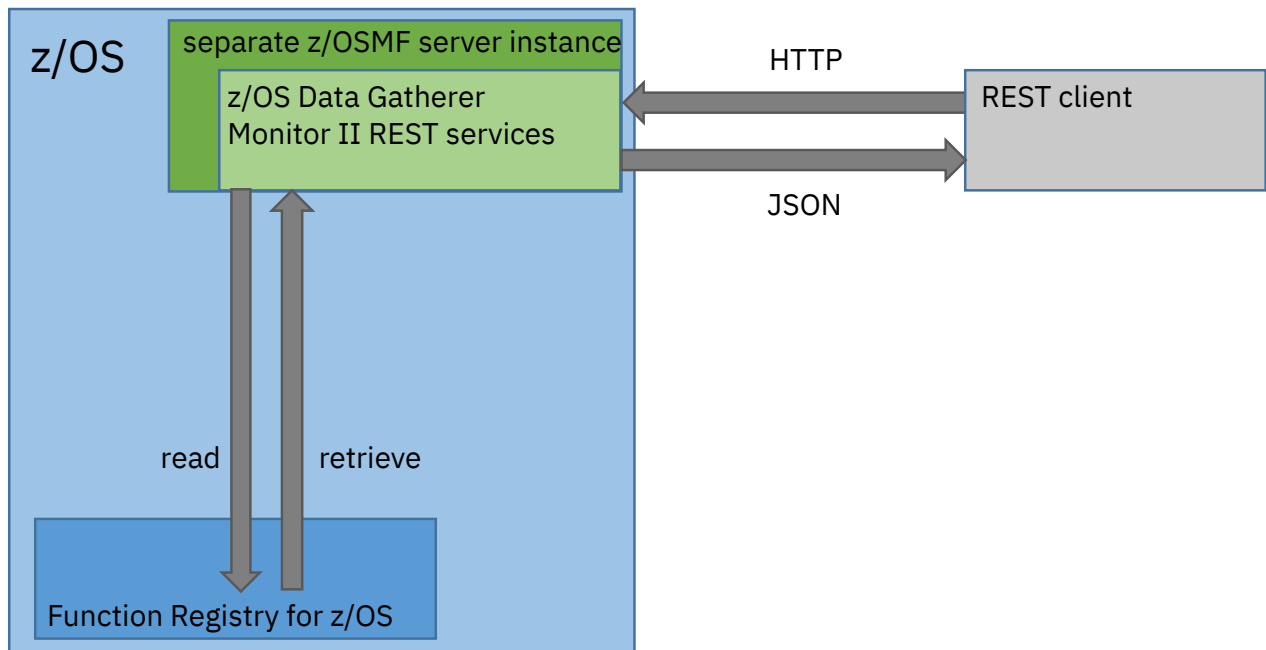


Figure 3. Overview of z/OS Data Gatherer Monitor II REST services

To retrieve Monitor II data on a call to Monitor II REST services, the caller's user ID must have authorization to access the Monitor II data.

For an overview of IBM Function Registry for z/OS, see [Using and registering functions with the IBM Function Registry for z/OS](#) in *z/OS MVS Programming: Authorized Assembler Services Guide*.

Monitor III REST services

The Monitor III REST services provide access to Monitor III VSAM data sets or local in-storage buffer via REST calls. As shown in Figure x on page x, Monitor III REST services run within the separate z/OSMF managed Liberty server instance. Monitor III REST services provide several REST endpoints which can be called by client applications. z/OS Data Gatherer Monitor III REST services read and retrieve data from any system with an active Monitor III session (local in-storage buffer) or from existing Monitor III VSAM data sets. Data is retrieved by using sysplex data server, ERB3XDRS. The application returns the requested Monitor III data in a JSON representation.

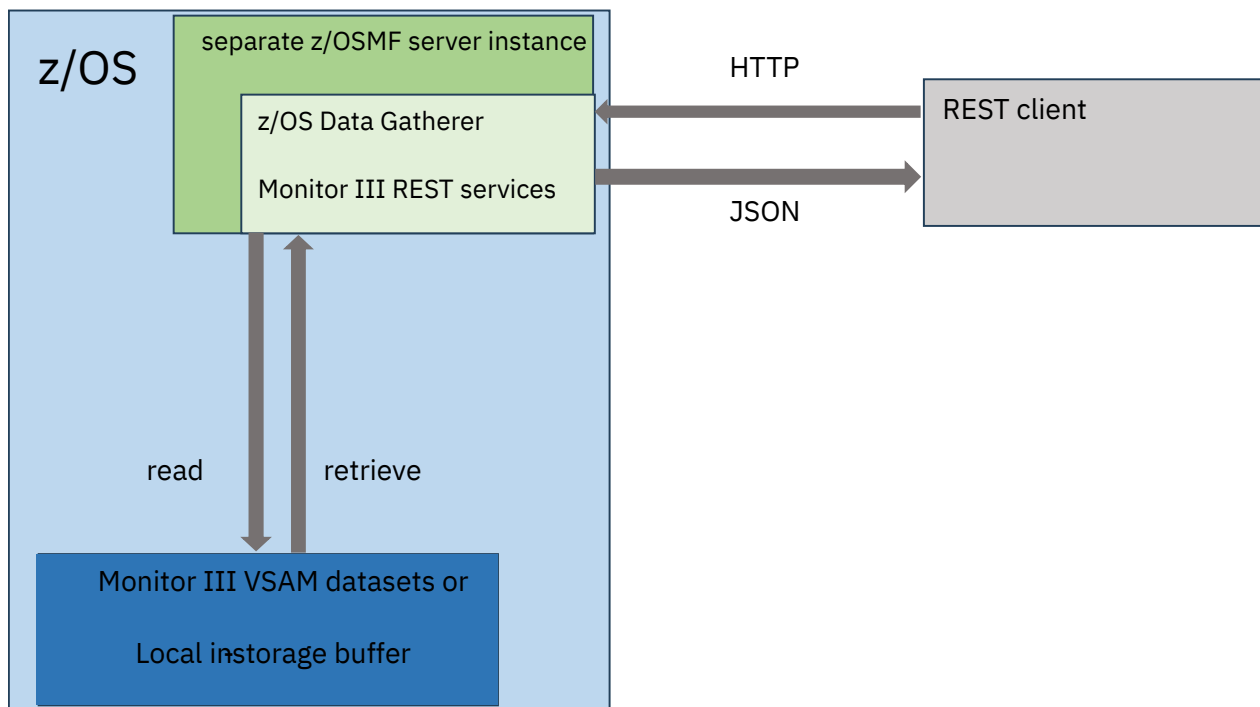


Figure 4. Overview of z/OS Data Gatherer Monitor III REST services

Retrieval of Monitor III data by using Monitor III REST services only works if the caller's user ID has authorization to access Monitor III data. If Monitor III data is read from VSAM data sets, the caller's user ID must have the authorization to access Monitor III data and Monitor III VSAM data sets.

For an overview of the z/OS Data Gatherer Monitor III sysplex data retrieval service, ERB3XDRS, see [ERB3XDRS — Monitor III sysplex data retrieval service in z/OS Data Gatherer Programmer's Guide](#).

Enhancing z/OS observability with z/OS OpenTelemetry Emitter

The z/OS OpenTelemetry Emitter provides a native z/OS interface that generates and forwards OpenTelemetry Protocol (OTLP) compliant distributed tracing data to OpenTelemetry collectors and observability backends. This enables CICS®, Db2®, IMS, MQ, and compatible vendor applications to fully participate in the observability ecosystem.

As shown in [Figure 5 on page 11](#), the z/OS OpenTelemetry Emitter is a standalone Java™ application that runs in its own address space. It reads SMF records containing encoded OTLP span information from SMF in-memory resources, transforms the data into an OTLP-compliant format, and sends the data out using OTLP/http or OTLP/gRPC to an endpoint.

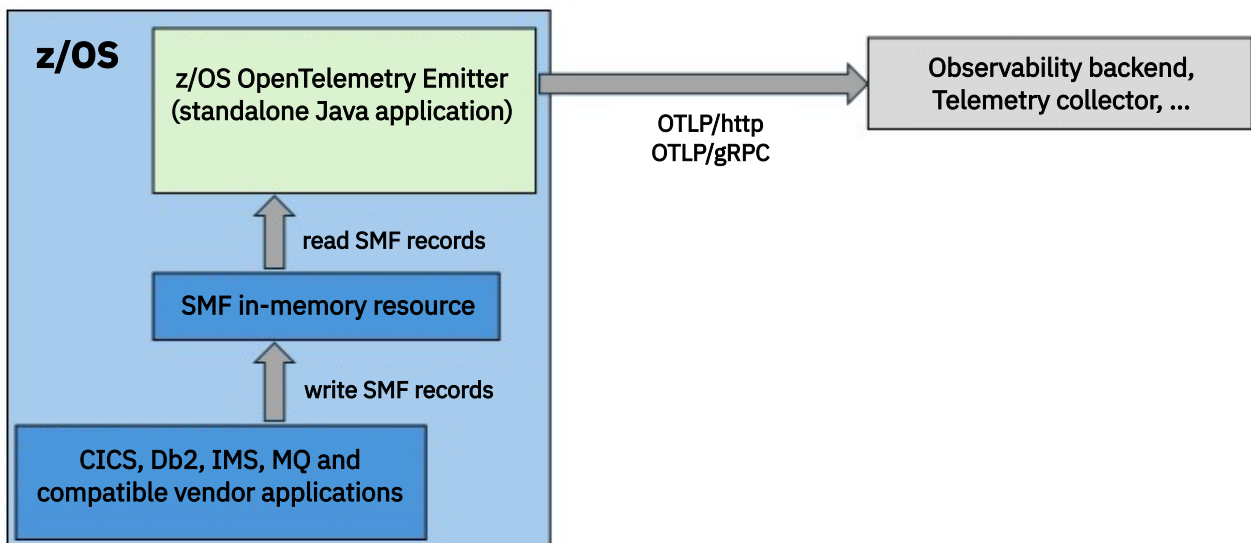


Figure 5. z/OS OpenTelemetry Emitter overview

For more information, see [Chapter 3, “Setting up z/OS OpenTelemetry Emitter,”](#) on page 35 and [z/OS OpenTelemetry Emitter](#) in [z/OS Data Gatherer Programmer's Guide](#).

Part 2. Administration

Administration is what you have to do after installing z/OS Data Gatherer and before you start using it for measuring resources. The administrator creates the prerequisites that the daily users will use, such as setting up job control procedures and defining standard data sets for the installation.

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

A data gatherer administrator can:

- Define *system parameters* and *access definitions* 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 data gatherer monitors store the data that they collect.
- Synchronize *SMF recording intervals* with data collection intervals of the data gatherer monitors, to obtain comparable measurements.
- Define *VSAM data sets* for storing data from Monitor III gatherer sessions.
- Define the *parmlib members* for the data gatherer 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 but can be modified as required. The options that can be included in the members are described in detail in [Part 4, “Data gathering reference,” on page 57](#).

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

Chapter 2. Setting up z/OS Data Gatherer

After installing z/OS Data Gatherer, you must perform the following administration tasks:

- Migrating from previous releases to the release described in this document
- Activating data gatherer functions, including:
 - Customizing the system environment
 - Specifying access definitions
- Setting up the RMF control session, including Monitor I and Monitor II
- Setting up the Monitor III gatherer session (RMFGAT), including defining VSAM data sets for the Monitor III gatherer
- Synchronizing SMF recording intervals
- Setting up parmlib members to store your standard options

Migrating from previous releases

With z/OS 2.5, the RMF Data Gatherer component is moved to a new z/OS base element, z/OS Data Gatherer, and can be ordered with or without the optional, priced feature, z/OS Advanced Data Gatherer. The RMF Reporter component continues to be delivered as the optional, priced feature, RMF. Regardless of whether you want to use RMF functionality in the same way as before, or order only the base element or the z/OS Advanced Data Gatherer feature or both, your SystemPac is built according to your order and you do not have to perform any migration action.

In most cases, it is not necessary to change anything in the procedures and parameters you are using. Typically, new data gatherer options will be activated automatically, and special considerations might be necessary only if you are using customized parmlib members for data gathering.

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

Customizing the system environment

The topics that follow describes the required tasks for customizing the system environment in order to ensure a properly functioning z/OS Data Gatherer on this system.

Define z/OS Data Gatherer library authorization

All z/OS Data Gatherer load modules reside in the two libraries SYS1.SGRBLINK and SYS1.SGRBLPA. If you are activating z/OS Data Gatherer for the first time, you have to define these libraries as APF authorized libraries. You can choose to do it either with or without an IPL.

To activate z/OS Data Gatherer with an IPL:

1. Add the SGRBLINK library to the link list.
2. Add the SGRBLINK library to the APF list.
3. Add the SGRBLPA library to the LPA list.
4. IPL the system.

To activate z/OS Data Gatherer without an IPL:

1. Add the SGRBLINK library to a dynamic link list.
2. Change the APF format to dynamic, if it is not already dynamic.
3. Add the SGRBLINK library to the dynamic APF list.

4. Add the SGRBLPA library to Dynamic LPA.
5. Issue SETPROG commands to make the changes effective.

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

Check the program properties table (PPT)

z/OS provides two default entries in the PPT for the modules, ERBMMFC 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 that you specify a PPT entry for ERBMMFC and ERB3GMFC in a SCHEDxx parmlib member, which must include the defaults and user overrides. Here is a SCHEDxx example:

```
/* PPT Entry for RMF (RMF Control/Monitor I) */
PPT PGMNAME(ERBMMFC) /*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 GRBMFRES

If you are installing z/OS Data Gatherer on a system that already has GRBMFRES (Memory Termination Resource) in the resource manager list, you should remove it, or you will experience performance degradation.

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

Global performance data control authority

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

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

Specifying access definitions

The topics that follow provide information about how to grant the required access rights for setting up z/OS Data Gatherer:

- [“Define user IDs and ensure access to z/OS UNIX System Services” on page 17](#)
- [“Assign started task procedures to user IDs” on page 17](#)
- [“Controlling access to data for sysplex data services” on page 17](#)

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

At first, you should define user IDs that are associated with z/OS data gathering. It is best to define two user IDs that relate to the two started tasks that z/OS Data Gatherer provides for data gathering. (Of course, you may decide to define only one user ID that is assigned to both started tasks).

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

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

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

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

Assign started task procedures to user IDs

In this step, you define the two data-gathering started tasks to RACF and assign them to the user IDs that you created in the previous step.

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

Controlling access to data for sysplex data services

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

z/OS Data Gatherer defines a RACF resource profile, ERBDS.SMFDATA, in the FACILITY class to control access to SMF data in the Sysplex Data Server's SMF buffers. Every user accessing the SMF records in this SMF buffer must be authorized.

ERBDS.SMFDATA

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

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

ERBDS.MON2DATA

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

ERBSDS.MON3DATA

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

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

Controlling the invocation of data reduction exit routines

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

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

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

Security server example

This information unit presents a security server coding example using RACF to achieve the required access and invocation control as described in [“Controlling access to data for sysplex data services”](#) on page 17 and [“Controlling the invocation of data reduction exit routines”](#) on page 18.

1. To activate the resource class FACILITY:

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

2. To define the profile:

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

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

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

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

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

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

4. Activate changes:

Checklist for access to sysplex data services

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

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

Setting up the RMF control session and monitoring sessions

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

- [“Customizing the RMF control session” on page 19](#)
- [“Specifying the priority for data gatherer started tasks” on page 20](#)
- [“Storing gatherer defaults” on page 20](#)

Customizing the RMF control session

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

The RMF control session is the base for data gathering through the different monitors, especially for Monitor I, Monitor II, and Monitor P. If you want to gather data with Monitor III, you also need the cataloged procedure, RMFGAT (see [“Setting up the Monitor III gatherer session RMFGAT” on page 20](#)).

This example shows the RMF procedure as supplied by IBM:

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

PARM

can be used for:

- Specifying the SMF buffer options to be used by the Sysplex Data Server. The format of this option is described in [“Controlling the SMF buffer” on page 42](#). The defaults mean that specifying PARM= ' SMFBUF ' is equivalent to:

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

You can override the values specified or defaulted here by using the SMFBUF option on the START RMF command when starting z/OS Data Gatherer.

- Providing automatic sysplex-wide management for the Distributed Data Server, PARM= ' DDS ' (see [Starting the Distributed Data Server in z/OS Resource Measurement Facility User's Guide](#)).
- Specifying the Monitor I gatherer session options, for example, PARM= ' MEMBER (10) , NOCACHE ' (see [Chapter 7, “Long-term data gathering with Monitor I,” on page 59](#)).
- Activating sysplex-wide optimized coupling facility hardware data collection, PARM= ' CFOPT ' (see [“Controlling sysplex-wide optimized coupling facility hardware data collection \(CFOPT\)” on page 43](#)).

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

To have z/OS Data Gatherer read the ERBRMFxx members from a specific single data set, use a cataloged procedure in the following form:

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

IEFPARM

Identifies the data set containing session options. If you specify an IEFARM DD statement in the procedure, z/OS Data Gatherer does not use the logical parmlib concatenation.

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

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

Specifying the priority for data gatherer started tasks

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

Storing gatherer defaults

The Monitor I and Monitor II gatherer sessions require several parameters to define the type of data to be gathered. These parameters are stored in parmlib members, and are used when you start the gatherer session. The parmlib members supplied with z/OS Data Gatherer 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 24](#).

Setting up the Monitor III gatherer session RMFGAT

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

- [“Defining VSAM data sets” on page 21](#)
- [“Ensuring common storage tracking” on page 22](#)

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

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

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

Since z/OS Data Gatherer runs with the NODSI setting in its PPT entry, be aware that the data gatherer does not hold any data set ENQUEUE (major name=SYSDSN, minor name=*dsname*) for the data sets allocated by the RMF procedure. A missing ENQUEUE can mislead functions, such as HSM, that rely on ENQUEUE for data set protection.

In a system without an active JES2 or JES3, you **must** make this preallocation before you start the Monitor III data gatherer (see [“Starting z/OS Data Gatherer without JES” on page 41](#)). You can use a DD DUMMY statement if you do not wish to store the gatherer messages.

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

Defining VSAM data sets

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

You can define up to 100 data sets for use with the data gatherer. You should define at least two data sets, because the gatherer deletes all data in a data set before writing to it, so a single data set would be emptied immediately after it was filled. Monitor III can keep about 1,100 sets of samples in one data set. Based on practical experience, it is best to define 6 VSAM data sets, each with 50 cylinders disk space. On small and medium systems, this will allow for about two days of data.

Note: If you need to change the amount of space later to meet the needs of your installation, add more data sets, but do not make the individual data sets larger. Increasing the size of the data sets may cause the data gatherer to run out of index entries and be unable to fill the additional space.

See the Data Index (see [Using the Data Index \(DI\)](#) in *z/OS Resource Measurement Facility User's Guide*), which tells you exactly the time range of the data that is available in your VSAM data sets. This can help you in defining the appropriate number of data sets.

Sysplex considerations

If you run the Monitor III data gatherer in a sysplex, it is a good practice to select names for the VSAM data sets with the z/OS system name as part of the data set name. Then you can easily use the capability of symbolic names to specify your parmlib members. See [“Generalizing parmlib members”](#) on page 25 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.SGRBCLS, to define the data sets.

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

where:

vsam_dsn

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

volume

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

num_tracks

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

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

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

[“Controlling data set recording”](#) on page 96 tells you how to specify which data sets are to be used for a particular Monitor III gatherer session.

Ensuring common storage tracking

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

```
SET DIAG=01
```

The defaults in the parmlib member DIAG01 are:

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

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

Synchronizing SMF recording intervals

All z/OS Data Gatherer monitors write SMF records if you specify the appropriate gatherer options. The RMF 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 the data gatherer to be synchronized, and for single-system reports, synchronization is recommended. Therefore, you should perform these tasks:

- [“Defining SMF record writing” on page 22](#)
- [“Defining SMF synchronization” on page 23](#)

Defining SMF record writing

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

SMF recording

SMF

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

- Monitor I and Monitor III write record types 70 – 78.
- Monitor II writes record type 79.
- Monitor P writes any SMF record type.

Example:

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

Monitor I / Monitor II

Option RECORD in ERBRMFxx parmlib member specifies SMF record collection.

Monitor III

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

SMF provides specific user exits to control data collection. Ensure that you do not suppress the writing of SMF type 70 – 79 records if you want to create Postprocessor or other reports.

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

Suppressing SMF record writing

If, by default, z/OS Data Gatherer writes an SMF record type or subtype that you do not want to be written, you can use one of the following methods to suppress SMF recording:

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

For more information, see *z/OS MVS Initialization and Tuning Reference*.

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

Defining SMF synchronization

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

INTVAL(mm)

SMF global recording interval - default is 30 minutes

SYNCVAL(mm)

Synchronization with the hour - default is 00

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

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

SYNC(SMF)

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

SYNC(RMF,mm)

Data gatherer synchronization with the hour

NOSYNC

No synchronization

INTERVAL(mm)

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

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

Monitor I active

Monitor III has the same synchronization as Monitor I

Monitor I inactive

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

Note:

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

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

Recommendation

Use the following values to synchronize SMF record writing:

SMF

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

Monitor I

SYNC(SMF)

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

Storing gatherer options

Perform the following tasks to tailor z/OS data gathering according to your requirements:

- [“Naming parmlib option members” on page 24](#)
- [“Generalizing parmlib members” on page 25](#)
- [“Defining parameters for Monitor I” on page 26](#)
- [“Defining parameters for Monitor II” on page 28](#)
- [“Defining parameters for Monitor III” on page 30](#)
- [“Defining parameters for Monitor P” on page 30](#)

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

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

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

Naming parmlib option members

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

If you do not specify a MEMBER option, z/OS Data Gatherer uses a particular default parmlib member for each type of gatherer session:

- ERBRMF00 for Monitor I
- ERBRMF01 for Monitor II
- ERBRMF04 for Monitor III
- ERBRMF05 for Monitor P

These members are supplied with z/OS Data Gatherer, 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 7, “Long-term data gathering with Monitor I,” on page 59](#), [Details of report commands in z/OS Resource Measurement Facility User's Guide](#), and [Chapter 9, “Short-term data gathering with Monitor III,” on page 83](#), respectively.

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

Syntax rules for ERBRMFxx members

The following syntax rules apply to the ERBRMFxx parmlib members:

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

Parmlib concatenation

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

If you modify members ERBRMF00 through ERBRMF05 according to your requirements, you should store them in a separate parmli to avoid that they will be overwritten unintentionally during the installation of an APAR or a follow-on release.

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

Generalizing parmli members

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

To make this adaptation automatic, z/OS Data Gatherer 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 Data Gatherer user. The **&SYSNAME** symbol resolves to the 8-character system name, and **&SYSCLONE** resolves to the last 2 non-blank characters of the system name.

z/OS Data Gatherer supports the use of symbolic names in the following places:

- All parmli members.

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

Example of using symbolic names for system-specific options:

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

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

Another example of using symbolic names:

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

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

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

Example that shows how to set up gathering options:

In an environment where several systems have access to one and the same storage subsystem, it is sufficient that the cache data gatherer is started just on one system. Running the gatherer on more

than one system creates several copies of identical SMF records type 74-5 (Monitor I) or VSAM records (Monitor III).

Since z/OS Data Gatherer has no sysplex control over the gatherer options, it cannot automatically deselect cache gathering on all but one system. To take advantage of shared parmlib in a sysplex environment, use the symbolic approach offered by z/OS.

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

```
IEASYM CA
```

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

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

- Create a shared 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 **&CACHEOPT** symbol is defined as 'NOCACHE', while on system SYS1, the symbol is resolved as 'CACHE'.

For details about defining your own symbols, refer to [z/OS MVS Initialization and Tuning Reference](#).

- In 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

- In the MODIFY command.

Again, the options can contain symbolic names, as in the parmlib members. The command is converted automatically during system command processing. The system responds to a MODIFY command that contains symbolic names as shown in the following example.

Example of command with symbolic name, and system response:

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

Defining parameters for Monitor I

z/OS Data Gatherer ships a default parmlib member (“ERBRMF00” on page 26) and an alternative one (“ERBRMF02” on page 27) to specify Monitor I gatherer options.

ERBRMF00

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

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

The options specified in ERBRMF00 are:

```

/*****
/* PART 1: MEASUREMENTS
/*****
CACHE                /* CACHE STATISTICS
CHAN                 /* CHANNEL STATISTICS
CPU                  /* CPU STATISTICS
CRYPTO                /* CRYPTO HARDWARE
DEVICE(DASD)         /* DIRECT ACCESS DEVICES MEASURED
DEVICE(NOTAPE)       /* NO TAPE DEVICES MEASURED
DEVICE(NOCHRDR)      /* NO CHARACTER READER DEVICES MEASURED
DEVICE(NOUNITR)      /* NO UNIT RECORD DEVICES MEASURED
DEVICE(NOCOMM)       /* NO COMMUNICATION DEVICES MEASURED
DEVICE(NOGRAPH)      /* NO GRAPHICS DEVICES MEASURED
DEVICE(NONMBR)       /* NO SELECTION BY DEVICE NUMBERS
DEVICE(NOSG)         /* NO SELECTION BY STORAGE GROUPS
NOENQ                /* NO ENQUEUES MEASURED
NOESS               /* NO ENTERPRISE DISK SYSTEMS MEASURED
NOFCD               /* NO FICON DIRECTOR MEASURED
IOQ(DASD)           /* DASD I/O QUEUEING MEASURED
IOQ(NOTAPE)         /* NO TAPE I/O QUEUEING MEASURED
IOQ(NOCHRDR)        /* NO CHARACTER READER I/O QUEUEING
IOQ(NOUNITR)        /* NO UNIT RECORD DEVICE I/O QUEUEING
IOQ(NOCOMM)         /* NO COMMUNICATION I/O QUEUEING
IOQ(NOGRAPH)        /* NO GRAPHICS DEVICE I/O QUEUEING
IOQ(NONMBR)         /* NO SELECTIVITY BY LCU NUMBERS
PAGESP             /* PAGE DATASET STATISTICS
PAGING             /* PAGING DATA
NOTRACE            /* NO TRACE REPORT
VSTOR(S)           /* VIRTUAL STORAGE SUMMARY DATA
WKLD               /* WORKLOAD MANAGER DATA
NOVMGUEST          /* NO CPU DISPATCH TIMES FOR Z/VM GUEST
/*****
/* PART 2: TIMING
/*****
CYCLE(1000)         /* SAMPLE EVERY SECOND (1000 MSEC)
NOSTOP              /* ACTIVE UNTIL OPERATOR ISSUES STOP
SYNC(SMF)           /* USE INTVAL/SYNVAL FROM SMFPRMXX
/*****
/* PART 3: REPORTING / RECORDING OF DATA
/*****
NOOPTIONS           /* OPTIONS NOT DISPLAYED, NO REPLY
RECORD              /* WRITE SMF RECORDS EVERY INTERVAL
NOREPORT            /* NO WRITTEN REPORTS TO SYSOUT
SYSOUT(A)           /* REPORTS TO CLASS A, IF REPORT
/*****
/* PART 4: USER EXITS
/*****
NOEXITS             /* DO NOT TAKE USER EXITS
/*****
```

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

ERBRMF02

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

The options specified in ERBRMF02 are:

```

/*****
/* PART 1: MEASUREMENTS
/*****
CACHE                /* CACHE STATISTICS
CHAN                 /* CHANNEL STATISTICS
CPU                  /* CPU STATISTICS
CRYPTO                /* CRYPTO HARDWARE
DEVICE(DASD)         /* DIRECT ACCESS DEVICES MEASURED
```

```

DEVICE(TAPE)           /* TAPE DEVICES MEASURED */
DEVICE(CHRDR)          /* CHARACTER READER DEVICES MEASURED */
DEVICE(UNITR)          /* UNIT RECORD DEVICES MEASURED */
DEVICE(COMM)           /* COMMUNICATION DEVICES MEASURED */
DEVICE(GRAPH)          /* GRAPHICS DEVICES MEASURED */
DEVICE(NONMBR)         /* NO SELECTION BY DEVICE NUMBERS */
DEVICE(NOSG)           /* NO SELECTION BY STORAGE GROUPS */
ENQ(SUMMARY)           /* ENQUEUES MEASURED */
ESS(LINK,RANK)         /* ENTERPRISE DISK SYSTEMS MEASURED */
FCD                    /* FICON DIRECTOR MEASURED */
IOQ(DASD)              /* DASD I/O QUEUEING MEASURED */
IOQ(TAPE)              /* TAPE I/O QUEUEING MEASURED */
IOQ(CHRDR)             /* CHARACTER READER I/O QUEUEING */
IOQ(UNITR)             /* UNIT RECORD DEVICE I/O QUEUEING */
IOQ(COMM)              /* COMMUNICATION I/O QUEUEING */
IOQ(GRAPH)             /* GRAPHICS DEVICE I/O QUEUEING */
IOQ(NONMBR)            /* NO SELECTIVITY BY LCU NUMBERS */
PAGESP                /* PAGE DATASET STATISTICS */
PAGING                /* PAGING DATA */
TRACE(RCVUICA,END)     /* TRACE 'UIC AVERAGE' */
TRACE(RVCPUA,END)     /* TRACE 'CPU USAGE*16' */
TRACE(RCVPTR,END)     /* TRACE 'PAGING RATE' */
VSTOR(D)              /* VIRTUAL STORAGE DETAIL DATA */
WKLD                  /* WORKLOAD MANAGER DATA */
NOVMGUEST              /* NO CPU DISPATCH TIMES FOR Z/VM GUEST */
/*****
/* PART 2: TIMING
*****/
CYCLE(250)            /* SAMPLE EVERY 250 MILLISECONDS */
STOP(8H)              /* STOP AFTER 8 HOURS */
SYNC(SMF)             /* USE INTVAL/SYNCVAL FROM SMFPRMXX */
/*****
/* PART 3: REPORTING / RECORDING OF DATA
*****/
OPTIONS               /* OPERATOR MAY EXAMINE/CHANGE OPTIONS */
RECORD                /* WRITE SMF RECORDS EVERY INTERVAL */
SYSOUT(A)             /* REPORTS TO CLASS A, IF REPORT */
/*****
/* PART 4: USER EXITS
*****/
NOEXITS               /* DO NOT TAKE USER EXITS */

```

Defining parameters for Monitor II

z/OS Data Gatherer ships a default parmlib member (“ERBRMF01” on page 28) and an alternative one (“ERBRMF03” on page 29) to specify Monitor II gatherer options.

ERBRMF01

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

The options specified in ERBRMF01 are:

```

/*****
/* PART 1: MEASUREMENTS
*****/
NOARD                 /* ADDRESS SPACE RESOURCE CONSUMPTION ? */
NOARDJ               /* ARD REPORT FOR PARTICULAR JOB ? */
  ASD                /* ADDRESS SPACE STATE DATA ? */
NOASDJ               /* ASD REPORT FOR PARTICULAR JOB ? */
NOASRM               /* ADDRESS SPACE SRM DATA ? */
NOASRMJ              /* ASRM REPORT FOR PARTICULAR JOB ? */
NOCHANNEL            /* CHANNEL DATA ? */
NODEV                /* DEVICE DATA ? */
NODEVV              /* DEVICE DATA FOR SPECIFIC DEVICE ? */

```

```

NOIOQUEUE          /* I/O QUEUING DATA ? */
NOPGSP             /* PAGE DATASET MEASUREMENTS ? */
NOSENQ            /* SYSTEM ENQUEUE CONTENTION ? */
NOSENQR           /* SYSTEM ENQUEUE RESERVE DATA ? */
NOSPAG            /* SYSTEM PAGING ACTIVITY ? */
NOSRCS            /* SYSTEM REAL STORAGE/CPU/SRM DATA ? */
/*****
/* PART 2: TIMING
*****/
SINTV(30S)         /* SESSION INTERVAL = 30 SECONDS */
STOP(30M)          /* STOP AFTER 30 MINUTES */
/*****
/* PART 3: REPORTING / RECORDING
*****/
NODELTA            /* TOTAL MODE */
NOOPTIONS          /* NO OPERATOR DISPLAY, NO REPLY */
RECORD             /* SMF RECORDING */
REPORT(DEFER)      /* REPORTS PRODUCED AFTER SESSION END */
SYSOUT(A)          /* INTERVAL REPORTS TO CLASS A */
/*****
/* PART 4: USER RECORDING/REPORTING
*****/
NOUSER             /* USER DATA ? */

```

ERBRMF03

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

The options specified in ERBRMF03 are:

```

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

```

Defining parameters for Monitor III

z/OS Data Gatherer ships a default parmlib member ("[ERBRMF04](#)" on page 30) to specify Monitor III gatherer options.

ERBRMF04

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

The options specified in ERBRMF04 are:

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

Defining parameters for Monitor P

z/OS Data Gatherer ships a default parmlib member ("[ERBRMF05](#)" on page 30) to specify Monitor P gatherer options.

ERBRMF05

This is the default parmlib member for the Monitor P data gatherer session. There is no IBM-supplied alternative member for this gatherer.

The options specified in ERBRMF05 are:

```
SYNC(SMF)           /* USE INTVAL/SYNVAL FROM SMFPRMXX   */
NOOPTIONS           /* OPTIONS NOT DISPLAYED, NO REPLY    */
```

Setting up z/OS Data Gatherer REST services

About this task

Complete the following steps to enable the z/OS Data Gatherer REST services.

Procedure

1. Follow the instructions in [Configure z/OS Data Gatherer: SMF REST Services](#), [Configure z/OS Data Gatherer: Monitor II REST Services](#), and [Configure z/OS Data Gatherer: Monitor III REST Services](#) in

IBM z/OS Management Facility Configuration Guide to set up and start a distinct z/OSMF server with its own user file system.

If you are setting up a z/OSMF server for the first time, examine the respective sections, especially the security setup for the z/OSMF server.

2. If you are configuring SMF REST services, enable the z/OSMF server to perform authorization checks for resources in the DATASET class:

- a) Create the SERVER class profile.

```
RDEFINE SERVER (BBG.SECCLASS.DATASET) UACC(NONE)
```

- b) Grant the z/OSMF server user ID access to the SERVER class profile.

```
PERMIT BBG.SECCLASS.DATASET CLASS(SERVER) ID(IZUSVR) ACCESS(READ)
```

IZUSVR is the default user ID for the z/OSMF server. If you assign a different user ID to the z/OSMF server started task, specify that user ID instead.

3. If you are configuring Monitor II or Monitor III REST services and you want to prevent unauthorized access to Monitor II or Monitor III data, you must enable user IDs to access that data:

- a) Create the FACILITY class profile.

```
RDEFINE FACILITY (ERBSDS.MON2DATA) UACC(NONE)
```

```
RDEFINE FACILITY (ERBSDS.MON3DATA) UACC(NONE)
```

- b) Grant the user IDs access to the ERBSDS.MONxDATA profile.

```
PERMIT ERBSDS.MON2DATA CLASS(FACILITY) ID(userID) ACCESS(READ)
```

```
PERMIT ERBSDS.MON3DATA CLASS(FACILITY) ID(userID) ACCESS(READ)
```

where *userID* is the user ID that needs to access Monitor II or Monitor III data.

Note: If the separate z/OSMF instance is already configured for your system, you can only enable another component by modifying the `zosmf.json` file, as described in step “1” on page 30.

4. If starting the separate z/OSMF instance for the first time, log in to z/OSMF and enable z/OSMF Data Gatherer REST Services (must be performed by an administrator):

- a) Open the z/OSMF General Settings task.

- b) Enable one or more of the following services:

- z/OS Data Gatherer SMF REST services
- z/OS Data Gatherer Monitor II REST services
- z/OS Data Gatherer Monitor III REST services

The following warning should appear:

```
If you enable z/OS Data Gatherer [SMF | Monitor II | Monitor III] REST Services, all
the other non Data Gatherer related services will be disabled after the z/OSMF server
gets restarted.
Do you want to continue?
```

- c) Click **Yes** to confirm.

5. Perform one or more of the following verification steps:

- To verify that the SMF REST services are running, open the Swagger interface at `https://host:port/zosmf/zosdg/smf/`.

Alternatively, access `https://host:port/zosmf/zosdg/smf/v3/api-docs` via a command line tool, such as **curl**, or in a web browser to retrieve the OpenAPI document for the service.

- To verify that the Monitor II REST Services are running, open the Swagger interface at `https://host:port/zosmf/zosdg/m2/`.

Alternatively access `https://host:port/zosmf/zosdgm2/v3/api-docs` via a command line tool, such as **curl**, or in a web browser to retrieve the OpenAPI document for the service.

- To verify that the Monitor III REST Services are running, open the Swagger interface at `https://host:port/zosmf/zosdgm3/`.

Alternatively access `https://host:port/zosmf/zosdgm3/v3/api-docs` via a command line tool, such as **curl**, or in a web browser to retrieve the OpenAPI document for the service.

What to do next

- To start and stop the z/OSMF server instance, see *IBM z/OS Management Facility Programming Guide* and *IBM z/OS Management Facility Configuration Guide*.
- For users to be able to use REST endpoints of the z/OSMF REST Services, they must be connected to the z/OSMF user group in your security product.

Application properties

The following tables list the available application properties. IBM recommends that you use the default settings and only change the properties if required to do so.

The properties can be set in the following files:

- `zdg-smf-rest.properties` file for SMF REST services
- `zdg-m2-rest.properties` file for Monitor II REST services
- `zdg-m3-rest.properties` file for Monitor III REST services

All the named property files reside in the `IZU_CONFIG_DIR` directory. The file encoding must be in UTF-8 format, *not* EBCDIC.

Shared properties

Table 3 on page 32 lists the available shared application properties.

Table 3. Available shared application properties			
Property	Description	Type	Default
<code>zdg.security.realm.name</code>	The realm name used by authentication providers.	string	SAFRealm
<code>zdg.security.authorization.groups</code>	A comma-separated list of SAF groups with OMVS segment and defined group ID (GID) that a user must be a member of in order to access the application functionality. The OMVS GID of the SAF group should be unique.	string	empty

SMF REST services properties

Table 4 on page 32 lists the available SMF REST services application properties.

Table 4. Available SMF REST services application properties			
Property	Description	Type	Default
<code>zdg.bridge.smf.cache.inactivity.lifetime</code>	If a cache entry was not referenced after the specified number of minutes, it will be erased in the next cleanup cycle (scheduled every 5 minutes). Not setting this value disables the cache cleanup.	integer	180

Configuring z/OS Data Gatherer REST services for high availability

Only one z/OSMF instance per Parallel Sysplex must be configured for z/OS Data Gatherer Services. However, it is possible to configure an additional z/OSMF instance for z/OS Data Gatherer REST services on another system in the sysplex to ensure high availability of these services. This can be done in a similar way as for other z/OSMF REST services that are supported for high availability, as described in [Configuring z/OSMF for high availability](#) in *IBM z/OS Management Facility Configuration Guide*.

Controlling the usage of SMF record producer service GRBSMFP

Users of applications that call the SMF record producer service, GRBSMFP, to write SMF record data must have RACF authorization. The z/OS Data Gatherer has defined a RACF resource profile, GRBSMFP.SMFDATA, in the FACILITY class to control usage of the SMF record producer service. Every user writing SMF record data by using GRBSMFP must be granted READ access to this RACF resource profile.

GRBSMFP.SMFDATA

Controls usage of the SMF record producer service, GRBSMFP, that writes SMF record data.

Chapter 3. Setting up z/OS OpenTelemetry Emitter

About this task

Instructions for setting up the z/OS OpenTelemetry Emitter are available in the file system in your z/OS installation at `/usr/lpp/grb/opentelemetry_emitter/dt/README`.

Before you set up the z/OS OpenTelemetry Emitter, IBM strongly suggests that you create a dedicated user ID (for instance, GRBOTELU) and group (for instance, GRBOTELG) to run the z/OS OpenTelemetry Emitter. Read access to the SMF in-memory resource should be restricted to this user ID and group to avoid interference with in-memory resource readout from other address spaces. The following procedure shows the RACF commands to accomplish this.

Note: Data producers (CICS, Db2, IMS, MQ, and compatible vendor applications) own the data written to the SMF in-memory resources and are solely responsible for data correctness in terms of schema and content, as well as for data security concerns, such as confidentiality. The z/OS OpenTelemetry Emitter assumes that input data has been generated by approved, well-known, and trusted producers. The z/OS OpenTelemetry Emitter does not validate, sanitize, or alter any data forwarded via the schema.

Procedure

Complete the following steps to secure the SMF in-memory resources for use with the z/OS OpenTelemetry Emitter.

1. Create the profile for the in-memory resource in the FACILITY class.

```
RDEFINE FACILITY IFASMF.OTEL.* UACC(NONE)
```

where `IFASMF.OTEL.*` is an SMF in-memory resource. The name of the in-memory resource must start with `IFASMF.` as the high-level qualifier.

Note: For more information about SMF in-memory resources, see [Defining in-memory resources in z/OS MVS System Management Facilities \(SMF\)](#).

2. Create a group.

```
ADDGROUP groupID
```

Replace *groupID* with an actual group name of your choice. (Example: GRBOTELG)

3. Permit READ access to the in-memory resource for the group.

```
PERMIT IFASMF.OTEL.* CLASS(FACILITY) ID(groupID) ACCESS(READ)
```

Replace *groupID* with the group name from step 2.

4. Create a user ID for the z/OS OpenTelemetry Emitter address space.

```
ADDUSER userID DFLTGRP(groupID) NOPASSWORD
```

Replace *userID* with an actual user ID (for example, GRBOTELU), and replace *groupID* with the group name from step 2.

5. Activate the changes by refreshing the RACLIST.

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

What to do next

IBM suggests that you classify the z/OS OpenTelemetry Emitter address space to WLM into a medium-importance STC service class, usually importance level 2 and execution velocity 50–60.

■ Proceed with setting up the z/OS OpenTelemetry Emitter.

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 z/OS Data Gatherer 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 gathered data is stored

The monitor options that you can specify on the system commands are described in detail in [Part 4, “Data gathering reference,”](#) on page 57, and [Reporting reference](#) in *z/OS Resource Measurement Facility User's Guide*.

Chapter 4. Starting and stopping z/OS Data Gatherer

The topics that follow explain the following tasks:

- How to start z/OS Data Gatherer:

Use the START system command to start the RMF control session, or to start the control session as well as a Monitor I and Monitor P 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:

z/OS 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 z/OS Data Gatherer:

Use the STOP system command to stop the RMF control session and all sessions.

- How to start and stop control sysplex-wide optimized coupling facility hardware data collection

Starting z/OS Data Gatherer

Enabling the RMF feature or the z/OS Advanced Data Gatherer feature:

The z/OS Advanced Data Gatherer feature and the RMF feature are optional z/OS features that are present in the system whether you have purchased them or not. If you have neither specifically ordered the RMF feature nor the z/OS Advanced Data Gatherer feature but you attempt to use Advanced Data Gatherer functionality, you will receive the following message:

ERB117I ADVANCED DATA GATHERER IS NOT ENABLED TO RUN ON THIS SYSTEM.

If the DDS parameter is specified on the START command and the RMF feature is disabled or not installed, the RMF control session issues the following message:

ERB148I DISTRIBUTED DATA SERVER CANNOT BE STARTED.

It is the responsibility of the system administrator to ensure that RMF or the z/OS Advanced Data Gatherer, if licensed, is enabled to run.

Use the START system command to invoke the RMF cataloged procedure to start z/OS Data Gatherer. Optionally, you can override specifications in the procedure JCL statements by specifying parameters on the START command.

The syntax of the START command is:

```
{START} RMF , , [parm] [ , REUSASID=YES ] [ , ASCBV31=YES ]
{S}
```

parm

Can serve the following purposes:

- Specifying the SMF buffer options (see [“Controlling the SMF buffer”](#) on page 42)
- Specifying the Monitor I gatherer session options (see [Chapter 7, “Long-term data gathering with Monitor I,”](#) on page 59)

If you specify options, each must have the form:

option[(*value*)]

For instance: DEVICE (DASD)

- Providing automatic sysplex-wide management for the Distributed Data Server (see [Starting the Distributed Data Server](#) in *z/OS Resource Measurement Facility User's Guide*)

- Activating sysplex-wide optimized coupling facility hardware data collection (see “Controlling sysplex-wide optimized coupling facility hardware data collection (CFOPT)” on page 43)

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

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

REUSASID=YES

An optional parameter that indicates that a reusable ASID should be assigned to the z/OS Data Gatherer address space.

Tip: IBM suggests that you specify the REUSASID=YES parameter when starting the z/OS Data Gatherer address space.

For more information, see the description of the REUSASID parameter in [Starting a system task from a console in z/OS MVS System Commands](#).

ASCBV31=YES

An optional parameter that indicates that the ASCB that is obtained for z/OS Data Gatherer is to be in 31-bit addressable storage.

Tip: IBM suggests that you specify the ASCBV31=YES parameter when starting the z/OS Data Gatherer address space.

For more information, see the description of the ASCBV31 parameter in [Starting a system task from a console in z/OS MVS System Commands](#).

By default, Monitor I is started along with the RMF control session. If options are specified for *parm*, they will be for Monitor I. To start a Monitor I session using options from the default parmlib member, ERBRMF00, or program defaults, omit this parameter. See [Chapter 6, “How session options are processed,” on page 51](#).

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

The Monitor P session does not start by default; it starts only if you specify **ZDG** for *parm*.

Examples:

- To start the RMF control session and Monitor I session, issue the following system command:
START RMF
- To start the RMF control session and the Monitor I and Monitor P sessions, issue the following command:
START RMF,,,ZDG
- To start the RMF control session and Monitor I session and specify Monitor I options, issue the following command:
START RMF,,,(DEVICE(DASD),CYCLE(500))
- To start RMF with a Monitor I session and an SMF buffer of 32 megabytes in which SMF record types 72–74 are to be stored, specify:
START RMF,,,(SMFBUF(RECTYPE(72:74)))

Starting z/OS Data Gatherer in the sysplex

If you want sysplex-wide reports, you must start the RMF control session and the data gatherer sessions separately on each system of the sysplex. However, the reporting can be done on one system.

IBM strongly suggests that you start z/OS Data Gatherer on all systems in a sysplex with the same options. This is essential for later sysplex reporting. The easiest way to do this is by using the ROUTE command, as shown in the following example:

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

Starting z/OS Data Gatherer without JES

Complete the following steps to run the z/OS data gatherers on a system on which a job entry subsystem (such as JES2) is not active:

1. Preallocate the 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 general messages

RMFSCZZ

For Monitor I session messages

RMFSCIII

For Monitor III session messages

RMFSCxx

For Monitor II background messages from session xx

RMFSCZDG

For SMF record producer instances from the Monitor P session

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 20](#).

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

2. Specify SUB=MSTR on the START command.

Enter the START command in the following format:

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

parm

You can specify other options as described in [“Starting z/OS Data Gatherer” on page 39](#).

SUB=MSTR

Use this specification if JES is not active on your system, and you want to run the z/OS data gatherers.

3. Suppress the printing of reports.

Start the gatherer sessions in the normal way (see [“Starting data gatherer sessions” on page 45](#)) but be sure to specify the NOREPORT option for both Monitor I and Monitor II background sessions.

Stopping z/OS Data Gatherer

Use the STOP system command to end the RMF control session and all active gatherer and background sessions. Any active Monitor II and Monitor III TSO/E sessions remain active. z/OS Data Gatherer issues a message to inform you that it has stopped. For information about stopping individual sessions, see [“Stopping data gatherer sessions” on page 49](#).

The syntax of the STOP command is:

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

Controlling the SMF buffer

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

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

- A size of 32 megabytes
- Collection of SMF record types 70 - 78, including all subtypes

You can override these values by specifying the SMFBUF option in any of the following 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 that starts the RMF control session (see “Setting up the RMF control session and monitoring sessions” on page 19). This overrides the default values.
2. By specifying the SMFBUF option on the START command for the RMF control session. This overrides any PARM specification and the defaults.
3. By specifying the SMFBUF option on the MODIFY command for the RMF control session. This overrides any specifications on the START command and in the cataloged procedure and the 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 the START command nor 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 in the range 0 - 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 type 79 subtype 15 records (for Monitor II IRLM long lock reporting) are written only if you define this explicitly, such as:

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

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

SMFPRMxx in SYS1.PARMLIB

To write SMF type 79 subtype 15 records (for Monitor II IRLM long lock reporting), the IEFU83, IEFU84, and IEFU85 exits must be defined, for example:

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

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

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

Specifying SMFBUF without any options on the MODIFY command 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:

1. Assume you have the following statements in your RMF cataloged procedure:

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

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

2. 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.

3. 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))
```

Controlling sysplex-wide optimized coupling facility hardware data collection (CFOPT)

Beginning with z/OS 2.5, z/OS Data Gatherer provides the ability to use an optimized data collection for coupling facility hardware statistics.

If optimized coupling facility hardware data collection is active, the data gatherer collects coupling facility hardware statistics only on one system in the sysplex. This is either the system that is acting as the Monitor III master gatherer system as determined by the Sysplex Data Server, or a helper system if the Monitor III data gatherer option MASTER with suboption CFHELP is used. (See [“MASTER” on page 88.](#)) This helps reduce the system overhead of coupling facility data collection.

Sysplex-wide optimized coupling facility hardware data collection is activated by using the CFOPT option. This option must be specified on only one system in the sysplex and is recognized by all sysplex members running z/OS 2.5 or later.

Starting optimized coupling facility hardware data collection

You can specify the CFOPT option as follows:

PARM='CFOPT'

To activate the sysplex-wide optimized coupling facility hardware data collection with the EXEC statement of the RMF procedure (see [“Customizing the RMF control session”](#) on page 19).

START RMF,,,CFOPT

To activate the sysplex-wide optimized coupling facility hardware data collection with the START RMF command.

MODIFY RMF,CFOPT

To activate the sysplex-wide optimized coupling facility hardware data collection with the MODIFY RMF command.

Stopping optimized coupling facility hardware data collection

You can stop the sysplex-wide optimized coupling facility hardware data collection by specifying the NOCFOPT option. For instance, the MODIFY RMF,NOCFOPT command deactivates the sysplex-wide optimized coupling facility hardware data collection, regardless of the system on which you enter the command.

Chapter 5. Starting and controlling Monitor sessions

The topics that follow explain how to start and stop z/OS Data Gatherer sessions, to specify and modify session options, and to display status for the following sessions:

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

You can find details about the options and commands for all data gatherer sessions in [Part 4, “Data gathering reference,”](#) on page 57 and in [Reporting reference](#) in *z/OS Resource Measurement Facility User's Guide*.

Specifying session options

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

- The PARM field of the EXEC statement in the RMF cataloged procedure (Monitor I session only. See [“Setting up the RMF control session and monitoring sessions”](#) on page 19 for details).
- The z/OS Data Gatherer parmlib member, or other equivalent data set member containing session options. See [“Storing gatherer options”](#) on page 24 for details.

If you do not specify an option in either the MODIFY command, the PARM field, or the parmlib member, z/OS Data Gatherer uses a program default. From the various specifications, z/OS Data Gatherer forms a list of options for the session. How it does this is described in [Chapter 6, “How session options are processed,”](#) on page 51.

Conflicts between session options

Some options cannot be used concurrently and may cause conflicts. Should any conflicts occur, z/OS Data Gatherer detects the mutually-exclusive options during input merge and selects compatible values for these options; the operator is notified of the selections made. The possible conflicts for each monitor are discussed in [Part 4, “Data gathering reference,”](#) on page 57 and in [Reporting reference](#) in *z/OS Resource Measurement Facility User's Guide*.

Starting data gatherer sessions

Session commands are issued as parameters on the MODIFY system command. Only one Monitor I session and one Monitor P session can be active at a time. Up to 32 non-interactive Monitor II sessions can be active concurrently.

z/OS Data Gatherer provides a cataloged procedure that starts a Monitor III data gatherer session, as described in [“Setting up the Monitor III gatherer session RMFGAT”](#) on page 20. 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 47.

Starting a specific monitor session

Once you have started the RMF control session, use the MODIFY system command to pass the START session command 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 the Monitor I session
- **Two alphanumeric characters other than ZZ** for a Monitor II background session
- **III** for the Monitor III gatherer session
- **ZDG** for the Monitor P 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 messages about that session.

parm

The options for the session. No options can be specified for the Monitor P 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 topics that follow on starting the respective sessions.

If you do not specify session options here, z/OS Data Gatherer takes all options from the parmlib member and program defaults. See [Chapter 6, “How session options are processed,” on page 51](#) for information about how options are processed when you start a Monitor 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 invoke the RMF cataloged procedure, ZZ is automatically assigned as the session ID.

If JES is not active in your system and you have started the RMF cataloged procedure 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 7, “Long-term data gathering with Monitor I,” on page 59](#).

Example: To start the Monitor I session and specify that processor activity is not to be measured and take all other options from other sources, issue the following 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 alphanumeric value, except ZZ.

If JES is not active in your system and you have started z/OS Data Gatherer with the SUB=MSTR option, you must specify the NOREPORT option when starting this gatherer. This and other options and values that you can specify for *parm* are listed in [Details of report commands in z/OS Resource Measurement Facility User's Guide](#).

Examples:

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

```
MODIFY RMF,START AB
```

2. To start a Monitor II background session and specify that reports be produced at the end of the session and that other options be taken from the ERBRMF07 parmlib member, issue the following command:


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

Starting a Monitor III gatherer session

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

z/OS Data Gatherer 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 20). 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 9, “Short-term data gathering with Monitor III,”](#) on page 83.

Examples:

1. To start a Monitor III gatherer session with all options taken from the parmlib member and the program defaults, issue the following command:

```
MODIFY RMF,START III
```

2. To start a Monitor III gatherer session 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)
```

Note: In case of a time change, such as for a switch to or from Daylight Saving Time, RMFGAT will be stopped and restarted automatically to reflect the correct time in the sampled records.

Starting a Monitor P session

The value of *session-id* for the Monitor P session is always ZDG.

Example: To start the Monitor P session, issue the following command:

```
MODIFY RMF,START ZDG
```

Modifying data gatherer session options

You can modify the options in effect for the data gatherer sessions by using the MODIFY command. A changed option remains in effect until the session ends or you issue another MODIFY command to change that option.

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 z/OS Data Gatherer processes options when you modify session options, see [Chapter 6, “How session options are processed,”](#) on page 51.

Examples:

- Monitor I session:

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

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

- Monitor II background session with the AB session ID:

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

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

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

Displaying the status of data gatherer sessions

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

```
{MODIFY} RMF, {DISPLAY} {ACTIVE}
{F}      {D}      {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 IDs and current options for all active sessions are to be displayed.

Examples:

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

```
F RMF,D ZZ
```

- To display the session identifiers and options for all active sessions, issue the following 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 data gatherer sessions

You can end data gatherer sessions in the following ways:

- By issuing the STOP system command, which stops all active background sessions. See [“Stopping z/OS Data Gatherer”](#) on page 41.
- By specifying a time value in the STOP option for a specific session. See [Part 4, “Data gathering reference,”](#) on page 57.
- By issuing a STOP session command to stop a specific session. All other active sessions continue processing. See [“Stopping a specific session”](#) on page 49.

Stopping a specific session

You can end any active session by issuing the following command:

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

session-id

The identifier assigned on the START command for your session.

Issuing the STOP session command forces an immediate end of interval. After interval processing is complete, z/OS Data Gatherer 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 sessions to continue processing, issue the following command:

```
MODIFY RMF,STOP ZZ
```

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

```
MODIFY RMF,STOP AB
```

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

```
MODIFY RMF,STOP III
```

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

```
MODIFY RMF,STOP ZDG
```

Chapter 6. How session options are processed

z/OS Data Gatherer processes session options from various sources in a certain order to create a list of options for a non-interactive session. z/OS Data Gatherer uses a list of options to control each of the following non-interactive sessions:

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

Session options are processed whenever you issue the following commands:

- A START system command to start Monitor I when you start z/OS Data Gatherer
- A START session command to start non-interactive sessions
- A MODIFY session command to modify non-interactive session options

The topics that follow describe how z/OS Data Gatherer processes session options in these situations.

Processing session options when you start a session

When you start a non-interactive session from the operator console, z/OS Data Gatherer processes the options from the following sources, listed 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 START system command to start a Monitor I session 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, z/OS Data Gatherer gives precedence to the options specified in the first member in the list. For instance, if you specify MEMBER(02,07), z/OS Data Gatherer first notes the options from the ERBRMF02 member, then processes those from the ERBRMF07 member. In case of conflicts, z/OS Data Gatherer uses the options from ERBRMF02. This means that if ERBRMF02 specifies ENQ(DETAIL) and ERBRMF07 specifies ENQ(SUMMARY), z/OS Data Gatherer establishes ENQ(DETAIL) for the session.

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

4. The default z/OS Data Gatherer parmlib members.

If you do not specify the MEMBER option in any of the ways listed previously, z/OS Data Gatherer uses the default parmlib members. z/OS Data Gatherer establishes options from the default parmlib members only if they were not specified in any of the higher-priority sources.

5. Program defaults (lowest priority).

z/OS Data Gatherer fills in any options that were not specified anywhere else with a program default. The program defaults for non-interactive session options are described in each respective topic in [Part 4, "Data gathering reference,"](#) on page 57.

If z/OS Data Gatherer 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 instance, z/OS Data Gatherer might detect the Monitor II background session options RECORD on the START command and

NORECORD in a parmlib member. Since z/OS Data Gatherer detected RECORD higher in the priority list, it takes that value.

If z/OS Data Gatherer detects invalid option values, it ignores them and uses the next valid value specified in priority source.

If z/OS Data Gatherer does not find any errors, it issues an informational message indicating that the session is active, and begins session processing.

Displaying the options in effect for a session

If z/OS Data Gatherer detects any errors while processing session options, it displays a list of options in effect for a non-interactive session on 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 in the following ways:

- Issue the DISPLAY session command from the operator console. For information about issuing a DISPLAY session command, see [“Displaying the status of data gatherer sessions” on page 48](#).
- Specify the OPTIONS session option.

Examples of processing options when you start a session

The following examples show how z/OS Data Gatherer processes session options when you start non-interactive sessions.

Example: When you start a Monitor I session

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

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

1. From the options specified on the START system command, z/OS Data Gatherer creates the following option list for the session:

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

2. z/OS Data Gatherer processes the MEMBER(10) option after it processes all other options on the START system command. The ERBRMF10 member contains the following options:

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

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

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

3. z/OS Data Gatherer processes the next option source, the PARM field of the RMF cataloged procedure. The START system command invokes the following user-modified cataloged procedure:

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

z/OS Data Gatherer processes the options specified in the PARM field of the RMF cataloged procedure and the option list is now:

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

z/OS Data Gatherer ignores CYCLE(2000) and DEVICE(NOTAPE) because these options have been filled in by a higher-priority source.

4. z/OS Data Gatherer processes the MEMBER(02) option after it processes all other options on the START system command. The ERBRMF02 member contains the following options:

```
OPTIONS
NOPAGESP
EXITS
```

z/OS Data Gatherer processes the member, and the option list is now:

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

5. z/OS Data Gatherer ignores EXITS specified in the ERBRMF02 member because it already filled that in from a higher priority source. z/OS Data Gatherer adds NOPAGESP from the ERBRMF02 member.
6. Because not all options have been filled in, z/OS Data Gatherer uses program defaults to complete the option list.

Example: 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)
```

1. z/OS Data Gatherer uses two of the three options from the START command to begin the list of session options:

```
DELTA
SINTV(30)
```

2. Because MEMBER (07) is specified on the START command, z/OS Data Gatherer generates the member name ERBRMF07 and locates it in SYS1.PARMLIB. Assume that the ERBRMF07 member contains the following options:

```
ASD          STOP(20)
SINTV(10)    SPAG
OPTIONS      SRCS
```

z/OS Data Gatherer adds all of these options except SINTV(10) to the option list. z/OS Data Gatherer does 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
```

3. To complete the option list, z/OS Data Gatherer proceeds to the IBM-supplied program defaults. (These defaults are indicated in the discussion of each option in [Details of report commands in z/OS Resource Measurement Facility User's Guide.](#))

After adding the defaults, z/OS Data Gatherer 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
```

Processing session options when you issue the modify command

When you use the MODIFY session command to modify the options for a non-interactive session, z/OS Data Gatherer processes the options in a different priority order than when you start a non-interactive session. z/OS Data Gatherer 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. z/OS Data Gatherer 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 on the MODIFY command.

When you specify more than one member, z/OS Data Gatherer 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, z/OS Data Gatherer generates the member names ERBRMF03 and ERBRMF07 and proceeds as follows:

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

z/OS Data Gatherer responds to errors in a MODIFY session command in the same way as in a START session command.

Examples of processing options when you issue the MODIFY command

The following examples show how z/OS Data Gatherer processes session options for non-interactive sessions when you use a MODIFY session command.

Example: 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 following command:

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

the options will be modified as you want.

If, however, the ERBRMF10 member includes the following options:

```
NOCHAN
CPU
NOSTOP
```

and you issue the following command:

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


z/OS Data Gatherer processes the options as follows:

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

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

Example: When you modify Monitor II session options

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 following command, the options will be modified as you want:

```
MODIFY RMF,MODIFY AB,ASD,NOSPAG,STOP(40M)
```

If, however, the ERBRMF09 member includes the following options:

```
ASD  
NOSPAG  
NOSTOP
```

and you issue the following command:

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

z/OS Data Gatherer processes the options as follows:

1. Replace NOSTOP in the current option list with STOP(40M).
2. Read the ERBRMF09 member, compare options from that member with the current options list, and replace NOASD with ASD, SPAG with NOSPAG, and STOP(40M) with NOSTOP.

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

Example: When you modify Monitor III data gatherer options

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

If you issue the following command, the options will be modified as you want:

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

If, however, the ERBRMF10 member includes the following options:

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

and you issue the following command:

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

z/OS Data Gatherer processes the options as follows:

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

In this particular case, the desired STOP(40M) option is not currently in effect. This particular command did not achieve the expected results because any option in a parmlib 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 following command:

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

Part 4. Data gathering reference

This part discusses the z/OS data gathering capabilities and 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 fully described in the appropriate chapters that follow.

Chapter 7. Long-term data gathering with Monitor I

The topics that follow describe the Monitor I 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 session” on page 45](#))
- The PARM field of the EXEC statement in the RMF cataloged procedure (see [“Setting up the RMF control session and monitoring sessions” on page 19](#))
- The Monitor I parmlib member, ERBRMF00 (see [“ERBRMF00” on page 26](#))

z/OS Data Gatherer merges the input to a final set of options for the session. See [Chapter 6, “How session options are processed,” on page 51](#) for details.

Summary of session options

Monitor I creates SMF records type 70 – 78. You can find an overview in [Table 2 on page 5](#).

[Table 5 on page 59](#) gives a summary of the available options, grouped by purpose. There are options for specifying:

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

Table 5. Summary of Monitor I Session Options		
Option	Description	Details on
Activity Options		
<u>CACHE</u> /NOCACHE	Cache activity	“CACHE” on page 61
<u>CHAN</u> /NOCHAN	Channel path activity	“CHAN” on page 62
<u>CPU</u> /NOCPU	Processor activity	“CPU” on page 62
<u>CRYPTO</u> /NOCRYPTO	Cryptographic hardware activity	“CRYPTO” on page 62
<u>DEVICE</u> (<i>type</i>)/NODEVICE	Device activity	“DEVICE” on page 62
$\{ \text{(SUMMARY)} \}$ ENQ $\{ \text{(DETAIL[majorname[,minorname]])} \}$ / NOENQ	Enqueue contention activity	“ENQ” on page 64
ESS[(<i>opt_list</i>)]/NOESS	Enterprise Disk System statistics	“ESS” on page 65
FCD/NOFCD	FICON® director activity	“FCD” on page 66

Table 5. Summary of Monitor I Session Options (continued)		
Option	Description	Details on
<u>IOQ</u> (<i>opt_list</i>)/NOIOQ	I/O queuing activity	“IOQ” on page 67
<u>PAGESP</u> /NOPAGESP	Page data set activity	“PAGESP” on page 70
<u>PAGING</u> /NOPAGING	System paging activity	“PAGING” on page 70
TRACE(<i>variable[,opt_list]</i>)/NOTRACE	Trace variables for the Trace Activity report	“TRACE” on page 72
$\begin{matrix} \{ (S \\ VSTOR \{ (D \quad [, jobname1, jobname2, \dots]) \} / \\ NOVSTOR \end{matrix}$	Virtual storage activity	“VSTOR” on page 77
<u>WKLD</u> /NOWKLD	Workload activity	“WKLD” on page 78
<u>VMGUEST</u> /NOVMGUEST	Processor activity	“VMGUEST” on page 78
Time-frame Options		
$\begin{matrix} \{ (1000) \} \\ CYCLE \{ (nnn) \} \end{matrix}$	The length of the cycle at the end of which Monitor I makes sampling observations	“CYCLE” on page 62
$\begin{matrix} (\{ 30M \}) \\ INTERVAL (\{ nnn [M] \}) \end{matrix}$	The length of the reporting interval in minutes used in combination with the options SYNC(RMF,mm), SYNC(RMF,mmM), or NOSYNC	“INTERVAL” on page 66
<u>STOP</u> (<i>value</i> $\begin{matrix} [M] \\ [H] \end{matrix}$)/NOSTOP	Desired duration of the Monitor I session, in minutes (M), or hours (H)	“STOP” on page 70
$\begin{matrix} \{ (SMF) \} \\ \{ (RMF, mm) \} \\ SYNC \{ (RMF, mmM) \} / NOSYNC \end{matrix}$	Interval synchronization with the SMF or z/OS Data Gatherer interval synchronization on the minute	“SYNC” on page 71
Reporting Options		
$\begin{matrix} \{ OPTIONS \} / \{ NOOPTIONS \} \\ \{ OPTN \} \quad \{ NOOPTN \} \end{matrix}$	Option list for the session to be printed at the operator console	“OPTIONS” on page 69
<u>RECORD</u> /NORECORD	Specifies whether measured data is to be written to SMF records	“RECORD” on page 70
SYSOUT(<i>class</i>)	SYSOUT class to which the formatted printed reports are directed	“SYSOUT” on page 71
Environment Options		
<u>EXITS</u> /NOEXITS	User exit routines to be executed during session processing to gather or report additional data	“EXITS” on page 65
MEMBER(<i>list</i>)	Parmlib member containing Monitor I session options	“MEMBER” on page 68

Default session options


Table 6 on page 61 lists the options that take effect by default. You need only specify contradictory ones:

Table 6. Monitor I Default Session Options	
Default Option	Description
CACHE	Measures cache activity
CHAN	Measures channel path activity
CPU	Measures processor activity
CRYPTO	Measures cryptographic hardware activity
CYCLE(1000)	Takes sample measurements once a second (1000 milliseconds)
DEVICE(DASD)	Measures DASD activity (not other classes of device)
IOQ(DASD)	Measures I/O queuing activity on logical control units for DASD
NOESS	Does not measure Enterprise Disk System statistics
NOFCD	Does not measure FICON director activity
PAGESP	Measures page data set activity
PAGING	Measures system paging activity
RECORD	Writes the measured data to SMF records
VSTOR(S)	Measures virtual storage activity for summary reports
WKLD	Measures system workload
INTERVAL(30M)	Combines data every 30 minutes (value is ignored for SYNC(SMF))
NOENQ	Does not measure contention activity
NOEXITS	Executes no user exits when gathering and reporting
NOTRACE	Does not trace any variables (no Trace Activity report)
NOVMGUEST	Does not measure CPU dispatch times and processor utilization for a z/OS system running as z/VM guest.
OPTIONS	Prints a list of session options at the operator console at the start of the session, allowing the operator to change options. For a fast start-up of Monitor I, we recommend that you specify NOOPTIONS unless changes at start-up are really necessary.
STOP(8H)	Ends the session 8 hours after it was started
SYNC(SMF)	Synchronizes the reporting interval with SMF

Descriptions of Monitor I session options

The following topics describe the Monitor I session options and the syntax for specifying them.

CACHE




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

Cache controller data is gathered by individual device address. There is no indication of which system in the sysplex initiates a recorded event. Therefore, the data can be gathered on any system sharing the cached devices. To avoid having duplicated data, you should gather cache activity data on one system only. Refer to "Example that shows how to set up gathering options" in ["Generalizing parmlib members"](#) on page 25, which shows how to set up gathering options.

To suppress the gathering of cache data, specify **NOCACHE**.


If you wish to monitor this activity, the RECORD option must also be in effect for the session. The RECORD option takes effect by default.

CHAN




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

CPU




Specifies processor activity measurement.

CRYPTO




Specifies cryptographic hardware activity measurement.

CYCLE




Specifies, in milliseconds, the length of the cycle at the end of which sampling observations are to be made, where *value* is the number of milliseconds. The valid range is 50 - 9999 milliseconds. If you specify a value less than 50, z/OS Data Gatherer increases the value to 50. If you specify more than 9999, z/OS Data Gatherer decreases the value to 9999. The default value is 1000 milliseconds. See [“INTERVAL and CYCLE options” on page 78](#) for considerations that apply to choosing a cycle length.

DEVICE



Type



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

Type is one of the following:

- One or more device numbers:

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

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

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

You can not exceed 32 characters, including commas and colons. When you specify a range of numbers, use a colon as a separator to indicate that the report is to consist of all numbers from *[t]bbbb* up to and including *[u]zzzz*.

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

- Any of the following classes:

CHRDR/NOCHRDR

Character reader devices

COMM/NOCOMM

Communications equipment

DASD/NODASD

Direct access storage devices

GRAPH/NOGRAPH

Graphics devices

TAPE/NOTAPE

Magnetic tape devices

UNITR/NOUNITR

Unit record devices

- One or more storage groups:

```
{SG} {([aaaaaaaa { })
       ([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 device data for all of the names from *bbbbbbbb* up to and including *zzzzzzzz* are collected.

If SG was used to specify one or more storage groups, NOSG can be used to clear the list of storage group names.

Monitor I always collects 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 specified. 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 collected.

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

Examples:

- To request data gathering 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 data gathering for magnetic tape devices and DASD you would specify:

```
DEVICE (TAPE)
```

To limit the gathering of DASD, you must specify NODASD and use the NMBR field to identify those devices you want to monitor.

ENQ



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


The optional *major* field can contain the 1- to 8-character major name of a serially reusable resource. Optionally, the major name is followed by a comma and a minor name. The *minor* field can contain the 1- to 44-character minor name of the resource.

Example:

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

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

ESS



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

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

LINK

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

RANK

extent pool statistics and rank statistics are gathered.

NOLINK

no link performance statistics are gathered.

NORANK

no extent pool and rank statistics are gathered.

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

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


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

Example:

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

- ESS
- ESS(LINK,RANK)

EXITS



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

If you have previously specified the **EXITS** option to gather SMF records with the Cache RMF Reporter (CRR) Program Offering (records type 245), this is no longer required with the Monitor I gatherer **CACHE** option. Therefore, you should specify **NOEXIT**, unless you have some other exit routines that you want to activate.

FCD




Specifies whether FICON director activities should be measured.

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

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

INTERVAL



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

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

Note:

z/OS Data Gatherer 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, z/OS Data Gatherer increases the value to one; if you specify more than 60 minutes, z/OS Data Gatherer decreases the value to 60. To synchronize the z/OS Data Gatherer SMF interval to any time within the hour, use the Monitor I SYNC option. See [“Synchronizing SMF recording intervals”](#) on page 22 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, z/OS Data Gatherer sets the STOP value to the INTERVAL value.
2. Monitor I extends INTERVAL in two situations:
 - When the system does not dispatch Monitor I at the end of the interval.

If Monitor I is executing and does not get control within the specified interval length, Monitor I extends the length to 99 minutes. If Monitor I still does not get control within the 99 minutes, data collection is skipped for that interval, and a message is issued to the operator. This can happen when the dispatching priority for z/OS Data Gatherer is too low; see [“Setting up the RMF control session and monitoring sessions”](#) on page 19 for information about 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 monitor session or RMF control session before stopping the processor.

IOQ



Options



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

- One or more LCU numbers:

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

NMBR requests specific logical control unit numbers, where *aaaa*, *bbbb*, and *zzzz* each represent hexadecimal logical control unit numbers. You can omit leading zeros. You can specify any combination of a single logical control unit number, in the format *aaaa*, a list of logical control unit numbers, in the format *aaaa,bbbb*, or a range of numbers in the format *bbbb:zzzz*, where *bbbb* is your first number and *zzzz* is your last number. You can not exceed 32 characters, including commas and colons. When you specify a range of numbers, use a colon as a separator to indicate that data for all LCU numbers from *bbbb* up to and including *zzzz* are collected.

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

- Any of the following classes:

CHRDR/NOCHRDR

Character reader

COMM/NOCOMM

Communications equipment

DASD/NODASD

Direct access storage

GRAPH/NOGRAPH

Graphics

TAPE/NOTAPE

Magnetic tape

UNITR/NOUNITR

Unit record

When you omit the IOQ option, the defaults are as underscored in the preceding list. If you specify IOQ, you must include an option. The option need include only the classes you want to either add to the default (DASD) or the specific LCU number you want data for. The definition of an LCU is model-dependent.

On all processors, an LCU is the set of devices attached to the same physical control unit (or group of control units that have one or more devices in common). Each device belongs to only one LCU, but the I/O processor (SAP), which is part of the channel subsystem, manages and schedules I/O work requests to the various devices within the LCU of the processor.

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

Note: When your system is running as a guest under z/VM, Monitor I cannot gather data for it. In this case, the RMF I/O Queuing Activity report shows only the static configuration data. When your system is running as a guest under an alternate virtual machine, IOQ data gathering is unsupported; the gathering option will be set to NOIOQ.

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

MEMBER



Specifies one or more parmlib members—up to five members can be specified—that contain Monitor I options for the session, where (xx) contains 1 - 5 members, separated by commas. Each member must be a 2-character alphanumeric value. z/OS Data Gatherer then forms the member name by adding the 2-character alphanumeric value to the ERBRMF prefix.


For the Monitor I session, the default is 00, indicating the ERBRMF00 parmlib member. The contents of the ERBRMF00 member are described in “Storing gatherer options” on page 24. If you have created your own parmlib data set, make sure you specify it in the RMF cataloged procedure. See “Setting up the RMF control session and monitoring sessions” on page 19.

If you specify more than one member, z/OS Data Gatherer processes the members' options in left-to-right priority order. For examples on how z/OS Data Gatherer processes session options, see [Chapter 6, “How session options are processed,”](#) on page 51.

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

Note: The MEMBER option must not be used in the z/OS Data Gatherer parmlib members, but is only valid together with a MODIFY console command.

OPTIONS



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

To avoid unnecessary console output and delay in starting the session, specify **NOOPTIONS**. However, if z/OS Data Gatherer detects any syntax errors while processing session options, **OPTIONS** is forced.

Figure 6 on page 69 shows a console output with the Monitor I option **OPTIONS** in effect. For each option, this console output shows the source where the option has been set, for example, -- **COMMAND** means that the option has been set using a **START** or **MODIFY** command.

Table 7 on page 69 explains all possible sources which may appear in a console output.

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


Figure 6. Console sample output with Monitor I **OPTIONS** in effect

Table 7. Where to specify Monitor I options

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


Table 7. Where to specify Monitor I options (continued)	
Source	Where Option is specified
-- DEFAULT	In the program defaults.
-- EXEC	On the EXEC statement in the RMF cataloged procedure.
-- CHANGED	z/OS Data Gatherer changed the option. A message describes the conflict and the change made.
-- MEMBER	In the z/OS Data Gatherer parmlib member.
-- REPLY	The option was changed from the operator console in reply to message ERB306I.

PAGESP




Specifies whether page data set activity is to be measured.

PAGING




Specifies whether system paging activity is to be measured.

RECORD



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

STOP




Specifies the desired duration for the Monitor I session in minutes (M) or hours (H). The valid range is from a minimum of 1 minute to a maximum of 1 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, z/OS Data Gatherer increases the value to one minute. If you specify more than 168 hours, z/OS Data Gatherer decreases 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; z/OS Data Gatherer 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.

SYNC



Specifies whether the interval is to be synchronized with SMF, or on the minute with the z/OS Data Gatherer interval synchronization mechanism.

SYNC(SMF) is the default and specifies that z/OS Data Gatherer will synchronize its interval using the SMF global interval and synchronization values.

The valid range is the number of minutes, 0 - 59 (*mm*), past the hour at which synchronization is to occur. If any value other than 0 - 59 is specified, or the value is omitted, z/OS Data Gatherer assigns a default value of 0. z/OS Data Gatherer synchronizes the interval by shortening the first interval. Subsequent intervals remain synchronized only when the length of the specified interval is a factor of 60. For example, if you specify an interval of 20 minutes synchronized on 10 minutes, reports are generated at 10, 30, and 50 minutes past the hour. Therefore, if you start your session at 9:05, the first interval is shortened so that a report is generated at 9:10. Similarly, if you start your session at 9:15, the first interval is shortened so that a report is generated at 9:30.


NOSYNC specifies that no synchronization is to be performed. Do not specify this if you want to generate sysplex reports.

Note:

1. If you specify SYNC(SMF), do not specify the INTERVAL option at the same time; if you do, z/OS Data Gatherer 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 22](#) for more information.

SYSOUT



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

TRACE



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

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

The valid *variable* values are:

Variable

Value

ASMERRS

Bad slots on local page data sets.

ASMIORQC

Count of I/O requests completed and returned to RSM.

ASMIORQR

Count of I/O requests received by I/O control.

ASMNVSC

Total local slots allocated for non-VIO private area pages.

ASMSLOTS

Total local slots (sum of slots in open local page data sets).

ASMVSC

Total local slots allocated for VIO private area pages.

CCVCPUCT

Number of online CPUs.

CCVEJST

This variable is no longer supported.

CCVENQCT

Number of users non-swappable for enqueue reasons.

CCVRBSTD

Recent base time of day.

CCVRBSWT

Recent base system wait time.

CCVUTILP

System CPU utilization.

LSCTCNT

Current number of logically swapped users for terminal wait.

LSCTMTE

Maximum think time allowed for logical swap candidate.

MCVFRCNT

Number of pages needed to be stolen by force steal routine.

MCVMGAGE

Expanded storage migration age.

MCVSB LTF

Long term percentage of eligible storage that is actually fixed.

MCVSI PR

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.

OMDGO REB

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.

RCEHSPER

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.

RCEHSPW

Total number of hiperspace pages written from central storage to expanded storage.

RCEHSPPI

Total number of hiperspace pages paged in from auxiliary storage.

RCEHSPPO

Total number of hiperspace pages paged out to auxiliary storage.

RCELPAPI

Number of PLPA and PLPA directory pages paged-in.

RCEMVBEL

Number of pages moved from below 16 megabytes in central storage.

RCENWSF

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

RCEPAGMV

Number of times a frame was moved from one frame to another.

RCEPOOL

Number of frames currently available to the system, including frames backing permanent storage (nucleus frames, hardware storage area frames, FLPA frames or fixed BLDL frames), bad frames and offline frames.

RCESPFR

Number of frames available by swap-out without requiring I/O.

RCESWPPI

Total number of pages requiring I/O to swap-in.

RCESWPPO

Total number of pages requiring I/O to swap-out.

RCETOTFX

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

RCETOTPI

Total number of pages paged-in excluding swap-in and VIO page-in.

RCETOTPO

Total number of pages paged-out, excluding swap-out, move-out of VIO pages, and page-out of VIO pages.

RCEVIOME

Number of VIO pages written to expanded storage.

RCEVIOMG

Number of VIO pages migrated from expanded storage to paging data sets.

RCEVIOPI

Total number of VIO pages paged-in, excluding swap-in.

RCEVIOPO

Total number of VIO pages, excluding swap-out, moved out or paged-out.

RCEVIORE

Number of VIO reads from extended storage.

RCEWSDNE

Total number of primary working set pages migrated to auxiliary storage.

RCVAFQA

Average available frame count.

RCVAVQC

AVQ low count.

RCVCPUA

CPU usage average * 16.

RCVFXIOP

Percentage of central storage that is fixed or allocated for paging.

RCVMFXA

Average number of fixed frames for the system.

RCVPAGRT

Total paging rate.

RCVPTR

Paging rate.

RCVSWPTM

Time (in milliseconds) used by ASM to process a request to transfer a group of pages to or from a data set.

RCVUICA

UIC average.

RMCAAWSC

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

RMCADWSC

Detected wait physical swap count.

RMCAEXSC

Exchange on recommendation value swap count.

RMCAFHLD

Number of swaps failed because of an outstanding HOLD SYSEVENT.

RMCAICSC

Improve central storage use.

RMCAIPSC

Improve system paging rate.

RMCALWSC

Long wait physical swap count.

RMCAMRSC

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

RMCANQSC

CPU enqueue exchange swap count

RMCAOISC

OMVS input wait.

RMCAOOSC

OMVS output wait.

RMCARSSC

Central storage shortage swap count.

RMCATISC

Terminal input swap count

RMCATOSC

Terminal output swap count.

RMCATSSC

Count of transition swaps.

RMCAUSSC

Unilateral swap out count.

RMCASSC

Auxiliary storage shortage swap count.

RMCTTRPC

Number of pages used for transaction elements.

SMCABFLS

Number of records lost because of a shortage of buffers.

SMCABFWT

Number of buffers written.

SMCADSCT

Number of records lost because of a full data set.

SMCANMFL

Current number of full buffers.

SMCARCWT

Number of records written.

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

MIN

Minimum sampled value of the variable over the sampling period.

MAX

Maximum sampled value of the variable over the sampling period.

AVG

Average value of the variable over the sampling period.

END

Snapshot of the last value in the sampling period.

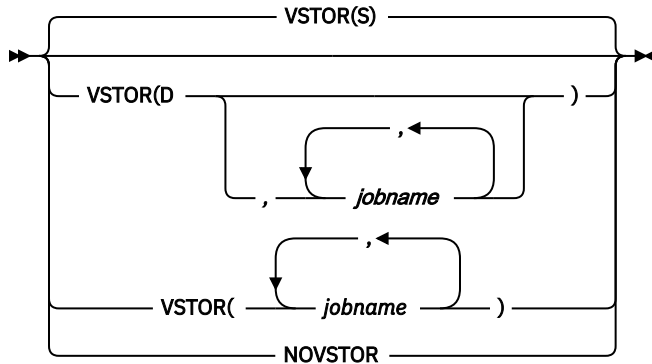
STDDEV

Standard deviation from the values sampled.

ALL

Default for *opt_list*, meaning all of the above.

VSTOR



Specifies whether virtual storage activity is to be measured. When you specify S, either explicitly or by default, summary data is collected in SMF records; when you specify D, both summary and detail data is collected in SMF records.

To collect private area data, replace *jobname* with the name of the job whose data is to be collected. z/OS Data Gatherer 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 z/OS Data Gatherer can gather meaningful data only for long-running jobs.

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

If you omit the VSTOR option, the default is VSTOR(S). Specifying VSTOR without any operands also collects summary data for common storage. The following examples show some other possible combinations.


Examples:

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

If you specify the name of a job that is not running when Monitor I begins measuring virtual storage activity, Monitor I issues a message indicating that it cannot gather data about the named job. For as long as the VSTOR option remains unchanged, Monitor I 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 Monitor I finds the job, it deletes the message from the operator console.


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

WKLD



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

VMGUEST



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

Special considerations

Specify Monitor I session options carefully. If z/OS Data Gatherer 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 for paging, page data set, processor, trace, virtual storage, CPU, channel, I/O queuing, and device activity measurements is statistically sampled. As the accuracy of sampled data increases with the number of random samples taken, you would expect to observe more precise results with decreased CYCLE time (for a fixed INTERVAL value), or with increased INTERVAL length (for a fixed CYCLE value). For example, 400 samples taken of random independent events provide a value that, with 90% confidence, should fall within 4% of the true value; 1,600 samples of random independent events decrease to 2% the expected range of error, with 90% confidence.

However, pure statistical predictions are not always applicable to a software measurement tool such as z/OS Data Gatherer 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 z/OS Data Gatherer samples internal indications of external system events. Thus, z/OS Data Gatherer 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.

Device class selection for the **DEVICE** option

Because Monitor I overhead is directly related to the number of devices being measured, the **DEVICE** option list should include only those devices that require measurement. To further reduce Monitor I overhead, select specific devices for data gathering rather than entire device classes. 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 8. Snapshot data gathering with Monitor II

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

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

All valid options are similar to those you can use during a Monitor II display session, so they are described in [Snapshot reporting with Monitor II](#) in *z/OS Resource Measurement Facility User's Guide*.

Chapter 9. Short-term data gathering with Monitor III

The topics that follow describe:

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

The detailed descriptions of the options are in alphabetical order.

Summary of gatherer session options

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

Before the session, use the following:

- The Monitor III gatherer session parmlib member. The default member is ERBRMF04. See [“Storing gatherer options” on page 24](#) for its contents. For a description of the MEMBER option and how to specify other parmlib members, see [“Descriptions of Monitor III data gatherer options” on page 85](#).
- The **parm** field of the START session command that starts the session. See [“Starting a specific monitor session” on page 45](#).

During the session, use the following:

- The PARM field of the MODIFY session command, to modify options already in effect. See [“Modifying data gatherer session options” on page 47](#).
- The response to the OPTIONS option.

Table 8 on page 83 gives a summary of the Monitor III gatherer session options. The referenced pages describe the options in detail.

Table 8. Monitor III data gatherer session options		
Option	Effect	Details on
CACHE(suboption...)	Defines cache data gathering.	“CACHE” on page 85
CFDETAIL	Defines level of detail for data gathering for the coupling facility.	“CFDETAIL” on page 85
CRYPTO	Controls data gathering for cryptographic hardware activity reports.	“CRYPTO” on page 86
CYCLE(nnnn)	Sets the length of the cycle at the end of which Monitor III samples data.	“CYCLE” on page 86
DATASET(suboption...)	Controls data set recording of sampled data.	“DATASET” on page 87
EADM	Controls data gathering for extended asynchronous data mover (EADM) activity report.	“EADM” on page 87
HFSNAME(suboption...)	Controls data set recording for z/OS UNIX file systems.	“HFSNAME” on page 88
IOSUB	Controls data set recording of I/O-subsystem and channel-path activity.	“IOSUB” on page 88
LOCK	Defines data gathering for lock reporting (spin locks and suspend locks).	“LOCK” on page 88
MASTER	Makes the system eligible/ineligible to be the master gatherer system.	“MASTER” on page 88
MEMBER(list)	Specifies parmlib members containing session options.	“MEMBER” on page 89
MINTIME(nnn)	Specifies the interval at which data samples are summarized.	“MINTIME” on page 90
OPD	Defines data gathering for OMVS process data.	“OPD” on page 90
OPTIONS	Controls display of the current options at the start of a session.	“OPTIONS” on page 90

Table 8. Monitor III data gatherer session options (continued)		
Option	Effect	Details on
PCIE	Controls data gathering for PCIE activity report.	“PCIE” on page 91
RESOURCE(...)	Specifies the job entry subsystem (JES) to be used.	“RESOURCE” on page 92
SGSPACE(suboption...)	Defines data gathering for storage group space and disk space monitoring.	“SGSPACE” on page 92
SCM	Deprecated. Use the EADM option instead.	“EADM” on page 87
STOP(value)	Sets the duration of the data gatherer interval.	“STOP” on page 93
SYNC	Synchronizes MINTIME within the hour.	“SYNC” on page 93
SYSOUT(class)	Specifies the SYSOUT class for gatherer messages.	“SYSOUT” on page 93
VSAMRLS(suboption...)	Controls data gathering for VSAM RLS activity.	“SYSOUT” on page 93
WSTOR	Sets the size of the Monitor III local storage buffer.	“WSTOR” on page 94
ZFS	Defines data gathering for monitoring zFS activity.	“ZFS” on page 95
ZIIPUSE	Specifies whether the Monitor III data gatherer is entitled to execute partially on IBM Z Integrated Information Processors (zIIPs).	“ZIIPUSE” on page 95

Default gatherer session options

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

Table 9. Monitor III default session options	
Default Option	Description
CACHE	Defines cache data gathering.
CFDETAIL	Defines partial data gathering for the coupling facility.
CRYPTO	Activity data is gathered for cryptographic hardware.
CYCLE(1000)	Takes data samples once a second (1000 milliseconds).
DATASET(STOP,NOSWITCH)	No data set recording will be done.
EADM	Activity data is gathered for extended asynchronous data mover (EADM).
IOSUB	Defines data gathering for the I/O subsystem and for channels.
MASTER	Makes the system eligible to be the master gatherer system.
MINTIME(100)	Builds a set of samples every 100 seconds.
NOLOCK	No data gathering for lock reporting (spin locks and suspend locks).
NOOPTIONS	Session options are not displayed at the operator console at the start of the session.
NOSGSPACE	No data gathering for storage group and disk space monitoring.
NOSTOP	The session does not stop automatically after a predefined time; you must use a STOP command.
OPD	Defines data gathering for OMVS process data.
PCIE	Activity data is gathered for PCI Express based functions.
RESOURCE(*JES2,JES2)	Assumes that JES2 is installed on the system.
SYNC(0M)	MINTIME is synchronized on the hour.
VSAMRLS	Activity data is gathered for VSAM RLS by storage class.
WSTOR(128)	Sets the Monitor III local storage buffer to 128 megabytes.
ZFS	Activity data is gathered about zFS.

Table 9. Monitor III default session options (continued)	
Default Option	Description
ZIIPUSE	Specifies whether the Monitor III data gatherer is entitled to execute partially on IBM Z Integrated Information Processors (zIIPs).

Monitor III creates two types of records:

Set of samples

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

SMF records


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

You can find detailed information about all record types in [“Monitored activities”](#) on page 5.

Descriptions of Monitor III data gatherer options

The following topics describe the Monitor III data gatherer options and the syntax for specifying them.

CACHE




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

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

Note: To avoid unnecessary high CPU utilization and duplicated data, you should gather cache activity data on one system only. See "Example that shows how to set up gathering options" in [“Generalizing parmlib members”](#) on page 25.

To suppress the gathering of cache data, specify NOCACHE.

CFDETAIL




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

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


With CFDETAIL, additional data is being gathered that enables you to get many details about the usage of each structure in the coupling facility. Consider that this data gathering is done only on the master gatherer system (see also [“MASTER”](#) on page 88).

CRYPTO



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

CYCLE




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

RMF sysplex reporting: Use the same CYCLE value for all systems in the sysplex to enable correct reporting in the RMF SYSPLEX report.


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:

- Increases the amount of processor time needed to sample data
- Causes Monitor III to fill the wrap-around storage buffer more quickly
- Uses more space in the user-defined VSAM data set

DATASET



Suboption




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

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

For detailed information on the DATASET option and its suboptions, see [“Controlling data set recording”](#) on page 96.

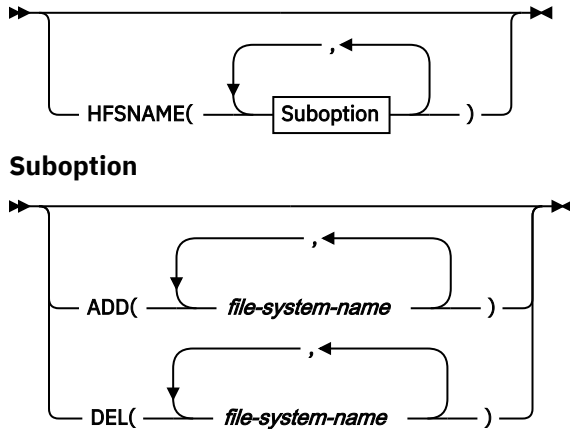
EADM



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

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

HFSNAME




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 portion of the RMF HFS Postprocessor report.


IOSUB



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


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

LOCK



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

MASTER



One member of a sysplex is selected by z/OS Data Gatherer to gather data that needs to be collected only once per sysplex. This is called *sysplex master gathering* and has been implemented to reduce workload

on non-master members and to reduce the amount of Monitor III and SMF record data. The *master gatherer* system is automatically determined by the Sysplex Data Server according to the following rules:

1. Monitor III gatherer is active.
2. System is running the highest z/OS release of the eligible systems in the sysplex.
3. Sysplex Data Server is running with SMF buffer (SMFBUF option).
4. MASTER option is specified.

You can use the MASTER parmlib option to refine the determination of which of the systems in a sysplex becomes the master gatherer system. If the MASTER option is specified for a certain system, this system is one candidate for the MASTER status. If NOMASTER is set, it will not be a MASTER candidate if there are other eligible systems having the MASTER option set. Thus you can use a combination of MASTER/NOMASTER options to select the *master gatherer* system, if there are several systems that fulfill priority rules “1” on page 89 through “3” on page 89 simultaneously.

Beyond using the MASTER option, the *master gatherer* status of a system can be changed dynamically by using the MODIFY commands, as described in “Modifying data gatherer session options” on page 47.

If optimized coupling facility hardware data collection is active, the CFHELP suboption can be used to specify a list of coupling facilities for a system that can act as helper systems to collect the coupling facility hardware statistics. The *cf-list* value is a comma-separated list of up to 8 coupling facility names. If more than 8 coupling facility names are specified, invalid option error message ERB225I is displayed. Message ERB225I lists the name of the coupling facility that was not accepted.


The helper system collects the coupling facility hardware statistics for the specified coupling facility only, if the system is running the highest z/OS release in the sysplex. As soon as there is a helper system collecting the coupling facility hardware statistics for a specific coupling facility, the coupling facility hardware statistics for that coupling facility will no longer be collected by the *master gatherer* system. If the CFHELP suboption is specified on the master system, the master system collects coupling facility hardware statistics for the specified coupling facility regardless of whether there is a helper system in the sysplex.

The CFHELP suboption allows the use of optimized coupling facility hardware data collection in the following sysplex environments:

- Sysplex environments with systems and coupling facilities in multiple sites. In this environment, there can be a long distance between a system and coupling facility if coupling facility hardware data is only collected on one sysplex system.
- Sysplex environments with an asymmetric sysplex configuration where the *master gatherer* system is not connected to all coupling facilities.

Note: Modifications of the CFHELP suboption are always cumulative. For instance, if the current status is MASTER(CFHELP(CF01)) and you specify MASTER(CFHELP(CF02)) via the MODIFY command, the result is MASTER(CFHELP(CF01,CF02)). Resetting a CFHELP parameter is only possible by specifying the NOMASTER option followed by another MASTER(CFHELP(*cf-list*)) specification.

MEMBER




Specifies 1 to 5 parmlib members that contain Monitor III gatherer options for the session. Each member is represented by a two-character alphanumeric value, to which z/OS Data Gatherer adds to the prefix ERBRMF to form the member name. The values in the xx list must be separated by commas.

For the Monitor III gatherer session, the default is 04, indicating the ERBRMF04 parmlib member. If you have created your own parmlib, make sure you specify it on the IEFRDER DD statement in the RMF cataloged procedure. See [“Setting up the RMF control session and monitoring sessions”](#) on page 19.

If you specify an option in more than one member, z/OS Data Gatherer uses the value specified in the leftmost member of the list.

MINTIME




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

RMF sysplex reporting: Use the same MINTIME value for all systems in the sysplex to enable correct reporting in the RMF SYSPLEX report.

Valid MINTIME values are in the range 10 - 999. The default is 100. If you specify a value outside the valid range, Monitor III 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 22 for more information about using MINTIME values to synchronize Monitor I and Monitor III recording intervals.

OPD



Specifies measurements for OMVS process data.

OPTIONS



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

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

Figure 7 on page 91 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 96.)

```

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


```

Figure 7. Console output with *OPTIONS* in effect

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


Table 10. Monitor III <i>OPTIONS</i> command sources	
Source keyword	Description of 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	z/OS Data Gatherer changed a conflicting option. A message describes the conflict and the change that was made.
MEMBER	The Monitor III parmlib member.
REPLY	From the operator console in reply to message ERB306I.

PCIE



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

RESOURCE



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

*JES2

Required if the installed primary JES is JES2.

*JES3

Required if the installed primary JES is JES3.

parm


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

The default is RESOURCE(*JES2,JES2).


SCM

Deprecated. Use the EADM option instead. See [“EADM ” on page 87](#).

SGSPACE



Suboption




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

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

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

STOP




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

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

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


SYNC



Specifies how the MINTIME interval is to be synchronized with the hour. This option must be specified if you want to generate RMF sysplex reports. See [“Synchronizing SMF recording intervals” on page 22](#) for more information. If you want synchronization, specify SYNC and the number of minutes (*mm*) after the hour (in the range 0 - 59) at which you want synchronization. If you specify a value that is not in the 0 - 59 range, Monitor III 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. Monitor III synchronizes the starting time of a set of samples by calculating how many sets of samples will fit in the time range up to the first synchronization point. This might mean that the MINTIME interval before the synchronization point is shortened. Subsequent sets of samples remain synchronized only when the MINTIME value is a factor of 60.


SYSOUT




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

The default value is A.

VSAMRLS



Suboption



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

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

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


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

A data set mask must apply to following rules:

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

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

WSTOR



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


The valid range for *nnn* is 4 - 999. If you do not specify a value, Monitor III uses a default of 128. If you specify a value outside the valid range, Monitor III uses 999 megabytes for a value above the range and 4 megabytes for a value below the range.

The Monitor III data gatherer stores the set of samples collected during a MINTIME in its own local storage buffer. If you specify data set recording during a session, the data gatherer 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 job name or device name) are held in tables to reduce the amount of local storage needed.

Note:


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

ZFS



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


ZIIPUSE




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

Controlling data set recording

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



Suboption



Specify at least one of the following suboptions:

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

ADD(*data-set-name*[,*data-set-name*])

DEL(*data-set-name*[,*data-set-name*])

Allows you to specify the name of the VSAM data set on which you want the data gatherer 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, the data gatherer issues a message.

ADD(*data-set-name*)

Allows the data gatherer 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 that the data gatherer 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, the data gatherer 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))
```

The data gatherer uses the empty data sets in the order in which they are defined. During data set recording, the data gatherer writes the samples from its local storage buffer to the data sets. When all the data sets are full, the data gatherer 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 the data gatherer does not reuse it, use the suboption DEL. This prevents the data gatherer from writing over data in the specified data set. To save data contained in the RMF.DS01 data set specified in the previous example, specify:

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

The data gatherer 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 or STOP at the beginning of a session on the START session command, or while the data gatherer is active with the MODIFY session command. If you do not want data set support for the data gatherer, use the default, which is DATASET(STOP).

The data gatherer handles the START and STOP suboptions only at the end of a MINTIME. At this point, the data gatherer 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 and STOP suboptions, the data gatherer avoids recording partial sets of samples in the data sets.

SWITCH NOSWITCH

Controls the selection of a data set for recording sampled data.

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

1. The data gatherer searches for an empty data set to record samples.
2. If there are no empty data sets, the data gatherer 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, the data gatherer chooses the active data set as follows:

1. The data gatherer 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, the data gatherer searches for an empty data set to record samples.
3. If there are no empty data sets, the data gatherer 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 data gatherer's local storage buffer. The valid value range 1 - 999. If you do not specify a value, the data gatherer uses

a default of 128. If you specify a value outside the valid range, the data gatherer uses 999 megabytes for a value above the range and 1 megabyte for a value below the range.

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

WHOLD works with the WSTOR option (see “WSTOR” on page 94) 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 the data gatherer 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 the data gatherer 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 6 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 the data gatherer through the DATASET option. The data set names must be identical to the names used to define the data sets, otherwise the data gatherer will not recognize them.

Because MEMBER(08) is specified in the START command, the data gatherer 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 the data gatherer to continue writing on the active data set of the previous session (in this case, RMF.DS05).

Assume the following facts 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 the data gatherer 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 the data gatherer can use all the data sets listed with the ADD suboption.

As a result, z/OS Data Gatherer produces the following list of options following the rules of processing session options:

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

For more information, see [Chapter 6, “How session options are processed,”](#) on page 51.

The RMF.DS06 data set is now available for data set recording. The RMF.DS05 data set cannot be used for recording during the session. The RMF.DS05 data set can be pre-allocated at the beginning of a TSO Monitor III reporter session, and the data on it can be displayed and analyzed. For more information, see [Transferring Monitor III VSAM data sets to other systems in z/OS Resource Measurement Facility User's Guide](#).

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

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

Modifying the data set support options

You can also modify the DATASET options while the data gatherer is active through the MODIFY session command. For more information, see [“Modifying data gatherer session options”](#) on page 47.

Example:

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

You want to:

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

The following command modifies the options:

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

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

Stopping data set support

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

- Specify the DATASET(STOP) option in the PARM field of the START session command.
- Specify the DATASET(STOP) option in the PARM field of the MODIFY session command.
- Specify the DATASET(STOP) option in a z/OS Data Gatherer 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 DATASET(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 DATASET(STOP) option in the ERBRMF04 member, and permits the recording of sampled data to the data sets defined by the DATASET(ADD) options.

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

Data set support for daylight saving time

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

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

There is a gap in local time where no data is selected. When the currently active data set is full, the data set with the oldest data is selected to store the current data.

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

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

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

Appendix A. Accessibility

Accessible publications for this product are offered through [IBM Documentation for z/OS \(www.ibm.com/docs/en/zos\)](http://www.ibm.com/docs/en/zos).

If you experience difficulty with the accessibility of any z/OS documentation see [How to Send Feedback to IBM](#) to leave documentation feedback.

Notices

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

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

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

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

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

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

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

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

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

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

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

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

*IBM Corporation
Site Counsel
2455 South Road*

Poughkeepsie, NY 12601-5400
USA

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Terms and conditions for product documentation

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

Applicability

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

Personal use

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

Commercial use

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

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

Rights

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

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

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

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

IBM Online Privacy Statement

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

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

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

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at ibm.com/privacy and IBM's Online Privacy Statement at ibm.com/privacy/details in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at ibm.com/software/info/product-privacy.

Policy for unsupported hardware

Various z/OS elements, such as DFSMSdfp, JES2, and MVS, contain code that supports specific hardware servers or devices. In some cases, this device-related element support remains in the product even after the hardware devices pass their announced End of Service date. z/OS may continue to service element code; however, it will not provide service related to unsupported hardware devices. Software problems related to these devices will not be accepted for service, and current service activity will cease if a problem is determined to be associated with out-of-support devices. In such cases, fixes will not be issued.

Minimum supported hardware

The minimum supported hardware for z/OS releases identified in z/OS announcements can subsequently change when service for particular servers or devices is withdrawn. Likewise, the levels of other software products supported on a particular release of z/OS are subject to the service support lifecycle of those

products. Therefore, z/OS and its product publications (for example, panels, samples, messages, and product documentation) can include references to hardware and software that is no longer supported.

- For information about software support lifecycle, see: [IBM Lifecycle Support for z/OS \(www.ibm.com/software/support/systemsz/lifecycle\)](http://www.ibm.com/software/support/systemsz/lifecycle)
- For information about currently-supported IBM hardware, contact your IBM representative.

Programming interface information

This information is intended to help you use z/OS Data Gatherer sessions. It contains a description of what z/OS Data Gatherer is, what it can do, and how to use the different sessions.

The information also describes intended programming interfaces that allow you to write programs to obtain the services of z/OS Data Gatherer.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at [Copyright and Trademark information \(www.ibm.com/legal/copytrade.shtml\)](http://www.ibm.com/legal/copytrade.shtml).

Java is a trademark of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Windows, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.

Other company, product, and service names may be trademarks or service marks of others.

Glossary

This glossary contains chiefly definitions of terms used in this book, but some more general z/OS Data Gatherer, RMF, and z/OS terms are also defined.

Words that are set in *italics* in the definitions are terms that are themselves defined in the glossary.

APPC/MVS

Advanced program-to-program communication

ASCH address space

APPC transaction scheduler address space

AS

Address space

address space

That part of z/OS main storage that is allocated to a job.

auxiliary storage (AUX)

All addressable storage, other than main storage, that can be accessed by means of an I/O channel; for example storage on direct access devices.

background session

A z/OS Data Gatherer monitor session that is started and controlled from the operator console. Contrast with *interactive session*

balanced systems

To avoid bottlenecks, the system resources (CP, I/O, storage) need to be balanced.

basic mode

A central processor mode that does not use logical partitioning. Contrast with *logically partitioned (LPAR) mode*.

bottleneck

A system resource that is unable to process work at the rate it comes in, thus creating a queue.

callable services

Parts of a program product that have a published external interface and can be used by application programs to interact with the product.

captured storage

See shared page group.

capture ratio

The ratio of reported CPU time to total used CPU time.

central processor (CP)

The part of the computer that contains the sequencing and processing facilities for instruction execution, initial program load, and other machine operations.

central processor complex (CPC)

A physical collection of hardware that consists of central storage, one or more central processors, timers, and channels.

channel path

The channel path is the physical interface that connects control units and devices to the CPU.

CICS

Customer Information Control System

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 z/OS 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 z/OS Data Gatherer, the time at the end of which one sample is taken. Varies between 50 ms and 9999 ms. See also *sample*.

data sample

See *sample*

DCM

See *Dynamic Channel Path Management*

delay

The delay of an address space represents a job that needs one or more resources but that must wait because it is contending for the resource(s) with other users in the system.

direct access storage device (DASD)

A device in which the access time is effectively independent of the location of the data. Usually: a magnetic disk device.

DLY

Delay

DP

Dispatching priority

dynamic channel path management

Dynamic channel path management provides the capability to dynamically assign channels to control units in order to respond to peaks in demand for I/O channel bandwidth. This is possible by allowing you to define pools of so-called floating channels that are not related to a specific control unit. With the help of the Workload Manager, channels can float between control units to best service the work according to their goals and their importance.

EMIF

ESCON multiple image facility

enclave

An enclave is a group of associated dispatchable units. More specifically, an enclave is a group of SRB routines that are to be managed and reported on as an entity.

EPDM

Enterprise Performance Data Manager/MVS

execution velocity

A measure of how fast work should run when ready, without being delayed for processor or storage access.

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.

GTF

generalized trace facility

high-speed buffer (HSB)

A cache or a set of logically partitioned blocks that provides significantly faster access to instructions and data than provided by central storage.

HS

hiperspace

HSB

High-speed buffer

HSM

Hierarchical Storage Manager

IBM Z® Application Assist Processor (zAAP)

A special purpose processor configured for running Java programming on selected zSeries machines.

IBM Z Integrated Information Processor (zIIP)

A special purpose processor designed to help free-up general computing capacity and lower overall total cost of computing for selected data and transaction processing workloads for business intelligence (BI), ERP and CRM, and selected network encryption workloads on the mainframe.

IMS

Information Management System

Information Management System (IMS)

A database/data communication (DB/DC) system that can manage complex databases and networks. Synonymous with IMS/VS.

interactive session

In RMF, a monitor display-session that is controlled from the display terminal. Contrast with *background session*.

JES

Job Entry Subsystem

LCU

Logical control unit. Logical control units are also called 'Control Unit Headers ' (CUH). For details about LCU/CUH please refer to the applicable *IBM Z Input/Output Configuration Program User's Guide for ICP IOCP* (SB10-7037).

logically partitioned (LPAR) mode

A central processor mode that is available on the Configuration frame when using the PR/SM feature. It allows an operator to allocate processor unit hardware resources among logical partitions. Contrast with *basic mode*.

logical partition (LP)

A subset of the processor hardware that is defined to support an operating system. See also *logically partitioned (LPAR) mode*.

LP

Logical partition

LPAR

Logically partitioned (mode)

LPAR cluster

An LPAR cluster is the subset of the systems that are running as LPARs on the same CEC. Based on business goals, WLM can direct PR/SM to enable or disable CP capacity for an LPAR, without human intervention.

migration rate

The rate (pages/second) of pages being moved from expanded storage through central storage to auxiliary storage.

mintime

The smallest unit of sampling in Monitor III. Specifies a time interval during which the system is sampled. The data gatherer combines all samples gathered into a set of samples. The set of samples can be summarized and reported by the reporter.

MPL

Multiprogramming level

OMVS

Reference to z/OS UNIX System Services

partitioned data set (PDS)

A data set in direct access storage that is divided into partitions, called members, each of which can contain a program, part of a program, or data.

PDS

partitioned data set

performance management

The activity which monitors and allocates data processing resources to applications according to goals defined in a service level agreement or other objectives.

The discipline that encompasses collection of performance data and tuning of resources.

PR/SM

Processor Resource/Systems Manager

Processor Resource/Systems Manager (PR/SM)

The feature that allows the processor to run several operating systems environments simultaneously and provides logical partitioning capability. See also *LPAR*.

range

The time interval you choose for your report.

Resident time

The time the address space was swapped in, in units of seconds.

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 z/OS 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 z/OS, z/VM®, and z/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 z/OS systems.

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 z/OS, a time-sharing system accessed from a terminal that allows user access to z/OS system services and interactive facilities.

UIC

Unreferenced interval count

uncaptured time

CPU time not allocated to a specific address space.

using

Jobs getting service from hardware resources (PROC or DEV) are *using* these resources.

velocity

A measure of how fast work should run when ready, without being delayed for processor or storage access. See also *execution velocity*.

VTOC

Volume table of contents

workload

A logical group of work to be tracked, managed, and reported as a unit. Also, a logical group of service classes.

WLM

Workload Manager

XCF

Cross-system coupling facility

XES/CF

See *Cross-system Extended Services/Coupling Facility*.

zAAP

see IBM Z Application Assist Processor.

zIIP

see IBM Z Integrated Information Processor.

Index

A

- accessibility
 - contact IBM [101](#)
- ACTIVE (active non-interactive sessions) parameter
 - DISPLAY command [48](#)
- activities, monitored [5](#)
- ADD|DEL suboption
 - Monitor III data gatherer session [96](#)
 - Monitor III DATASET option [96](#)
- administration, data gatherer [13](#)
- ALL (all non-interactive sessions) parameter
 - DISPLAY command [48](#)
- ASMERRS (bad slots on local page data set)
 - Trace Activity report variable [72](#)
- ASMIORQC (I/O requests completed and returned)
 - Trace Activity report variable [72](#)
- ASMIORQR (I/O request received)
 - Trace Activity report variable [72](#)
- ASMNVSC (total local slots for non-VIO private area pages)
 - Trace Activity report variable [72](#)
- ASMSLOTS (total slots in open local page data sets)
 - Trace Activity report variable [72](#)
- ASMVSC (total local slots allocated for VIO private area pages)
 - Trace Activity report variable [72](#)
- assistive technologies [101](#)

B

- background session
 - starting Monitor II [53](#)
- BPX.WLMSEVER [17](#)

C

- CACHE
 - Monitor III data gatherer session
 - description [85](#)
- cache activity
 - Monitor I session (CACHE/NOCACHE)
 - summary [59](#)
- CACHE/NOCACHE (cache activity)
 - Monitor I session
 - summary [59](#)
- CANCEL command
 - system command [41](#)
- cataloged procedure
 - Monitor III data gatherer session
 - RMFGAT [47](#)
- CCVCPUCT (number of online CPUs)
 - Trace Activity report variable [72](#)
- CCVEJST (low priority user job step time)
 - Trace Activity report variable [72](#)

- CCVENQCT (number of users non-swappable for enqueue reasons)
 - Trace Activity report variable [72](#)
- CCVRBSTD (recent base time of day)
 - Trace Activity report variable [72](#)
- CCVRBSWT (recent base system wait time)
 - Trace Activity report variable [72](#)
- CCVUTILP (system CPU utilization)
 - Trace Activity report variable [72](#)
- central processor unit (CPU) activity
 - Monitor I session (CPU/NOCPU)
 - summary [59](#)
- CFDETAIL (coupling facility)
 - Monitor III data gatherer session
 - description [85](#)
 - effect on system performance [85](#)
- CFOPT option [43](#)
- CHAN/NOCHAN (channel path activity)
 - Monitor I session
 - summary [59](#)
- channel path activity
 - Monitor I session (CHAN/NOCHAN)
 - summary [59](#)
- command
 - syntax diagrams [xii](#)
- common storage
 - detail report
 - Monitor I VSTOR/NOVSTOR option [77](#)
 - summary report
 - Monitor I VSTOR/NOVSTOR option [77](#)
- console output from OPTIONS
 - Monitor III data gatherer session [90](#)
- contact
 - z/OS [101](#)
- coupling facility hardware data collection [43](#)
- CPU activity
 - Monitor I session (CPU/NOCPU) [62](#)
- CPU/NOCPU (central processor unit activity)
 - Monitor I session
 - summary [59](#)
- CRYPTO
 - Monitor I session [86](#)
- CRYPTO/NOCRYPTO (cryptographic hardware activity)
 - Monitor I session
 - summary [59](#)
- cryptographic hardware activity
 - Monitor I session (CRYPTO/NOCRYPTO) [62](#)
- customizing
 - system environment [15](#)
- CYCLE (cycle length)
 - Monitor I session
 - overhead reduction [79](#)
 - summary [60](#)
 - undesirable results [78](#)
 - Monitor III data gatherer session
 - description [86](#)
 - effect on system performance [86](#)

D

- data collection
 - skipped interval
 - Monitor I session [66](#)
 - statistical sampling
 - Monitor I session [78](#)
- data collection, optimized coupling facility hardware [43](#)
- data gatherer
 - Monitor III session
 - modifying [55](#)
 - session identifier [47](#)
- data gatherer administration [13](#)
- data gatherer sessions
 - displaying status of [48](#)
- data gatherer started tasks
 - specifying priority for [20](#)
- data gatherer)
 - non-interactive session
 - starting [45](#)
- data gathering [4](#)
- data gathering capabilities [57](#)
- data gathering for SMF record producer instances with Monitor P [5](#)
- data set support
 - daylight saving time [100](#)
- DATASET (record samples)
 - ADD|DEL suboption
 - Monitor III data gatherer session [96](#)
 - Monitor III data gatherer session
 - description [87](#)
 - START|STOP suboption
 - Monitor III data gatherer session [97](#)
 - SWITCH|NOSWITCH suboption
 - Monitor III data gatherer session [97](#)
 - WHOLD suboption
 - Monitor III data gatherer session [97](#)
- daylight saving time [100](#)
- default session options [60](#)
- defining SMF synchronization [23](#)
- device number (NMBR/NONMBR) sub-option
 - Monitor I session device activity option [63](#)
- DEVICE/NODEVICE (I/O device activity)
 - classes sub-option [63](#)
 - Monitor I session
 - overhead reduction [79](#)
 - summary [59](#)
 - NMBR/NONMBR sub-option [63](#)
 - SG/NOSG sub-option [63](#)
- dispatching priority
 - Monitor I session [66](#)
- DISPLAY command
 - session command [48](#)
- displaying
 - options in effect for a session [52](#)

E

- EADM (Extended Asynchronous Data Mover)
 - Monitor III data gatherer session
 - description [87](#)
- Emitter, z/OS OpenTelemetry [10](#), [35](#)
- ENQ/NOENQ (enqueue activity)

ENQ/NOENQ (enqueue activity) (*continued*)

- Monitor I session
 - description [64](#)
 - summary [59](#)
- enqueue activity
 - Monitor I session (ENQ/NOENQ)
 - description [64](#)
 - summary [59](#)
- ERBRMF00 (parmlib member)
 - Monitor I session
 - contents [26](#)
- ERBRMF01 (parmlib member)
 - Monitor II session
 - contents [28](#)
- ERBRMF02 (parmlib member)
 - Monitor I session
 - contents [27](#)
- ERBRMF03 (parmlib member)
 - Monitor II session
 - contents [29](#)
- ERBRMF04 (parmlib member)
 - Monitor III data gatherer session
 - content [89](#)
 - session option specification [83](#)
 - Monitor III session
 - contents [30](#)
- ERBRMF05 (parmlib member)
 - Monitor P session
 - contents [30](#)
- ERBVSDEF CLIST [21](#)
- ESS activity
 - Monitor I session (ESS/NOESS)
 - description [65](#)
- ESS/NOESS (Enterprise Storage Server)
 - Monitor I session
 - summary [59](#)
- ESS/NOESS (ESS activity)
 - Monitor I session
 - description [65](#)
- EXITS/NOEXITS (user exit)
 - Monitor I session
 - description [65](#)
 - summary [60](#)

F

- FACILITY class
 - RACF resource profiles [17](#)
- FCD/NOFCD (FICON director activity)
 - Monitor I session
 - description [66](#)
 - summary [59](#)
- FICON director activity
 - Monitor I session (FCD/NOFCD)
 - description [66](#)

G

- gatherer options
 - storing [24](#)
- global performance data control authority [16](#)
- glossary [107](#)
- GRBMFRES

GRBMFRES (*continued*)
removing [16](#)

H

HFSNAME (UNIX file system statistics)
Monitor III data gatherer session
description [88](#)

I

I/O device activity
Monitor I session (DEVICE/NODEVICE)
classes sub-option [63](#)
NMBR/NONMBR sub-option [63](#)
SG/NOSG sub-option [63](#)
summary [59](#)
Monitor I session (I/O DEVICE/
NODEVICE)
overhead reduction [79](#)
I/O queuing activity
Monitor I session
description [67](#)
Monitor I session (IOQ/NOIOQ)
classes sub-option [67](#)
NMBR/NONMBR sub-option [67](#)
summary [60](#)
IBM Z Application Assist Processor (zAAP)
definition [109](#)
IBM Z Integrated Information Processor (zIIP)
definition [109](#)
IEFPARM DD statement
definition [20](#)
IEFPROC EXEC statement
definition [19](#)
INTERVAL (interval length)
Monitor I session
description [66](#)
overhead reduction [79](#)
summary [60](#)
undesirable results [78](#)
Monitor II background session
conflicting options [45](#)
interval report
REPORT/NOREPORT option
Monitor II background session [45](#)
interval synchronization
Monitor I session (SYNC/NOSYNC)
description [71](#)
summary [60](#)
Monitor III data gatherer session option [93](#)
IOQ/NOIOQ (I/O queuing activity)
Monitor I session
classes sub-option [67](#)
description [67](#)
NMBR/NONMBR sub-option [67](#)
summary [60](#)
IOSUB (I/O subsystem)
Monitor III data gatherer session
description [88](#)
IRLM long lock reporting [43](#)

J

JCL (job control language)
Monitor III data gatherer session [47](#)
RMF cataloged procedure [19](#)
JES
starting z/OS Data Gatherer without
[41](#)
JES not active [39](#), [46](#)

K

keyboard
navigation [101](#)
PF keys [101](#)
shortcut keys [101](#)

L

LOCK (lock reporting)
Monitor III data gatherer session
description [88](#)
long-term data gathering
with Monitor I [4](#)
with Monitor III [4](#)
LSCTCNT (logically swapped users for terminal wait)
Trace Activity report variable [72](#)
LSTCTMTE (think time for logical swap candidate)
Trace Activity report variable [72](#)

M

major name
Monitor I session
ENQ/NOENQ option [64](#)
MASTER
Monitor III data gatherer session [88](#)
master gatherer system [88](#), [89](#)
MCVFRcnt (pages to be stolen by force steal routine)
Trace Activity report variable [72](#)
MCVMGAGE (expanded storage migration age)
Trace Activity report variable [72](#)
MCVSBTLF (percentage of eligible storage that is actually
fixed)
Trace Activity report variable [73](#)
MCVSIPR (common page-in rate)
Trace Activity report variable [73](#)
MCVSTCRI (highest system UIC)
Trace Activity report variable [73](#)
MCVTWSS (common target working set size)
Trace Activity report variable [73](#)
MEMBER (parmlib member)
Monitor I session
description [68](#)
summary [60](#)
Monitor II background session option
processing order [54](#)
Monitor III data gatherer session
description [89](#)
migrating from previous releases [15](#)
minor name
Monitor I session
ENQ/NOENQ option [64](#)

- MINTIME (interval length)
 - Monitor III data gatherer session
 - description [90](#)
- MODIFY command
 - Monitor III data gatherer session [99](#)
 - session command
 - example [48](#)
 - modifying gatherer session options [47](#)
 - syntax [47](#)
 - system command
 - non-interactive session [45](#)
 - passing DISPLAY session command [48](#)
 - passing MODIFY session commands [47](#)
 - passing STOP session command [49](#)
- Monitor I
 - long-term data gathering with [4](#)
 - session options [61](#)
- Monitor I session
 - cache activity
 - CACHE/NOCACHE option summary [59](#)
 - central processor unit activity
 - CPU/NOCPU option summary [59](#)
 - channel path activity
 - CHAN/NOCHAN option summary [59](#)
 - CPU/NOCPU option [62](#)
 - cryptographic hardware activity
 - CRYPTO/NOCRYPTO [62](#)
 - CRYPTO/NOCRYPTO option summary [59](#)
 - cycle length
 - CYCLE [62](#)
 - CYCLE option summary [60](#)
 - overhead reduction [79](#)
 - data collection
 - skipped interval [66](#)
 - statistical sampling [78](#)
 - data recording
 - RECORD/NORECORD option description [70](#)
 - RECORD/NORECORD option summary [60](#)
 - device activity
 - DEVICE/NODEVICE option description [63](#)
 - DEVICE/NODEVICE option summary [59](#)
 - overhead reduction [79](#)
 - dispatching priority [66](#)
 - enqueue activity
 - ENQ/NOENQ option description [64](#)
 - ENQ/NOENQ option summary [59](#)
 - ESS activity
 - ESS/NOESS option description [65](#)
 - ESS/NOESS option summary [59](#)
 - FICON director activity
 - FCD/NOFCD option description [66](#)
 - FCD/NOFCD option summary [59](#)
 - I/O device activity
 - DEVICE/NODEVICE option description [63](#)
 - DEVICE/NODEVICE option summary [59](#)
 - I/O queuing activity
 - classes sub-option [67](#)
 - device number suboption [67](#)
 - IOQ/NOIOQ option description [67](#)
 - IOQ/NOIOQ option summary [60](#)
 - interval length
 - INTERVAL option description [66](#)
 - INTERVAL option summary [60](#)
 - overhead reduction [79](#)
- Monitor I session (*continued*)
 - interval synchronization
 - SYNC/NOSYNC option description [71](#)
 - SYNC/NOSYNC option summary [60](#)
 - NOZZ option [39](#)
 - options
 - command source (input source) [69](#)
 - conflicting options [78](#)
 - OPTIONS/NOOPTIONS option description [69](#)
 - OPTIONS/NOOPTIONS option summary [60](#)
 - specifying [45](#)
 - undesirable results [78](#)
 - page data set activity
 - PAGESP/NOPAGESP option description [70](#)
 - PAGESP/NOPAGESP option summary [60](#)
 - parmlib member
 - MEMBER option description [68](#)
 - MEMBER option summary [60](#)
 - printed report
 - SYSOUT option description [71](#)
 - SYSOUT option summary [60](#)
 - private area data summary and detail reports [77](#)
 - report interval
 - interval length option summary [60](#)
 - session duration
 - STOP/NOSTOP option description [70](#)
 - STOP/NOSTOP option summary [60](#)
 - session identifier [46](#)
 - session option
 - CACHE/NOCACHE (cache activity) [59](#)
 - CHAN/NOCHAN (channel path activity) [59](#)
 - CPU/NOCPU (central processor unit activity) [59](#)
 - CRYPTO/NOCRYPTO (cryptographic hardware activity) [59](#)
 - CYCLE (cycle length) [60](#), [78](#), [79](#)
 - DEVICE/NODEVICE (I/O device activity) [59](#), [63](#), [79](#)
 - displaying [48](#)
 - ENQ/NOENQ (enqueue activity) [59](#), [64](#)
 - ESS/NOESS (Enterprise Storage Server) [59](#)
 - ESS/NOESS (ESS activity) [65](#)
 - EXITS/NOEXITS (user exit routine) [60](#), [65](#)
 - FCD/NOFCD (FICON director activity) [59](#), [66](#)
 - INTERVAL (interval length) [60](#), [66](#), [78](#), [79](#)
 - IOQ/NOIOQ (I/O queuing activity) [60](#), [67](#)
 - MEMBER (parmlib member) [60](#), [68](#)
 - modifying [48](#)
 - OPTIONS/NOOPTIONS (option list) [60](#), [69](#)
 - PAGESP/NOPAGESP (page data set activity) [60](#), [70](#)
 - PAGING/NOPAGING (system paging activity) [60](#), [70](#)
 - RECORD/NORECORD (record to SMF data set) [60](#), [70](#)
 - STOP/NOSTOP (session duration) [60](#), [70](#)
 - SYNC/NOSYNC (interval synchronization) [22](#), [60](#), [71](#)
 - SYSOUT (sysout class) [60](#), [71](#)
 - TRACE/NOTRACE (trace activity) [60](#), [72](#)
 - VSTOR/NOVSTOR (virtual storage activity) [60](#), [77](#)
 - WKLD/NOWKLD (workload activity) [60](#), [78](#)
 - SMF record
 - synchronization with Monitor III [22](#)
 - writing data to [60](#), [70](#)
 - starting
 - concurrent with RMF control session [39](#)
 - stopping [49](#)
 - sysout class

- Monitor I session (*continued*)
 - sysout class (*continued*)
 - SYSOUT option description [71](#)
 - SYSOUT option summary [60](#)
 - system paging activity
 - PAGING/NOPAGING option description [70](#)
 - PAGING/NOPAGING option summary [60](#)
 - trace activity
 - TRACE/NOTRACE [72](#)
 - TRACE/NOTRACE option summary [60](#)
 - user exit routine
 - EXITS/NOEXITS option description [65](#)
 - EXITS/NOEXITS option summary [60](#)
 - virtual storage activity
 - VSTOR/NOVSTOR option description [77](#)
 - VSTOR/NOVSTOR option summary [60](#)
 - VMGUEST/NOVMGUEST [78](#)
 - workload activity
 - WKLD/NOWKLD option description [78](#)
 - WKLD/NOWKLD option summary [60](#)
- Monitor I session option
 - CACHE/NOCACHE (cache activity) [61](#)
 - CHAN/NOCHAN (channel path activity) [62](#)
 - CPU/NOCPU (CPU activity) [62](#)
 - CRYPTO [86](#)
 - CRYPTO/NOCRYPTO (cryptographic hardware activity) [62](#)
 - CYCLE (cycle length) [62](#)
 - default [59](#)
 - summary [59](#)
- Monitor II
 - background session [55](#)
 - snapshot monitoring with [4](#)
- Monitor II background session
 - maximum number [45](#)
 - modifying options [48](#)
 - options
 - conflicting [45](#)
 - session identifier [46](#)
 - starting [46](#), [53](#)
 - stopping [49](#)
- Monitor II data gatherer session
 - session option [81](#)
- Monitor II REST services [8](#)
- Monitor III
 - long-term data gathering with [4](#)
 - short-term data collection with [4](#)
- Monitor III data gatherer session
 - cataloged procedure [47](#)
 - console output [48](#)
 - controlling recording [96](#)
 - description [83](#)
 - local storage buffer size [94](#)
 - modifying options [48](#)
 - option list [90](#)
 - parmlib member
 - description [83](#), [89](#)
 - member name generation [89](#)
 - option [89](#)
 - RMFGAT [20](#)
 - sample collection
 - time interval [90](#)
 - user defined data set [87](#)
 - session identifier [47](#)

- Monitor III data gatherer session (*continued*)
 - session option
 - CACHE (cache subsystem) [85](#)
 - CFDETAIL (coupling facility) [85](#)
 - CYCLE (cycle length) [86](#)
 - DATASET (record samples) [87](#), [96](#)
 - defaults [84](#)
 - EADM (Extended Asynchronous Data Mover) [87](#)
 - HFSNAME (UNIX file system statistics) [88](#)
 - IOSUB (I/O subsystem) [88](#)
 - LOCK (lock reporting) [88](#)
 - MASTER [88](#)
 - MEMBER (parmlib member) [89](#)
 - MINTIME (interval length) [90](#)
 - OPD (OMVS process data) [90](#)
 - OPTIONS (option list) [90](#)
 - PCIE (PCI Express based functions) [91](#)
 - RESOURCE (JES resource) [92](#)
 - SGSPACE [92](#)
 - specification [83](#)
 - STOP (session duration) [93](#)
 - SYNC (interval synchronization) [93](#)
 - SYSOUT (sysout class) [93](#)
 - VSAMRLS [94](#)
 - WSTOR (buffer size) [94](#)
 - ZFS [95](#)
 - ZIIPUSE [95](#)
 - session options
 - descriptions [85](#)
 - starting [45](#), [47](#)
 - stopping [49](#)
 - VSAM data set support
 - data set selection [97](#)
 - defining data sets [21](#)
 - modifying [99](#)
 - page release control [97](#)
 - starting [97](#), [98](#)
 - stopping [100](#)
 - ZFS/NOZFS [95](#)
 - ZIIPUSE/NOZIIPUSE [95](#)
- Monitor III gatherer session
 - data set support [20](#)
- Monitor III REST services [9](#)
- Monitor P
 - data gathering for SMF record producer instances with [5](#)
- Monitor P data gatherer session
 - stopping [49](#)
- Monitor P session
 - starting [39](#), [47](#)
- monitored activities [5](#)

N

- naming parmlib option members [24](#)
- navigation
 - keyboard [101](#)
- NMBR/NONMBR (device number) sub-option
 - Monitor I session device activity option [63](#), [67](#)
- non-interactive session
 - description [45](#)
 - identifier [48](#)
 - options
 - modifying [47](#), [48](#)
 - specifying [45](#)

- non-interactive session (*continued*)
 - session control
 - session option processing [51](#)
 - session identifiers
 - displaying [48](#)
 - session options
 - displaying [48](#)
 - starting [45](#), [47](#)
 - stopping [49](#)
- non-TSO/E session
 - starting
 - RMF cataloged procedure [19](#)
- NOZZ (no Monitor I session)
 - Monitor I session [39](#)

O

- OMDGAMRE (messages on the action message retention facility queue)
 - Trace Activity report variable [73](#)
- OMDGCMDI (commands issued per second)
 - Trace Activity report variable [73](#)
- OMDGOEB (operator reply entries on system reply queue)
 - Trace Activity report variable [73](#)
- OMDGWQEB (WTO queue elements on system output queue)
 - Trace Activity report variable [73](#)
- OMDGWTI (write-to-logs issued per second)
 - Trace Activity report variable [73](#)
- OMDGWTOI (total lines of messages, write-to-operators issued per second)
 - Trace Activity report variable [73](#)
- OPD (OMVS process data)
 - Monitor III data gatherer session
 - description [90](#)
- OpenTelemetry Emitter, z/OS [10](#), [35](#)
- optimized sysplex-wide coupling facility hardware data collection [43](#)
- option list
 - Monitor I session (OPTIONS/NOOPTIONS)
 - description [69](#)
 - summary [60](#)
- options
 - default session [60](#)
 - Monitor I session
 - conflicting options [78](#)
 - undesirable results [78](#)
 - Monitor III data gatherer session
 - descriptions [85](#)
 - session options
 - MODIFY command [47](#)
- OPTIONS/NOOPTIONS (option list)
 - Monitor I session
 - description [69](#)
 - summary [60](#)
 - Monitor III data gatherer session
 - description [90](#)

P

- page data set activity
 - Monitor I session (PAGESP/NOPAGESP)
 - description [70](#)

- page data set activity (*continued*)
 - Monitor I session (PAGESP/NOPAGESP) (*continued*)
 - summary [60](#)
- PAGESP/NOPAGESP (page data set activity)
 - Monitor I session
 - description [70](#)
 - summary [60](#)
- PAGING/NOPAGING (system paging activity)
 - Monitor I session
 - description [70](#)
 - summary [60](#)
- parmlib member
 - ERBRMF00 (Monitor I session) [26](#)
 - ERBRMF01 (Monitor II session) [28](#)
 - ERBRMF02 (Monitor I session) [27](#)
 - ERBRMF03 (Monitor II session) [29](#)
 - ERBRMF04 (Monitor III session) [30](#)
 - ERBRMF05 (Monitor P session) [30](#)
 - member name generating
 - Monitor I session [68](#)
 - Monitor I session (MEMBER)
 - description [68](#)
 - summary [60](#)
- parmlib naming [25](#)
- parmlib option members
 - naming [24](#)
- PCIE (PCI Express based functions)
 - Monitor III data gatherer session
 - description [91](#)
- performance data
 - accessing across the sysplex [6](#)
 - sysplex
 - accessing data across [6](#)
- performance management, z/OS [3](#)
- PR/SM
 - processor activity [62](#)
- priority for data gatherer started tasks
 - specifying [20](#)
- private area
 - detail report
 - Monitor I VSTOR/NOVSTOR option [77](#)
 - summary report
 - Monitor I VSTOR/NOVSTOR option [77](#)
- procedure
 - RMF [19](#)
- processing session options
 - Monitor I session [55](#)
 - Monitor II background session [55](#)
 - when you modify a session
 - Monitor II background session [54](#)
- program default
 - Monitor I session options [59](#)
 - Monitor II session options [81](#)
- program properties table (PPT) [16](#)
- programming interface information [106](#)

R

- RACF FACILITY
 - BPX.WLMSEVER [17](#)
- RACF resource profiles [17](#)
- RAXESCT (common storage pages on expanded storage)
 - Trace Activity report variable [73](#)
- RAXFMCT (frames allocated to common)

RAXFMCT (frames allocated to common) (*continued*)
Trace Activity report variable [73](#)

RCEAEC (total expanded storage E frames currently on ESTE queue)
Trace Activity report variable [73](#)

RCEAECLO (available expanded storage low threshold)
Trace Activity report variable [73](#)

RCEAECOK (available expanded storage satisfactory threshold)
Trace Activity report variable [73](#)

RCEAFC (total frames currently on all available frame queues)
Trace Activity report variable [73](#)

RCEAFCLO (available central storage low threshold)
Trace Activity report variable [73](#)

RCEAFCOK (available central storage low threshold)
Trace Activity report variable [73](#)

RCEBELFX (total number of fixed pages below 16 megabytes)
Trace Activity report variable [73](#)

RCECOMPI (common area pages paged-in)
Trace Activity report variable [73](#)

RCECOMPO (common area pages paged-out)
Trace Activity report variable [74](#)

RCEDFRS (number of times a deferred frame allocation has been satisfied)
Trace Activity report variable [74](#)

RCEESINU (in-use expanded storage frames)
Trace Activity report variable [74](#)

RCEESREA (non-VIO pages read from expanded storage)
Trace Activity report variable [74](#)

RCEESWRT (pages written to expanded storage frames)
Trace Activity report variable [74](#)

RCEHSPERM (total hiperspace pages migrated from auxiliary storage) expanded to
Trace Activity report variable [74](#)

RCEHSPER (total hiperspace pages read from expanded to central storage)
Trace Activity report variable [74](#)

RCEHSPEW (total hiperspace pages written from central to expanded storage)
Trace Activity report variable [74](#)

RCEHSPPI (total hiperspace pages paged in from auxiliary storage)
Trace Activity report variable [74](#)

RCEHSPPO (total hiperspace pages paged out to auxiliary storage)
Trace Activity report variable [74](#)

RCELPAPI (PLPA and PLPA directory pages paged-in)
Trace Activity report variable [74](#)

RCEMBEL (pages moved from below 16M in central storage)
Trace Activity report variable [74](#)

RCENWSF (total number secondary and non-working pages migrated)
Trace Activity report variable [74](#)

RCEPAGMV (number of times a frame was moved)
Trace Activity report variable [74](#)

RCEPOOL (frames currently available to the system)
Trace Activity report variable [74](#)

RCESPFRR (frames available by swap-out without requiring I/O)
Trace Activity report variable [74](#)

RCESWPPI (total number of pages requiring I/O to swap-in)
Trace Activity report variable [74](#)

RCESWPPO (pages requiring I/O to swap-out)
Trace Activity report variable [74](#)

RCETOTFX (total number of pages currently fixed)
Trace Activity report variable [74](#)

RCETOTPI (total pages paged-in excluding swap-in and VIO page-in)
Trace Activity report variable [74](#)

RCETOTPO (total number pages paged out)
Trace Activity report variable [74](#)

RCEVIOME (VIO pages written to expanded storage)
Trace Activity report variable [74](#)

RCEVIOMG (VIO pages migrated from expanded storage to paging data sets)
Trace Activity report variable [74](#)

RCEVIOPI (total number VIO pages paged-in)
Trace Activity report variable [75](#)

RCEVIOPO (total number VIO pages paged-out)
Trace Activity report variable [75](#)

RCEVIORE (VIO reads from extended storage)
Trace Activity report variable [75](#)

RCEWSDNE (primary working set pages migrated to auxiliary storage)
Trace Activity report variable [75](#)

RCVAFQA (average available frame count)
Trace Activity report variable [75](#)

RCVAVQC (AVQ low count)
Trace Activity report variable [75](#)

RCVCPUA (CPU usage average * 16)
Trace Activity report variable [75](#)

RCVFXIOP (percentage central storage fixed or allocated)
Trace Activity report variable [75](#)

RCVMFXA (average number fixed frames)
Trace Activity report variable [75](#)

RCVPAGRT (total paging rate)
Trace Activity report variable [75](#)

RCVPTR (paging rate)
Trace Activity report variable [75](#)

RCVSWPTM (time used by ASM to transfer pages to/from data set)
Trace Activity report variable [75](#)

RCVUICA (UIC average)
Trace Activity report variable [75](#)

RECORD/NORECORD (record to SMF data set)
Monitor I session
description [70](#)
summary [60](#)
Monitor II background session
conflicting options [45](#)

REGION parameter
parameter adjustment [19](#)

report interval
Monitor I session
length [66](#)
synchronization [22](#), [60](#), [71](#)

REPORT/NOREPORT (interval report)
Monitor II background session
conflicting options [45](#)

RESOURCE (JES resource)
Monitor III data gatherer session
description [92](#)

- resource name
 - format
 - Monitor I session [64](#)
- REST services
 - application properties [32](#)
 - configuring for high availability [33](#)
 - Monitor II [8](#)
 - Monitor III [9](#)
 - setting up [30](#)
 - SMF [8](#)
- RMCAAWSC (APPC/MVS transaction scheduler wait swap count)
 - Trace Activity report variable [75](#)
- RMCADWSC (detected wait physical swap count)
 - Trace Activity report variable [75](#)
- RMCAEXSC (exchange on recommendation value swap count)
 - Trace Activity report variable [75](#)
- RMCAFHLD (failed swaps because of outstanding HOLD SYSEVENT)
 - Trace Activity report variable [75](#)
- RMCAICSC (improve central storage use)
 - Trace Activity report variable [75](#)
- RMCAIPSC (improve system paging rate)
 - Trace Activity report variable [75](#)
- RMCALWSC (long wait physical swap count)
 - Trace Activity report variable [75](#)
- RMCAMRSC (make room to swap in a user swapped out too long)
 - Trace Activity report variable [75](#)
- RMCANQSC (CPU enqueue exchange swap count)
 - Trace Activity report variable [75](#)
- RMCAOISC (OMVS input wait)
 - Trace Activity report variable [75](#)
- RMCAOOSC (OMVS output wait)
 - Trace Activity report variable [75](#)
- RMCARSSC (central storage shortage swap count)
 - Trace Activity report variable [76](#)
- RMCATISC (terminal input swap count)
 - Trace Activity report variable [76](#)
- RMCATOSC (terminal output swap count)
 - Trace Activity report variable [76](#)
- RMCATSSC (count of translation swaps)
 - Trace Activity report variable [76](#)
- RMCAUSSC (unilateral swap out count)
 - Trace Activity report variable [76](#)
- RMCAXSSC (auxiliary storage shortage swap count)
 - Trace Activity report variable [76](#)
- RMCTTRPC (pages used for transaction elements)
 - Trace Activity report variable [76](#)
- RMF cataloged procedure
 - JCL description [19](#)
 - session option processing [51](#)
 - starting non-TSO/E session [19](#)
- RMF control session
 - RMF cataloged procedure [19](#)
 - starting
 - without concurrent Monitor I session [39](#)
- RMFGAT (Monitor III cataloged procedure) [47](#)

S

session

- session (*continued*)
 - displaying options in effect for [52](#)
- session command
 - DISPLAY [48](#)
 - MODIFY command
 - example [48](#)
 - modifying gatherer session options [47](#)
 - syntax [47](#)
 - START command
 - Monitor III data gatherer session [98](#)
 - non-interactive session [45](#), [47](#)
 - starting sessions [45](#)
 - STOP command
 - Monitor III data gatherer session [93](#)
 - non-interactive session [49](#)
 - stopping sessions [45](#)
- session control
 - non-interactive session
 - session option processing [51](#)
 - overview [45](#)
 - starting [45](#)
- session duration
 - Monitor I session (STOP/NOSTOP)
 - description [70](#)
 - summary [60](#)
 - Monitor II background session (STOP/NOSTOP)
 - conflicting options [45](#)
- session identifier
 - STOP session command [49](#)
 - use in DISPLAY session command [48](#)
 - use in MODIFY session command [47](#)
- session option
 - format [46](#)
 - MODIFY command [47](#)
 - Monitor I session
 - displaying [48](#)
 - specifying [46](#)
- session options
 - default [60](#)
 - displaying those in effect [52](#)
 - Monitor I [61](#)
 - Monitor III data gatherer
 - descriptions [85](#)
- setting up z/OS Data Gatherer [15](#)
- SG/NOSG (storage group) sub-option
 - Monitor I session device activity option [63](#)
- SGSPACE/NOSGSPACE
 - data gatherer [92](#)
 - Monitor III data gatherer session [92](#)
- short-term data collection with Monitor III [4](#)
- shortcut keys [101](#)
- SMCABFLS (records lost because of a shortage of buffers)
 - Trace Activity report variable [76](#)
- SMCABFWT (number of buffers written)
 - Trace Activity report variable [76](#)
- SMCADSCT (records lost because of full data set)
 - Trace Activity report variable [76](#)
- SMCANMFL (current number of full buffers)
 - Trace Activity report variable [76](#)
- SMCARCWT (number of records written)
 - Trace Activity report variable [76](#)
- SMF buffer [7](#), [17](#), [19](#), [42](#)
- SMF record
 - Monitor I and Monitor III synchronization [22](#)

- SMF record (*continued*)
 - Monitor I session
 - RECORD/NORECORD option description [70](#)
 - RECORD/NORECORD option summary [60](#)
- SMF record writing
 - activating [22](#)
 - suppressing [22](#)
- SMF REST services [8](#)
- SMF synchronization
 - defining [23](#)
- SMFBUF option [19](#), [42](#), [49](#)
- Snapshot monitoring with Monitor II [4](#)
- specifying priority for data gatherer started tasks [20](#)
- START command
 - Monitor III data gatherer session [100](#)
 - session command
 - Monitor III data gatherer session [98](#)
 - non-interactive session [45](#), [47](#)
 - session option processing [51](#)
 - system command
 - session option processing [51](#)
- START|STOP suboption
 - DATASET option
 - Monitor III data gatherer session [96](#)
 - Monitor III data gatherer session [97](#)
- starting
 - Monitor II background session [53](#)
- starting a session [52](#)
- statistical sampling
 - Monitor I session [78](#)
- status
 - displaying for data gatherer sessions [48](#)
- STOP command
 - session command
 - Monitor III data gatherer session [93](#)
 - system command [41](#), [49](#)
- STOP/NOSTOP (session duration)
 - data gatherer [93](#)
 - Monitor I session
 - description [70](#)
 - summary [60](#)
 - Monitor II background session
 - conflicting options [45](#)
 - Monitor III data gatherer session
 - description [93](#)
- stopping
 - non-interactive session [49](#)
- storage group (SG/NOSG) sub-option
 - Monitor I session device activity option [63](#)
- storing data [5](#)
- summary of changes [xvii](#)
- SWITCH|NOSWITCH suboption
 - DATASET option
 - Monitor III data gatherer session [96](#)
 - Monitor III data gatherer session [97](#)
- symbolic substitution [25](#)
- SYNC/NOSYNC (interval synchronization)
 - Monitor I session
 - description [71](#)
 - summary [60](#)
 - synchronizing with Monitor III [22](#)
 - Monitor III data gatherer session
 - description [93](#)
- syntax diagrams

- syntax diagrams (*continued*)
 - how to read [xii](#)
- SYS1.PROCLIB system procedure library
 - RMF cataloged procedure [19](#)
- SYSOUT (sysout class)
 - Monitor I session
 - description [71](#)
 - summary [60](#)
 - Monitor III data gatherer session
 - description [93](#)
- sysplex
 - starting z/OS Data Gatherer in [40](#)
- Sysplex
 - considerations for [21](#)
- Sysplex Data Server
 - accessing data [7](#)
 - activating [7](#)
 - data flow [7](#)
 - overview [6](#)
 - programming interface services [7](#)
 - services for Monitor II data [8](#)
 - services for Monitor III data [8](#)
 - services for SMF data [7](#)
 - SMF buffer [7](#), [17](#)
- sysplex-wide coupling facility hardware data collection, optimized [43](#)
- system commands
 - CANCEL command [41](#)
 - MODIFY command
 - non-interactive session [45](#)
 - passing DISPLAY session command [48](#)
 - passing MODIFY session commands [47](#)
 - passing STOP session command [49](#)
 - START command [39](#)
 - STOP command [41](#), [49](#)
- system environment
 - customizing [15](#)
- system paging activity
 - Monitor I session (PAGING/NOPAGING)
 - description [70](#)
 - summary [60](#)

T

- this document
 - who should use [xi](#)
- TRACE/NOTRACE (trace activity)
 - Monitor I session
 - summary [60](#)
- trademarks [106](#)

U

- UNIX System Services
 - access to UNIX System Services [17](#)
 - OMVS process data [90](#)
- user exit routine
 - Monitor I session (EXITS/NOEXITS)
 - description [65](#)
 - summary [60](#)
- user interface
 - ISPF [101](#)

user interface (*continued*)
TSO/E [101](#)

V

virtual storage activity
 Monitor I session (VSTOR/NOVSTOR)
 description [77](#)
 summary [60](#)
VMGUEST/NOVMGUEST
 Monitor I session [78](#)
VSAM (Virtual Storage Access Method) data set
 Monitor III data gatherer session
 data set definition [21](#)
 data set selection [97](#)
 modifying [99](#)
 page release control [97](#)
 starting [97](#), [98](#)
 stopping [100](#)
VSAMRLS
 Monitor III data gatherer session
 description [94](#)
VSTOR/NOVSTOR (virtual storage activity)
 Monitor I session
 description [77](#)
 summary [60](#)

W

who should use this document [xi](#)
WHOLD (page release storage value) suboption
 DATASET option
 Monitor III data gatherer session [96](#)
 Monitor III data gatherer session [97](#)
 relationship to WSTOR suboption [98](#)
WKLD/NOWKLD (workload activity)
 Monitor I session
 description [78](#)
 summary [60](#)
WSTOR (buffer size)
 Monitor III data gatherer session
 description [94](#)
WHOLD suboption
 relationship in controlling page space [98](#)

Z

z/OS Data Gatherer
 cataloged procedure
 RMFGAT [47](#)
 controlling session options [55](#)
 dispatching priority [66](#)
 gathering data [4](#)
 non-interactive session
 option specification [45](#)
 session control [51](#)
 starting [47](#)
 overhead reduction [79](#)
 overview of capabilities [1](#)
 REGION parameter [19](#)
 session option processing
 conflicts [51](#)
 controlling [55](#)

z/OS Data Gatherer (*continued*)
 session option processing (*continued*)
 invalid option value [52](#)
 source order [51](#)
 sessions
 control [45](#)
 description [45](#)
 starting [45](#)
 stopping [49](#)
 setting up [15](#)
 skipped data collection [66](#)
 starting
 session option processing [51](#)
 starting in a sysplex [40](#)
 starting without JES [41](#)
 stopping
 non-interactive session [49](#)
 storing gatherer options [24](#)
z/OS OpenTelemetry
 Emitter
 setting up [35](#)
ZFS/NOZFS [95](#)
ZIIPUSE/NOZIIPUSE [95](#)
ZZ (Monitor I session identifier) [46](#)



Product Number: 5655-ZOS

SC31-5703-70

