

z/OS



# Resource Measurement Facility Report Analysis

*Version 2 Release 2*

**Note**

Before using this information and the product it supports, read the information in "Notices" on page 505.

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

© **Copyright IBM Corporation 1990, 2016.**

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

# Contents

**Figures . . . . . ix**

**Tables . . . . . xiii**

**About this document . . . . . xvii**

Who should use this document . . . . . xvii

How this document is organized . . . . . xvii

z/OS information . . . . . xviii

**How to send your comments to IBM . . . . . xix**

If you have a technical problem . . . . . xix

**Summary of changes . . . . . xxi**

Summary of changes for z/OS RMF Report Analysis for Version 2 Release 2, as updated December 2016 . . . . . xxi

Changes for z/OS Version 2 Release 2 . . . . . xxi

Summary of changes for z/OS RMF Report Analysis for Version 2 Release 1, as updated February 2015 . . . . . xxii

Changes made in z/OS Version 2 Release 1 . . . . . xxiii

Statistics about CF structures residing in

Storage Class Memory . . . . . xxiii

Monitoring PCIe function and zEDC activity . . . . . xxiii

Support of Group Capacity enhancements and absolute LPAR capacity limits . . . . . xxiii

New channel path details in Monitor III and Postprocessor coupling facility reports . . . . . xxiv

Enhanced Postprocessor Crypto Hardware Activity report . . . . . xxiv

Additional Postprocessor reports in XML format . . . . . xxiv

Cross platform monitoring support for Windows . . . . . xxiv

SMF Recording Facility for AIX, Linux and Windows performance data . . . . . xxv

Monitoring of pageable large pages activity . . . . . xxv

Support of Storage Class Memory for paging z/OS Version 2 Release 1 summary of changes . . . . . xxv

**Chapter 1. Introducing RMF data gathering and reporting . . . . . 1**

Gathering data . . . . . 1

Short-term data collection with Monitor III . . . . . 1

Snapshot monitoring with Monitor II . . . . . 1

Long-term data gathering with Monitor I and Monitor III . . . . . 1

Long-term overview reporting with the Postprocessor . . . . . 2

Report analysis with the Spreadsheet Reporter . . . . . 2

Monitoring on the workstation . . . . . 2

What you can gather and report . . . . . 3

Activity monitoring . . . . . 3

Delay monitoring . . . . . 4

Long-term performance analysis with RMF XP . . . . . 5

Reporting of other SMF data . . . . . 5

**Chapter 2. Interactive performance analysis with Monitor III . . . . . 7**

Using Monitor III reports . . . . . 7

System activities measured . . . . . 7

Where to start . . . . . 7

Using cursor-sensitive control . . . . . 11

Common Monitor III report measurements . . . . . 12

Monitor III MINTIME and range . . . . . 15

Monitor III report options . . . . . 17

Monitor III sysplex support in different time zones . . . . . 17

The reporter session . . . . . 18

Starting and stopping a Monitor III reporter session . . . . . 18

The Primary Menu . . . . . 19

Selecting a report . . . . . 19

Monitor III report commands . . . . . 20

Header for single-system reports . . . . . 23

Header for sysplex reports . . . . . 23

The Sysplex Report Selection Menu . . . . . 24

The Overview Report Selection Menu . . . . . 25

The Job Report Selection Menu . . . . . 25

The Resource Report Selection Menu . . . . . 26

The Subsystem Report Selection Menu . . . . . 27

The User Report Selection Menu . . . . . 27

The Data Index . . . . . 28

How to request the Data Index . . . . . 28

Contents of the Data Index . . . . . 28

CACHDET - Cache Detail Report . . . . . 35

How to request this report . . . . . 35

Contents of the report . . . . . 35

CACHSUM - Cache Summary Report . . . . . 39

How to request this report . . . . . 39

Contents of the report . . . . . 40

CFACT - Coupling Facility Activity Report . . . . . 42

How to request this report . . . . . 43

Special considerations . . . . . 43

Contents of the report . . . . . 43

CFOVER - Coupling Facility Overview Report . . . . . 50

How to request this report . . . . . 50

Contents of the report . . . . . 50

CFSYS - Coupling Facility Systems Report . . . . . 52

How to request this report . . . . . 52

Contents of the report . . . . . 52

CHANNEL - Channel Path Activity Report . . . . . 56

How to request this report . . . . . 56

Special considerations of report output . . . . . 56

Contents of the report . . . . . 57

CPC - CPC Capacity Report . . . . . 60

How to request this report . . . . . 60

Contents of the report . . . . . 60

DELAY - Delay Report . . . . . 63

How to request this report . . . . . 63

Contents of the report . . . . .	64	Operator Message and Mount Delay variations	121
DEV - Device Delays Report . . . . .	70	Processor Delay variation . . . . .	122
How to request this report . . . . .	70	Quiesce variation . . . . .	123
Contents of the report . . . . .	71	Storage Delay variation . . . . .	123
DEVN - Device Activity Report . . . . .	72	XCF variation . . . . .	124
How to request this report . . . . .	73	LOCKSP - Spin Lock Report . . . . .	126
Contents of the report . . . . .	73	How to request this report . . . . .	126
DEVR - Device Resource Delays Report . . . . .	75	Contents of the report . . . . .	126
How to request this report . . . . .	75	LOCKSU - Suspend Lock Report . . . . .	128
Contents of the report . . . . .	76	How to request this report . . . . .	128
DEVT - Device Activity Trend Report . . . . .	80	Contents of the report . . . . .	128
How to request this report . . . . .	80	OPD - OMVS Process Data Report . . . . .	129
Contents of the report . . . . .	80	How to request this report . . . . .	130
DSD - Detailed Storage Delays Report . . . . .	81	Contents of the report . . . . .	130
How to request this report . . . . .	81	PCIE - PCIE Activity Report . . . . .	133
Contents of the report . . . . .	81	How to request this report . . . . .	133
DSND - Data Set Delays Report . . . . .	82	Contents of the report . . . . .	133
How to request this report . . . . .	83	Cursor sensitive control . . . . .	135
Special considerations . . . . .	83	PROC - Processor Delays Report . . . . .	136
Contents of the report . . . . .	84	How to request this report . . . . .	136
DSNJ - Data Set Delays - Job Report . . . . .	85	Contents of the report . . . . .	137
How to request this report . . . . .	86	PROCU - Processor Usage Report . . . . .	139
Contents of the report . . . . .	86	How to request this report . . . . .	139
DSNV - Data Set Delays - Volume Report . . . . .	87	Contents of the report . . . . .	140
How to request this report . . . . .	87	RG - Resource Group Data Report . . . . .	141
Contents of the report . . . . .	88	How to request this report . . . . .	141
ENCLAVE - Enclave Report . . . . .	89	Contents of the report . . . . .	141
How to request this report . . . . .	90	RLSLRU - VSAM LRU Overview Report . . . . .	143
Contents of the report . . . . .	90	How to request this report . . . . .	143
Field descriptions . . . . .	90	Contents of the report . . . . .	143
Cursor-sensitive control on the Enclave Report	91	RLSSC/RLSDS - VSAM RLS Activity Report . . . . .	145
Report options . . . . .	93	How to request this report . . . . .	145
ENQ - Enqueue Delays Report . . . . .	95	Contents of the report . . . . .	145
How to request this report . . . . .	95	SCM - Storage Class Memory (SCM) Activity	
Contents of the report . . . . .	95	Report . . . . .	150
ENQR - Enqueue Resource Delays Report . . . . .	99	How to request this report . . . . .	151
How to request this report . . . . .	99	Contents of the report . . . . .	151
Contents of the report . . . . .	99	SPACED - Disk Space Report . . . . .	152
GROUP - Group Response Time Report . . . . .	101	How to request this report . . . . .	152
How to request this report . . . . .	101	Contents of the report . . . . .	153
Contents of the report . . . . .	102	SPACEG - Storage Space Report . . . . .	153
HSM - Hierarchical Storage Manager Delays		How to request this report . . . . .	153
Report . . . . .	109	Contents of the report . . . . .	154
How to request this report . . . . .	109	STOR - Storage Delays Report . . . . .	154
Contents of the report . . . . .	110	How to request this report . . . . .	154
IOQUEUE - I/O Queuing Activity Report . . . . .	111	Contents of the report . . . . .	155
How to request this report . . . . .	112	STORC - Common Storage Report . . . . .	156
Special considerations of report output . . . . .	112	How to request this report . . . . .	157
Contents of the report . . . . .	112	Contents of the report . . . . .	158
JES - Job Entry Subsystem Delays Report . . . . .	114	STORCR - Common Storage Remaining Report . . . . .	160
How to request this report . . . . .	114	How to request this report . . . . .	160
Contents of the report . . . . .	115	Contents of the report . . . . .	161
JOB - Job Delay Report . . . . .	117	STORF - Storage Frames Report . . . . .	162
How to request this report . . . . .	117	How to request this report . . . . .	162
Contents of the report . . . . .	117	Contents of the report . . . . .	162
Displaying the Job Delays report . . . . .	117	STORM - Storage Memory Objects Report . . . . .	164
The top and bottom parts of the Job Delays		How to request this report . . . . .	164
report . . . . .	118	Contents of the report . . . . .	164
Device Delay variation . . . . .	118	STORR - Storage Resource Delays Report . . . . .	166
Enqueue Delay variation . . . . .	120	How to request this report . . . . .	166
HSM and JES variations . . . . .	120	Contents of the report . . . . .	167

STORS - Storage Delay Summary Report . . . . .	170	Report options . . . . .	255
How to request this report . . . . .	170	CHANNEL - Channel Path Activity report . . . . .	255
Contents of the report . . . . .	171	How to request this report . . . . .	255
SYSENQ - Sysplex Enqueue Delays Report . . . . .	174	Special considerations of report output . . . . .	256
How to request this report . . . . .	174	Contents of the report . . . . .	256
Contents of the report . . . . .	175	DEV/DEVV - Device Activity report . . . . .	258
SYSDINFO - System Information Report . . . . .	176	How to request this report . . . . .	259
How to request this report . . . . .	176	Special considerations of report output . . . . .	259
Contents of the report . . . . .	177	Contents of the report . . . . .	261
SYSRTD - Response Time Distribution Report . . . . .	181	HFS - Hierarchical File System Statistics report . . . . .	266
How to request this report . . . . .	181	How to request this report . . . . .	266
Contents of the report . . . . .	181	Contents of the report . . . . .	267
SYSSUM - Sysplex Summary report . . . . .	185	ILOCK - IRLM Long Lock Detection report . . . . .	269
How to request this report . . . . .	186	How to request this report . . . . .	269
Contents of the report . . . . .	187	Special considerations . . . . .	269
SYSTREND - System Trend Report . . . . .	196	Contents of the report . . . . .	270
How to request this report . . . . .	196	IOQUEUE - I/O Queuing Activity report . . . . .	271
Contents of the report . . . . .	196	How to request this report . . . . .	272
SYSWKM - Work Manager Delays Report . . . . .	198	Special considerations of report output . . . . .	272
How to request this report . . . . .	198	Contents of the report . . . . .	273
Contents of the report . . . . .	199	LLI - Library List report . . . . .	276
USAGE - Monitor III Job USAGE Report . . . . .	206	How to request this report . . . . .	276
How to request this report . . . . .	206	Contents of the report - Link Library List . . . . .	276
Contents of the report . . . . .	206	Contents of the report - LPA Library List . . . . .	277
Field descriptions . . . . .	207	Contents of the report - APF Library List . . . . .	277
Monitor III Utility fields . . . . .	208	OPT - OPT Settings report . . . . .	278
Cursor-sensitive control on the USAGE Report . . . . .	209	How to request this report . . . . .	278
WFEX - Workflow/Exceptions Report . . . . .	210	Contents of the report . . . . .	279
How to request this report . . . . .	211	PGSP - Page Data Set Activity report . . . . .	280
Contents of the report . . . . .	211	How to request this report . . . . .	280
Workflow/Exceptions graphic Report . . . . .	225	Special considerations of report output . . . . .	280
XCF - Cross-System Coupling Facility Delays Report . . . . .	226	Contents of the report . . . . .	281
How to request this report . . . . .	226	SDS - Sysplex Data Server report . . . . .	282
Contents of the report . . . . .	226	How to request this report . . . . .	282
ZFSFS - zFS File System . . . . .	228	Contents of the report . . . . .	283
How to request this report . . . . .	228	SENQ - System Enqueue Contention report . . . . .	284
Contents of the report . . . . .	229	How to request this report . . . . .	284
ZFSKN - zFS Kernel report . . . . .	232	Different types of SENQ reports . . . . .	284
How to request this report . . . . .	232	Contents of the report . . . . .	285
Contents of the report . . . . .	233	SENQR - System Enqueue Reserve report . . . . .	288
ZFSOVW - zFS Overview Report . . . . .	233	How to request this report . . . . .	288
How to request this report . . . . .	234	Contents of the report . . . . .	289
Contents of the report . . . . .	234	SPAG - Paging Activity report . . . . .	290
		How to request this report . . . . .	290
		Contents of the report . . . . .	291
		SRCS - Central Storage/Processor/SRM report . . . . .	292
		How to request this report . . . . .	292
		Contents of the report . . . . .	293
<b>Chapter 3. Snapshot reporting with Monitor II . . . . .</b>	<b>241</b>		
Monitor II sessions . . . . .	241	<b>Chapter 4. Real-time reporting with Monitor I . . . . .</b>	<b>295</b>
Structure of Monitor II reports . . . . .	242		
Contents of the Monitor II report header . . . . .	242	<b>Chapter 5. Long-term overview reporting with the Postprocessor . . . . .</b>	<b>297</b>
Different formats of Monitor II reports . . . . .	243	Interval and duration reports . . . . .	298
Different modes of Monitor II reports . . . . .	243	Single-system report header . . . . .	300
ARD/ARDJ - Address Space Resource Data report . . . . .	246	Sysplex report header . . . . .	301
How to request this report . . . . .	246	Duration report header . . . . .	301
Contents of the report . . . . .	246	Overview condition names . . . . .	301
ASD/ASDJ - Address Space State Data report . . . . .	250	CACHE - Cache Subsystem Activity report . . . . .	301
How to request this report . . . . .	250		
Contents of the report . . . . .	250		
ASRM/ASRMJ - Address Space SRM Data report . . . . .	253		
How to request this report . . . . .	253		
Contents of the report . . . . .	253		

How to request this report . . . . .	302	Missing data in report fields . . . . .	404
Different report levels . . . . .	302	Messages . . . . .	405
Cache Summary reporting . . . . .	303	Contents of the report . . . . .	406
Subsystem-level reporting . . . . .	304	Spreadsheet and Overview reference . . . . .	413
Device-level reporting . . . . .	311	OMVS - OMVS Kernel Activity report . . . . .	413
Spreadsheet and Overview reference . . . . .	314	How to request this report . . . . .	414
CF - Coupling Facility Activity report . . . . .	316	Contents of the report . . . . .	414
How to request this report . . . . .	316	PAGESP - Page Data Set Activity report . . . . .	416
Contents of the report . . . . .	317	How to request this report . . . . .	416
Spreadsheet and Overview reference . . . . .	333	Using the information given in the report . . . . .	416
CHAN - Channel Path Activity report . . . . .	334	Contents of the report . . . . .	417
How to request this report . . . . .	335	Overview reference . . . . .	418
Contents of the report . . . . .	336	PAGING - Paging Activity report . . . . .	418
Spreadsheet and Overview reference . . . . .	339	How to request this report . . . . .	418
CPU - CPU Activity report . . . . .	340	Using the information given in the report . . . . .	419
How to request this report . . . . .	341	Contents of the report . . . . .	419
Contents of the report . . . . .	341	PCIE - PCIE Activity Report . . . . .	430
Using the information in the CPU Activity report . . . . .	344	How to request this report . . . . .	430
Using the information in the Partition Data Report. . . . .	350	Contents of the report . . . . .	430
Using the information in the LPAR Cluster Report. . . . .	356	Spreadsheet and Overview reference . . . . .	432
Using the information in the Group Capacity Report. . . . .	357	SCM - SCM Activity Report . . . . .	433
Using the information in the Hardware Group Report. . . . .	359	How to request this report . . . . .	433
Duration report. . . . .	360	Contents of the report . . . . .	434
CRYPTO - Crypto Hardware Activity report . . . . .	361	SDELAY - Serialization Delay report. . . . .	435
How to request this report . . . . .	362	How to request this report . . . . .	435
Contents of the report . . . . .	362	Contents of the report . . . . .	436
Spreadsheet and Overview reference . . . . .	366	SDEVICE - Shared Device Activity report . . . . .	442
DEVICE - Device Activity report . . . . .	367	How to request this report . . . . .	443
How to request this report . . . . .	368	Using the information given in the report . . . . .	443
Contents of the report . . . . .	368	Shared Direct Access Device Activity report . . . . .	443
Spreadsheet and Overview reference . . . . .	376	Shared Magnetic Tape Device Activity report . . . . .	444
DOMINO - Lotus Domino Server report . . . . .	377	Spreadsheet and Overview reference . . . . .	449
How to request this report . . . . .	378	TRACE - Trace Activity report. . . . .	450
Contents of the report . . . . .	378	How to request this report . . . . .	450
ENQ - Enqueue Activity report . . . . .	381	Using the information given in the report . . . . .	450
How to request this report . . . . .	382	Contents of the report . . . . .	450
Different report levels . . . . .	382	VSTOR - Virtual Storage Activity report . . . . .	454
Contents of the report . . . . .	383	How to request this report . . . . .	455
Messages . . . . .	385	Contents of the report . . . . .	455
ESS - Enterprise Disk Systems report . . . . .	386	Data gathering considerations . . . . .	455
How to request this report . . . . .	387	Overview reference . . . . .	458
Contents of the report . . . . .	387	Common Storage Detail section . . . . .	458
FCD - FICON Director Activity report . . . . .	392	Private Area Summary section. . . . .	459
How to request this report . . . . .	392	Private Area Detail section . . . . .	462
Contents of the report . . . . .	393	High Virtual Memory Usage section. . . . .	463
HFS - Hierarchical File System Statistics report . . . . .	394	WLMGL - Workload Activity report . . . . .	463
How to request this report . . . . .	394	How to request this report . . . . .	463
Contents of the report . . . . .	395	The WLMGL option list . . . . .	464
Special considerations . . . . .	398	Contents of the report . . . . .	464
HTTP - HTTP Server report . . . . .	398	Service Class Period report . . . . .	464
How to request this report . . . . .	398	Service Class report . . . . .	467
Contents of the report . . . . .	399	Workload Group report . . . . .	468
IOQ - I/O Queuing Activity report . . . . .	403	Workload Group and Service Class Period report . . . . .	468
How to request this report . . . . .	404	Policy Summary report . . . . .	469
Using the information given in the report . . . . .	404	Report Class report . . . . .	469
Data gathering considerations . . . . .	404	Report Class Period report . . . . .	470
		Service Policy page . . . . .	470
		Field descriptions for all reports . . . . .	471
		Spreadsheet and Overview reference . . . . .	483

XCF - Cross-System Coupling Facility Activity report . . . . .	486
How to request this report . . . . .	486
Contents of the report . . . . .	486
Exception report . . . . .	490
How to request this report . . . . .	490
Contents of the report . . . . .	490
Improved exception reporting . . . . .	492
Overview report . . . . .	494
How to request this report . . . . .	494
Report description. . . . .	494
Summary report . . . . .	498
How to request this report . . . . .	498
Generating a report . . . . .	499
Special considerations of report output . . . . .	499
Spreadsheet reference. . . . .	500

<b>Appendix. Accessibility . . . . .</b>	<b>501</b>
Accessibility features . . . . .	501
Consult assistive technologies . . . . .	501
Keyboard navigation of the user interface . . . . .	501
Dotted decimal syntax diagrams . . . . .	501
<b>Notices . . . . .</b>	<b>505</b>
Policy for unsupported hardware. . . . .	506
Minimum supported hardware . . . . .	507
Programming interface information . . . . .	507
Trademarks . . . . .	507
<b>Glossary . . . . .</b>	<b>509</b>
<b>Index . . . . .</b>	<b>515</b>





# Figures

1. Suggested Sequence for Using Monitor III Reports . . . . .	10	50. ENCLAVE Report . . . . .	90
2. Monitor III Primary Menu . . . . .	19	51. ENCLAVE Report - Enclave Details . . . . .	92
3. Header of Monitor III Single-System Reports	23	52. ENCLAVE Report - Enclave Classification Attributes (1) . . . . .	93
4. Header of Monitor III Sysplex Reports . . . . .	24	53. ENCLAVE Report Options . . . . .	94
5. Monitor III Sysplex Report Selection Menu	24	54. ENQ Report . . . . .	95
6. Monitor III Overview Report Selection Menu	25	55. ENQR Report . . . . .	99
7. Monitor III Job Report Selection Menu . . . . .	26	56. ENQR Report Options Panel . . . . .	100
8. Monitor III Resource Report Selection Menu	27	57. GROUP Report . . . . .	102
9. Monitor III Subsystem Report Selection Menu	27	58. GROUP Report - Response Time Components	103
10. Monitor III User-written Report Selection Menu . . . . .	28	59. GROUP Report Options Panel . . . . .	109
11. Data Index . . . . .	29	60. HSM Report . . . . .	110
12. Data Index with Preallocated Data Sets - Detailed View . . . . .	30	61. IOQUEUE Report . . . . .	112
13. Data Index - Condensed Version . . . . .	31	62. JES Delays report . . . . .	115
14. Data Index with Preallocated Data Sets - Condensed Version . . . . .	31	63. Top Part of Job Delays report . . . . .	118
15. Data Index Options Panel . . . . .	33	64. Bottom Part of Job Delay report . . . . .	118
16. Data Index - Sort Order Descend . . . . .	34	65. Device Delay variation of the Job Delay report . . . . .	119
17. CACHDET Report . . . . .	35	66. Enqueue Delay variation of the Job Delay report . . . . .	120
18. CACHDET Report - Volume Details . . . . .	36	67. HSM Delay variation of Job Delay report	121
19. CACHDET Report - SSID Details . . . . .	36	68. Operator Message Delay variation of the Job Delay report . . . . .	122
20. CACHDET Report Options . . . . .	39	69. Processor Delay variation of the Job Delay report . . . . .	122
21. CACHSUM Report . . . . .	40	70. Quiesce Delay variation of the Job Delay report . . . . .	123
22. CACHSUM Report - SSID Details . . . . .	40	71. Storage Delay variation of Job Delay report	123
23. CFACT Report . . . . .	43	72. XCF Delay variation of Job Delay report	124
24. CFACT Report - Details for a Lock Structure	44	73. Job Report Options Panel . . . . .	125
25. CFACT Report - Details for a List Structure	44	74. Spin Lock Report . . . . .	127
26. CFACT Report - Details for a Cache Structure	45	75. Lock Report Options . . . . .	127
27. CFACT Report - Details for a List Structure using SCM storage (1) . . . . .	45	76. Suspend Lock Report . . . . .	128
28. CFACT Report - Details for a List Structure using SCM storage (2) . . . . .	46	77. OPD Report . . . . .	130
29. Coupling Facility Report Options Panel	49	78. OPD Report - Details for Server Process	131
30. CFOVER Report . . . . .	50	79. OPD Report Options Panel . . . . .	132
31. CFSYS Report . . . . .	52	80. PCIE Activity Report . . . . .	133
32. CFSYS Report - Subchannels and Paths	53	81. RMF Hardware Accelerator And Compression Activity panel . . . . .	135
33. CHANNEL Report . . . . .	57	82. PROC - Processor Delay Report . . . . .	137
34. CPC Capacity report . . . . .	61	83. PROCU - Processor Usage Report . . . . .	140
35. DELAY Report . . . . .	64	84. RG Report . . . . .	142
36. DELAY Report Options Panel . . . . .	68	85. RLRLRU Report . . . . .	143
37. DELAY Report Job Selection/Exclusion Panel	70	86. VSAM LRU Overview - Buffer Counts by Pool . . . . .	144
38. DEV Report . . . . .	71	87. VSAM RLS Activity by Storage Class - Sysplex Total View . . . . .	146
39. DEV Report Options Panel . . . . .	72	88. VSAM RLS Activity by Storage Class - System/CF Structure View . . . . .	147
40. DEVN Report . . . . .	73	89. VSAM RLS Activity by Data Set - Sysplex Total View . . . . .	148
41. DEVR Report . . . . .	76	90. VSAM RLS Activity by Data Set - System/CF Structure View . . . . .	149
42. DEVR Report Options Panel . . . . .	79	91. SCM Activity Report . . . . .	151
43. DEVT Report . . . . .	80	92. Disk Space Report . . . . .	153
44. Modified STOR Report Showing all Storage Delays in Detail . . . . .	82	93. Storage Space Report . . . . .	154
45. DSND Report . . . . .	84		
46. DSND Report Options Panel . . . . .	85		
47. DSNJ Report . . . . .	86		
48. DSNV Report . . . . .	88		
49. DSNV Report Options Panel . . . . .	89		

94.	STOR Report . . . . .	155	150.	DEV Report Options Panel . . . . .	265
95.	STORC Report . . . . .	158	151.	DEVV Report Options Panel . . . . .	266
96.	STORC Report Options Panel . . . . .	160	152.	HFS Report . . . . .	267
97.	STORCR Report. . . . .	161	153.	HFS Report Options Panel . . . . .	268
98.	STORF Report . . . . .	162	154.	ILOCK Report . . . . .	270
99.	STORM Report . . . . .	164	155.	IOQUEUE Report . . . . .	273
100.	STORR Report . . . . .	167	156.	IOQUEUE Report Options Panel . . . . .	275
101.	STORS Report . . . . .	171	157.	LLI Report - Link Library List . . . . .	276
102.	STORS Report Options Panel . . . . .	173	158.	LLI Report - LPA Library List . . . . .	277
103.	SYSENQ Report. . . . .	175	159.	LLI Report - APF Library List . . . . .	277
104.	SYSINFO Report . . . . .	177	160.	OPT Settings. . . . .	279
105.	SYSRTD Report - With Response Time Data	181	161.	PGSP Report. . . . .	281
106.	SYSRTD Report - Without Response Time Data . . . . .	182	162.	SDS Report . . . . .	283
107.	SYSSUM Report - GO Mode . . . . .	187	163.	SENQ Summary Report . . . . .	285
108.	SYSSUM Report - STOP Mode . . . . .	188	164.	SENQ GRS Report . . . . .	286
109.	SYSSUM Report for Workload Groups	190	165.	SENQ Detail Report . . . . .	286
110.	SYSTREND Report. . . . .	197	166.	SENQ Report Options Panel . . . . .	287
111.	SYSWKM Report for Subsystem CICS	199	167.	SENQR Report . . . . .	289
112.	SYSWKM Report for Subsystem IMS	200	168.	SENQR Report Options Panel . . . . .	290
113.	SYSWKM Report Options Panel . . . . .	205	169.	SPAG Report. . . . .	291
114.	Job Usage report . . . . .	207	170.	SRCS Report. . . . .	293
115.	WFEX Report . . . . .	211	171.	Header of a Single-System Report. . . . .	300
116.	WFEX Report Options Action Panel . . . . .	216	172.	Header of a Sysplex report . . . . .	301
117.	WFEX Definition and Criteria Panel . . . . .	217	173.	Header of a Single-System Duration Report	301
118.	WFEX Definition of UNAVAIL Panel	220	174.	Cache Subsystem Summary Report - Subsystem Summary . . . . .	303
119.	WFEX Criteria Names for Class: Job . . . . .	222	175.	Cache Subsystem Activity Report - Top-20 Device Lists . . . . .	304
120.	WFEX Criteria Names for Class: Service Class	223	176.	Cache Subsystem Activity Report - Status and Overview . . . . .	305
121.	WFEX Criteria Names for Class: Processor	223	177.	Cache Subsystem Activity Report - Device Overview . . . . .	310
122.	WFEX Criteria Names for Class: Storage	224	178.	Cache Subsystem Activity Report - Cache Device Activity (device-level reporting) . . . . .	312
123.	WFEX Criteria Names for Class: Device	224	179.	Cache Subsystem Activity Report - Cache Offline . . . . .	314
124.	WFEX Graphic Report . . . . .	225	180.	Structure of the Coupling Facility Activity report . . . . .	318
125.	XCF Report . . . . .	226	181.	Coupling Facility Activity Report - Usage Summary . . . . .	319
126.	ZFSFS - Report Options . . . . .	228	182.	Coupling Facility Activity Report - Structure Activity . . . . .	326
127.	ZFSFS Report . . . . .	229	183.	Coupling Facility Activity Report - Subchannel Activity . . . . .	329
128.	ZFSFS Report - File system details . . . . .	230	184.	Coupling Facility Activity Report - CF to CF Activity . . . . .	332
129.	ZFSFS Report - zFS File System Details	230	185.	Channel Path Activity report . . . . .	336
130.	zFS Kernel Report . . . . .	233	186.	CPU Activity Report . . . . .	345
131.	zFS Overview Report . . . . .	234	187.	CPU Activity Report - Partition Data Report	352
132.	I/O Details by Type . . . . .	236	188.	CPU Activity Report - LPAR Cluster Report	356
133.	User Cache Details. . . . .	237	189.	CPU Activity Report - Group Capacity Report	358
134.	Vnode Cache Details . . . . .	238	190.	CPU Activity Report - Hardware Group Report . . . . .	359
135.	Metadata Cache Details . . . . .	239	191.	Crypto Hardware Activity Report. . . . .	362
136.	Header of a Monitor II TSO/E session report	242	192.	Direct Access Device Activity Report	372
137.	Format of printed reports from a Monitor II display session . . . . .	244	193.	Communication Equipment Activity Report	372
138.	Format of printed reports from a Monitor II background session . . . . .	245	194.	Magnetic Tape Device Activity Report	372
139.	ARD Report in z/Architecture . . . . .	247	195.	Lotus Domino Server Report - Summary	378
140.	ARDJ Report for a system running in 31-bit mode . . . . .	247	196.	Lotus Domino Server Report - Details	379
141.	ARD, ASD, and ASRM Report Options Panel	249	197.	Enqueue Detail Activity Report . . . . .	383
142.	ARDJ, ASDJ, and ASRMJ Report Options Panel . . . . .	249	198.	ESS Link Statistics . . . . .	388
143.	ASD Report . . . . .	251			
144.	ASDJ Report . . . . .	251			
145.	ASRM Report . . . . .	254			
146.	ASRMJ Report . . . . .	254			
147.	CHANNEL Report. . . . .	256			
148.	DEV Report . . . . .	261			
149.	DEVV Report . . . . .	261			

199.	ESS Extent Pool Statistics . . . . .	389	220.	Shared Magnetic Tape Device Activity Report	444
200.	ESS Rank Statistics . . . . .	390	221.	TRACE Report . . . . .	451
201.	FICON Director Activity Report . . . . .	393	222.	VSTOR report - Common Storage Summary	457
202.	HFS Global Statistics Report . . . . .	395	223.	VSTOR report - Common Storage Detail	459
203.	HFS File System Statistics Report . . . . .	397	224.	VSTOR report - Private Area Summary	460
204.	HTTP Server Report - Summary . . . . .	399	225.	VSTOR report - Private Area Detail and High Virtual Memory Usage (above 2GB) . . . . .	462
205.	HTTP Server Report - Details . . . . .	400	226.	WLMGL - Service Class Period report - with execution velocity goal for Period 1 and response time distributions . . . . .	465
206.	I/O Queuing Activity Report . . . . .	406	227.	WLMGL - Service Class Period report - with response time distribution (percentile) . . . . .	466
207.	OMVS Kernel Activity Report . . . . .	414	228.	WLMGL - Service Class Period report - with subsystem data and response time distribution for response time goal . . . . .	467
208.	PAGESP Report . . . . .	417	229.	WLMGL - Service Class Report . . . . .	468
209.	PAGING Report - Central Storage Paging Rates . . . . .	420	230.	WLMGL - Workload Group with associated service classes . . . . .	469
210.	PAGING Report - Central Storage Movement and Request Rates / Frame and Slot Counts . . . . .	423	231.	WLMGL Report - Service Policy Page	471
211.	PAGING Report - Memory Objects and High Virtual Storage Frames . . . . .	427	232.	XCF Activity Report - Usage by System	487
212.	Postprocessor PCIE Activity Report . . . . .	431	233.	XCF Activity Report - Usage by Member	488
213.	Storage Class Memory (SCM) Activity Report	434	234.	XCF Activity Report - Path Statistics (Coupling Facility and Channel-to-Channel) . . . . .	489
214.	SDELAY Report - Serialization Delay Summary . . . . .	437	235.	Exception Report - Low CPU Utilization	491
215.	SDELAY Report - Serialization Delay Details - CMS Lock Details . . . . .	439	236.	Exception Report - CPU Utilization . . . . .	491
216.	SDELAY Report - Serialization Delay Details - CML and Local Lock Details . . . . .	440	237.	Overview Report - Exception Version	493
217.	SDELAY Report - Serialization Delay Details - GRS Latch Details . . . . .	441	238.	Overview Report - Summary Version	496
218.	SDELAY Report - Serialization Delay Details - GRS Enqueue Details . . . . .	442	239.	Summary Report . . . . .	499
219.	Shared DASD Activity Report . . . . .	444			



## Tables

1. Monitored activities and SMF record types	3	45. Fields in the XCF Delay variation of the Job Delay report	124
2. Report Commands	21	46. Additional Fields in the Job Delay report	125
3. User-Written Report Commands	23	47. Fields in the LSP Report	127
4. Monitor III Heading Information	23	48. Fields in the LSU Report	129
5. Fields in the Data Index	31	49. Fields in the OPD Report	130
6. Fields in the CACHDET Report	37	50. Fields in the OPD Details Report	131
7. Fields in the CACHDET Report - Volume and SSID Details	37	51. Fields in the PCIE Activity Report	133
8. Fields in the CACHSUM Report	41	52. Monitor III Utility fields	134
9. Fields in the CACHSUM Report - SSID Details	41	53. Fields on pop-up panels	135
10. Fields in the CFACT Report	46	54. Fields in the PROC Report	137
11. Fields in the CFOVER Report	50	55. Additional Fields in the PROC Report	139
12. Fields in the CFSYS Report	53	56. Fields in the PROCU Report	140
13. Fields in the CFSYS Report - Subchannels and Paths	54	57. Additional Fields in the PROCU Report	141
14. Fields in the CHANNEL Report	57	58. Fields in the RG Report	142
15. Additional Fields in the CHANNEL Report	59	59. Fields in the RLSLRU Report	144
16. Fields in the CPC Capacity Report	61	60. Fields in the VSAM RLS Activity report	149
17. Fields in the DELAY Report	64	61. Fields in the SCM Activity Report	151
18. Additional Fields in the DELAY Report	67	62. Fields in the Disk Space Report	153
19. Fields in the DEV Report	71	63. Fields in the Storage Space Report	154
20. Fields in the DEVN Report	74	64. Fields in the STOR Report	155
21. DEVN Report - Cursor-sensitive Control for Select and Sort	74	65. Additional Fields in the STOR Report	156
22. Fields in the DEVR Report	76	66. Fields in the STORC Report - System Information Section	158
23. Additional Fields in the DEVR Report	78	67. Fields in the STORC Report - Job Information Section	159
24. Fields in the DEVT Report	81	68. Additional Fields in the STORC Report	160
25. Fields in the DSD Report	82	69. Fields in the STORCR Report	161
26. Fields in the DSND Report	84	70. Fields in the STORF Report	162
27. Fields in the DSNJ Report	86	71. Fields in the STORM Report - System Summary Section	164
28. Fields in the DSNV Report	88	72. Fields in the STORM Report - Address Space Section	165
29. Fields in the ENCLAVE Report	90	73. Fields in the STORR Report - Central Storage Summary Section	168
30. Fields in the ENCLAVE Report - Enclave Details	92	74. Fields in the STORR Report - Page/Swap Activity Section	168
31. Fields in the ENQ Report	95	75. Additional Fields in the STORR Report	170
32. Fields in the ENQR Report	99	76. Fields in the STORS Report	171
33. Fields in the GROUP Report	103	77. Additional Fields in the STORS Report	172
34. Fields in the HSM Report	110	78. Fields in the SYSENG Report	175
35. Fields in the IOQUEUE Report	113	79. Fields in the SYSINFO Report	178
36. Fields in the JES Report	115	80. Additional Fields in the SYSINFO Report	180
37. Fields in the Job Delays report	118	81. Fields in the SYSRTD Report	185
38. Fields in the Device Delay variation of the Job Delay report	119	82. Fields in the SYSSUM Report	191
39. Fields in the Enqueue Delay variation of the Job Delay report	120	83. Fields in the SYSSUM Report - Continuation	194
40. Fields in the HSM/JES delay variation of the Job Delay report	121	84. Additional Fields in the SYSSUM Report	195
41. Fields in the Operator Message and Mount Delay variations of the Job Delay report	122	85. Fields in the SYSTREND Report	197
42. Fields in the Processor Delay variation of the Job Delay report	122	86. SYSTREND Report - Cursor-sensitive Control for Navigation	198
43. Fields in the Quiesce Delay variation of the Job Delay report	123	87. Fields in the SYSWKM Report	202
44. Fields in the Storage Delays variation of the Job Delay report	124	88. Fields in the Job Usage report	207
		89. Additional fields in the Job Usage report	208
		90. Cursor sensitivity on Job Usage report	209
		91. Fields in the Speed Section of the WFEX Report	212

92. Fields in the Exceptions Section of the WFEX Report . . . . .	213	136. Fields in the Coupling Facility Activity Report - Subchannel Activity . . . . .	329
93. Fields in the WFEX Action Panel . . . . .	217	137. Fields in the Coupling Facility Activity Report - Subchannel Activity - Channel Path Details . . . . .	331
94. Fields in the WFEX Definition and Criteria Panel . . . . .	218	138. Fields in the CF to CF Activity Section	332
95. Fields in the Graphic WFEX Report . . . . .	226	139. Fields in the Coupling Facility Activity Report - CF to CF Activity - Channel Path Details . . . . .	333
96. Fields in the XCF Report . . . . .	227	140. Overview names in the Coupling Facility Activity Report . . . . .	333
97. zFS File System report field descriptions	230	141. Fields in the Channel Path Activity report	336
98. zFS File System report field descriptions - zFS File Systems Details panel . . . . .	231	142. Overview names in the Channel Path Activity report . . . . .	339
99. zFS Kernel report field descriptions . . . . .	233	143. Fields in the CPU Activity Report . . . . .	345
100. zFS Overview report field descriptions	234	144. Overview names in the CPU Activity Report	349
101. Fields in the zFS Overview report - I/O details . . . . .	236	145. Fields in the Partition Data Report . . . . .	352
102. Fields in the zFS Overview Report - User Cache Details . . . . .	237	146. Overview names in the Partition Data Report	355
103. Fields in the zFS Overview Report - Vnode Cache Details . . . . .	238	147. Fields in the LPAR Cluster Report . . . . .	356
104. Fields in the zFS Overview Report - Metadata Cache/Backing Cache Details . . . . .	239	148. Overview names in the LPAR Cluster Report	357
105. Monitor II Display Session Areas . . . . .	243	149. Fields in the Group Capacity Report . . . . .	358
106. Fields in the ARD and ARDJ reports . . . . .	247	150. Overview names in the Group Capacity Report . . . . .	359
107. Fields in the ASD and ASDJ Report . . . . .	251	151. Fields in the Hardware Group Report	359
108. Fields in the ASRM and ASRMJ Report	254	152. Overview names in the Hardware Group Report . . . . .	360
109. Fields in the CHANNEL Report . . . . .	257	153. Fields in the CRYPTO Hardware Activity Report . . . . .	362
110. Fields in the DEV and DEVV Report	262	154. Overview names in the CRYPTO Hardware Activity Report . . . . .	366
111. Fields in the HFS Report . . . . .	267	155. Fields in the Device Activity Reports	372
112. Fields in the HFS Report Options Panel	268	156. Overview names in the DASD Activity Report . . . . .	376
113. Fields in the ILOCK Report . . . . .	270	157. Fields in the Domino Server Summary Report	379
114. Fields in the IOQUEUE Report . . . . .	274	158. Fields in the Domino Server Details Report	380
115. Fields in the LLI Report . . . . .	278	159. Fields in the Enqueue Activity Report	383
116. Fields in the OPT Settings report . . . . .	279	160. Overview names in the Enqueue Activity Report . . . . .	385
117. Fields in the PGSP Report . . . . .	281	161. Common Fields in the Enterprise Disk Systems Report . . . . .	387
118. Fields in the SDS Report . . . . .	283	162. Fields in the ESS Link Statistics . . . . .	388
119. Fields in the SENQ Report . . . . .	286	163. Fields in the ESS Extent Pool Statistics	389
120. Fields in the SENQR Report . . . . .	289	164. Fields in the ESS Rank Statistics . . . . .	390
121. Fields in the SPAG Report . . . . .	291	165. Overview names in the ESS Link Statistics section . . . . .	391
122. Fields in the SRCS Report . . . . .	293	166. Overview names in the ESS Extent Pool Statistics section . . . . .	392
123. Available formats for Postprocessor reports	297	167. Overview names in the ESS Rank Statistics section . . . . .	392
124. Interval and Duration Reports . . . . .	299	168. Fields in the FICON Director Activity Report.	393
125. Monitor II Interval Reports . . . . .	300	169. Overview names in the FICON Director Activity Report . . . . .	394
126. Fields in the Cache Subsystem Summary Report . . . . .	304	170. Fields in HFS Global Statistics Report	395
127. Fields in the Cache Subsystem Activity Report - Header . . . . .	305	171. Fields in the HFS File System Statistics Report	397
128. Fields in the Cache Subsystem Activity Report - Subsystem Status . . . . .	306	172. Fields in the HTTP Server Summary Report	400
129. Fields in the Cache Subsystem Activity Report - Subsystem Overview . . . . .	307	173. Fields in the HTTP Server Details Report	401
130. Fields in the Cache Subsystem Activity Report - Device Overview . . . . .	310	174. Fields in the I/O Queuing Activity Report	407
131. Fields in the Cache Subsystem Activity Report - RAID Rank Activity . . . . .	311	175. Overview names in the I/O Queuing Activity Report . . . . .	413
132. Fields in the Cache Subsystem Activity Report - Cache Device Status . . . . .	312	176. Fields in the OMVS Kernel Activity Report	415
133. Overview names in the Cache Subsystem Activity Report . . . . .	314	177. Fields in the Page Data Set Activity report	417



178. Overview names in the Page Data Set Activity report . . . . .	418	198. Variables in the Trace Activity Report	451
179. Fields in the Paging Activity report - Central Storage Paging Rates . . . . .	420	199. Fields in the Virtual Storage Activity Report - Common Storage Summary . . . . .	457
180. Overview names in the Paging Activity report - Central Storage Paging Rates . . . . .	422	200. Overview names in the Virtual Storage Activity Report . . . . .	458
181. Fields in the Paging Activity report - Central Storage Movement and Request Rates . . . . .	423	201. Fields in the Virtual Storage Activity Report - Common Storage Detail Section . . . . .	459
182. Overview names in the Paging Activity report - Central Storage Movement and Request Rates . . . . .	424	202. Fields in the Virtual Storage Activity Report - Private Area Summary . . . . .	460
183. Fields in the Paging Activity report - Frame and Slot Counts . . . . .	425	203. Fields in the Virtual Storage Activity Report - Private Storage Map . . . . .	461
184. Overview names in the Paging Activity report - Frame and Slot Counts . . . . .	426	204. Fields in the Virtual Storage Activity Report - Bottom Half . . . . .	461
185. Fields in the Paging Activity report - Memory Objects and High Virtual Storage Frames . . . . .	428	205. Fields in the Virtual Storage Activity report - Private Area Detail section . . . . .	462
186. Overview names in the Paging Activity report - Memory Objects and High Virtual Storage Frames . . . . .	429	206. Fields in the Virtual Storage Activity report - High Virtual Memory Usage section . . . . .	463
187. Fields in the PCIE Activity Report . . . . .	431	207. Fields in the Workload Activity Report	471
188. Overview conditions in the PCIE Activity Report . . . . .	432	208. Fields in the WLMGL Report - POLICY	482
189. Fields in the SCM Activity Report . . . . .	434	209. Exception and Overview names in the Workload Activity Report . . . . .	484
190. Fields in the Serialization Delay Summary section . . . . .	437	210. Fields in the XCF Activity Report - Usage by System . . . . .	487
191. Fields in the Serialization Delay Details section - CMS Lock Details . . . . .	439	211. Fields in the XCF Activity Report - Usage by Member . . . . .	489
192. Fields in the Serialization Delay Details section - CML and Local Lock Details . . . . .	440	212. Fields in the XCF Activity Report - XCF Path Statistics . . . . .	489
193. Fields in the Serialization Delay Details section - GRS Latch Details . . . . .	441	213. Fields in the Exception Report . . . . .	492
194. Fields in the Serialization Delay Details section - GRS Enqueue Details . . . . .	442	214. Fields in the Overview Report . . . . .	493
195. Fields in the Shared Device Activity Reports	444	215. Overview Header Record - Prefix Section	497
196. Overview names in the Shared DASD Activity Report . . . . .	449	216. Overview Header Record - Header Section	497
197. Fields in the Trace Activity Report . . . . .	451	217. Overview Header Record - Report Column Names Section . . . . .	497
		218. Overview Data Record - Prefix Section	498
		219. Overview Data Record - Data Section	498
		220. Overview Data Record - Report Data Section	498
		221. Fields in the Summary Report . . . . .	499





---

## About this document

The Resource Measurement Facility (RMF™) is a performance management tool that measures selected areas of system activity and presents the data collected in the form of System Management Facility (SMF) records, formatted printed reports, or formatted display reports. You can use this data to evaluate system performance and identify reasons for performance problems.

This document describes all RMF reports in detail, how to generate them, what they contain, their options, and how to use them.

For information about starting RMF and session options, see *z/OS RMF User's Guide*.

### About special purpose processors:

Throughout this document, zIIP refers to IBM System z9® Integrated Information Processors or to IBM System z10™ Integrated Information Processors. zAAP refers to IBM System z Application Assist Processors.

---

## Who should use this document

This document is intended for the system programmer and performance analyst responsible for measuring and improving system performance. Because RMF is a tool for measuring z/OS system performance, this document assumes that the reader has extensive knowledge of the z/OS system. For an overview of RMF, see *z/OS RMF User's Guide*.

---

## How this document is organized

This document contains the following chapters:

**Chapter 1, "Introducing RMF data gathering and reporting," on page 1**

This chapter explains how RMF is divided into monitors, and what sessions run under the different monitors. It also describes what data you can collect using the different monitors and sessions.

**Chapter 2, "Interactive performance analysis with Monitor III," on page 7**

This chapter gives an example of how you can navigate through the Monitor III reports, explains how cursor-sensitive control works, describes some common Monitor III measurements, and explains all reports in detail.

**Chapter 3, "Snapshot reporting with Monitor II," on page 241**

This chapter describes the Monitor II reports, includes example reports, and provides a detailed description of the report fields.

**Chapter 4, "Real-time reporting with Monitor I," on page 295**

This chapter gives you a table of reports you can request when using a Monitor I session. Since all Monitor I reports are also Postprocessor reports, the detailed description of these reports is located in Chapter 5, "Long-term overview reporting with the Postprocessor," on page 297

**Chapter 5, “Long-term overview reporting with the Postprocessor,” on page 297**

This chapter describes the reports you can request using the Postprocessor. The descriptions include report examples and detailed descriptions of the report fields.

---

## **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 V2R2 Information Roadmap*.

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

---

## How to send your comments to IBM

We appreciate your input on this documentation. Please provide us with any feedback that you have, including comments on the clarity, accuracy, or completeness of the information.

Use one of the following methods to send your comments:

**Important:** If your comment regards a technical problem, see instead “If you have a technical problem.”

- Send an email to [mhvrcfs@us.ibm.com](mailto:mhvrcfs@us.ibm.com).
- Send an email from the "Contact us" web page for z/OS (<http://www.ibm.com/systems/z/os/zos/webqs.html>).

Include the following information:

- Your name and address
- Your email address
- Your phone or fax number
- The publication title and order number:  
z/OS V2R2 RMF Report Analysis  
SC34-2665-03
- The topic and page number or URL of the specific information to which your comment relates
- The text of your comment.

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

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

---

## If you have a technical problem

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

- Visit the IBM Support Portal ([support.ibm.com](http://support.ibm.com)).
- Contact your IBM service representative.
- Call IBM technical support.



---

## Summary of changes

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

---

### Summary of changes for z/OS RMF Report Analysis for Version 2 Release 2, as updated December 2016

This edition includes the following topics that contain new and changed information:

#### **New**

“Using the information in the Hardware Group Report” on page 359

#### **Changed**

Table 10 on page 46

Table 15 on page 59

Figure 34 on page 61

Table 16 on page 61

“PCIE - PCIE Activity Report” on page 133

ASD report: “Contents of the report” on page 250

OPT Settings report: “Contents of the report” on page 279

“Coupling Facility Usage Summary section” on page 318

“Coupling Facility Structure Activity section” on page 325

“CF to CF Activity section” on page 332

“Using the information in the Partition Data Report” on page 350

“Using the information in the Group Capacity Report” on page 357

“IOQ - I/O Queuing Activity report” on page 403

Figure 206 on page 406

PCIE Activity Report: “Contents of the report” on page 430

“Spreadsheet and Overview reference” on page 432

“Service Class report” on page 467

“Workload Group and Service Class Period report” on page 468

“Field descriptions for all reports” on page 471

“Spreadsheet and Overview reference” on page 483

Table 47 on page 127

“Contents of the report” on page 362

---

### Changes for z/OS Version 2 Release 2

This edition contains information previously presented in *z/OS V2R1 RMF Report Analysis* (SC34-2665-01)

## New information

This edition includes the following new information:

- A new Monitor III PCIE Activity report is described in “PCIE - PCIE Activity Report” on page 133, Table 2 on page 21, and Figure 8 on page 27
- A new Monitor III SCM Activity report is described in “SCM - Storage Class Memory (SCM) Activity Report” on page 150, Table 2 on page 21, and Figure 8 on page 27
- A new Monitor III USAGE report is described in “USAGE - Monitor III Job USAGE Report” on page 206, Table 2 on page 21, and “The Overview Report Selection Menu” on page 25
- A new Monitor III ZFSFS report is described in “ZFSFS - zFS File System” on page 228, Table 2 on page 21, and “The Sysplex Report Selection Menu” on page 24
- A new Monitor III ZFSKN report is described in “ZFSKN - zFS Kernel report” on page 232, Table 2 on page 21, and “The Sysplex Report Selection Menu” on page 24
- A new Monitor III ZFSOVW report is described in “ZFSOVW - zFS Overview Report” on page 233, Table 2 on page 21, and “The Sysplex Report Selection Menu” on page 24
- “High Virtual Memory Usage section” on page 463

## Changed information

This edition includes the following topics that contain changed information:

- Table 1 on page 3 has been updated.
- “ENQ - Enqueue Activity report” on page 381
- “Private Area Detail section” on page 462
- An new GCMSUAV overview has been added to “Spreadsheet and Overview reference” on page 355 for the Partition Data Report.

## Deleted information

This these topics have been deleted from this edition:

- Monitor III reports ZFSACT and ZFSSUM have been removed. They are replaced by new Monitor III reports ZFSFS, ZFSKN, and ZFSOVW.
- Monitor III Utility fields section of “STORF - Storage Frames Report” on page 162 has been removed.

---

## Summary of changes for z/OS RMF Report Analysis for Version 2 Release 1, as updated February 2015

### Changed information

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

- Table 12 on page 53 has been updated.
- “CPC - CPC Capacity Report” on page 60 has been updated.
- “ENCLAVE - Enclave Report” on page 89 has been updated.
- “OPD - OMVS Process Data Report” on page 129 has been updated.
- “PROC - Processor Delays Report” on page 136 has been updated.

- “PROCU - Processor Usage Report” on page 139 has been updated.
- “SYSINFO - System Information Report” on page 176 has been updated.
- “CPU - CPU Activity report” on page 340 has been updated.
- “CRYPTO - Crypto Hardware Activity report” on page 361 has been updated.
- “TRACE - Trace Activity report” on page 450 has been updated.
- “WLMGL - Workload Activity report” on page 463 has been updated.
- Table 137 on page 331 has been updated.
- “CF to CF Activity section” on page 332 has been updated.

---

## Changes made in z/OS Version 2 Release 1

This document contains information previously presented in *z/OS RMF Report Analysis*, SC33-7991-19, which supports z/OS Version 1 Release 13.

### Statistics about CF structures residing in Storage Class Memory

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

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

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

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

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

### Monitoring PCIe function and zEDC activity

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

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

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

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

### Support of Group Capacity enhancements and absolute LPAR capacity limits

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

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

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

## **New channel path details in Monitor III and Postprocessor coupling facility reports**

New channel path detail information is available for CIB and CFP channels paths. RMF provides this information in the *Subchannel Activity* and the *CF to CF Activity* sections of the Postprocessor *Coupling Facility Activity* report.

Also, the Monitor III *Coupling Facility Systems* report is enhanced to provide a new *Channel Path Details* section in the *Subchannels and Paths* pop-up.

In addition, RMF stores the newly gathered channel path detail information for coupling facilities into SMF record 74-4.

## **Enhanced Postprocessor Crypto Hardware Activity report**

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

- Cryptographic CCA coprocessor
- Cryptographic PKCS11 coprocessor
- Cryptographic accelerator

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

## **Additional Postprocessor reports in XML format**

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

- *Cache Subsystem Activity*
- *Channel Path Activity*
- *Coupling Facility Activity*
- *Enqueue Activity*
- *Hierarchical File System Statistics*
- *I/O Queuing Activity*
- *Page Data Set Activity*
- *PCIE Activity Report*
- *Shared Device Activity*
- *Virtual Storage Activity*
- *XCF Activity*

## **Cross platform monitoring support for Windows**

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



## SMF Recording Facility for AIX, Linux and Windows performance data

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

## Monitoring of pageable large pages activity

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

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

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

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

## Support of Storage Class Memory for paging

RMF provides measurements about storage type SCM (storage class memory) in the following reports:

- The Postprocessor as well as the Monitor II *Page Data Set Activity* reports provide information about SCM blocks used by the Auxiliary Storage Manager (ASM).
- The Postprocessor *Paging Activity* report provides information about shared and high virtual shared and common frames backed on SCM and also provides information about SCM blocks used by ASM.

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

## z/OS Version 2 Release 1 summary of changes

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

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



---

## Chapter 1. Introducing RMF data gathering and reporting

This document provides you with detailed information about the RMF reports, which are grouped together as follows:

- Interactive Performance Analysis with Monitor III
- Snapshot Reporting with Monitor II
- Real-time Reporting with Monitor I
- Long-term Overview Reporting with the Postprocessor

---

### Gathering data

RMF gathers data using three monitors:

- Short-term data collection with Monitor III
- Snapshot monitoring with Monitor II
- Long-term data gathering with Monitor I and Monitor III

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

#### Short-term data collection with Monitor III

A typical Monitor III gatherer session has a 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 using Monitor III reports. You get actual performance data (response times, execution velocity) on a very detailed level for comparison with goals defined in your service policy.

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 provide this information under the heading **using**.

There is also information about delays, which are important indicators of performance problems.

#### 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

## Introduction

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 is 30 minutes. You can select these options according to your requirements and change them whenever the need arises. Because Monitor I runs in the background and requires little overhead, it can run continuously to provide data for long-term analyses.

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

---

## Long-term overview reporting with the Postprocessor

The Postprocessor offers different types of reports:

*Interval report:* Draws a picture of the sysplex performance for each interval for which data has been gathered.

*Duration report:* The data is summarized over longer periods of time with a maximum value of 100 hours — practically no time limitation.

*Summary report:* Presents an overview of system activity over a specified reporting period.

*Exception report:* Presents a summary of the values that exceeded installation-defined thresholds over a specified period of time.

*Overview report:* This report provides enhanced exception and summary reporting, and offers records for further processing, for example spreadsheet applications on the workstation.

---

## Report analysis with the Spreadsheet Reporter

RMF reports are presented in tabular form, and one very efficient way of handling data in tables is to use a spreadsheet. The Spreadsheet Reporter, a component of RMF that runs on the workstation, converts Postprocessor listings and Overview records into spreadsheets. At your workstation, independent of the systems you are monitoring, you can use one of several familiar spreadsheet applications to manipulate the data as you wish. In addition, the Spreadsheet Reporter provides sample macros to help you in presenting and analyzing performance data at a glance. You find a detailed description in the *z/OS RMF User's Guide*.

Do not hesitate to install and to use this function; you will see that you get a lot of powerful reporting capabilities that help you in running the performance management tasks for your system.

---

## Monitoring on the workstation

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

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

For an introduction to z/OSMF, refer to *z/OS RMF User's Guide*, or for detailed information, refer to *IBM z/OS Management Facility Configuration Guide*.

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

## What you can gather and report

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

### Activity monitoring

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

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

Table 1. Monitored activities and SMF record types

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

## Introduction

Table 1. 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	★	★		★
74.10	★			SCM I/O Activity	★			★
			76	System counters			★	★
74.3/6	★			UNIX	★	★		★
	★		78.2	Virtual storage	★		★	★
	★		72.3	Workload Service classes and report classes	★			★
74.2	★			XCF	★			★
	★			zFS	★			

## Delay monitoring

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

### Delayed address spaces and groups

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

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

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

### Delay reasons for address spaces

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

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

## Long-term performance analysis with RMF XP

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

### Subtype 1-12

AIX on System p performance data

### Subtype 20-31

Linux on System x performance data

### Subtype 40-53

Linux on System z performance data

### Subtype 60-64

Windows on System x performance data

For information on the metric categories provided in the subtypes and how to request the collection of SMF record type 104 from the systems of all or selected supported platforms, refer to the topic *Cross platform monitoring with RMF XP* in *z/OS RMF User's Guide*.

## Reporting of other SMF data

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

### WebServer performance reporting

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

### Lotus Domino support

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





---

## Chapter 2. Interactive performance analysis with Monitor III

This information unit:

- guides you through a Monitor III reporter session
- provides a suggested sequence of reports
- explains how to navigate using cursor-sensitive control
- explains some common Monitor III report measurements
- introduces some Monitor III concepts
- explains how a Monitor III reporter session works
- describes the Monitor III menus
- describes each Monitor III report in detail

---

### Using Monitor III reports

Read the following topics to learn how to use Monitor III reports efficiently:

- “System activities measured”
- “Where to start”
- “Using cursor-sensitive control” on page 11
- “Common Monitor III report measurements” on page 12
- “Monitor III MINTIME and range” on page 15
- “Monitor III report options” on page 17
- “Monitor III sysplex support in different time zones” on page 17

### System activities measured

Monitor III reports can provide delay information for any single job and for any of the following job groups or classes:

- System (all jobs)
- Workload groups and service classes
- TSO
- Batch
- Started tasks
- ASCH
- OMVS
- Enclaves

For each job or group of jobs, Monitor III reports the delay experienced during the report interval and identifies the primary cause of the delay. For any service class period, Monitor III provides a breakdown of response time and displays the information on the **Group Response Time** report.

See chapter “DELAY - Delay Report” on page 63 for more details.

### Where to start

This chapter shows how Monitor III can be used for system monitoring and performance analysis, and helps a new or unexperienced user to find his way through the various RMF reports.

It is necessary to:

## Using Monitor III reports

- Report goal values versus actual values
- Combine data from the entire sysplex to give you an overview at a glance
- Provide accessibility to reports for each system in the sysplex from a single point of control.

Monitor III addresses these needs by:

- Providing sysplex reports
- Arranging the reports in a hierarchy that allows stepping from an overview screen down to address space or resource specific reports.

### Suggested sequence of reports

1. To monitor a sysplex, start with the Sysplex Summary (SYSSUM) report. Use the report options to set the Performance Index threshold to a value of, for example, 0.8 as a warning level, and select a type, so that service class periods are included in the report.

Start the report in GO mode and let it run.

As long as everything is running well, the performance status line at the top of the report shows only green. When the "warning" level for a goal is reached, the corresponding service class with the respective period appears on the report in yellow. And when a goal is not met, the corresponding service class appears on the report in red, followed by the service class period that missed the goal.

2. To find out what is causing the red line, leave GO mode and put the cursor on the line where the goal was not met. If several goals have been missed, the performance index can be of help to find out which goal was "missed most".

Depending on the type of service class, different detailed reports are shown:

- For service classes, a response time breakdown is shown on the Response Time Distribution (SYSRTD) report.
- For subsystem service classes, the transaction states are shown on the Work Manager Delays (SYSWKM) report.

The SYSRTD report has a sysplex view in the upper part of the screen, and provides a single-system breakdown in a scrollable list on the bottom part of the screen.

Furthermore, you can step from the SYSWKM report to the SYSRTD report using cursor-sensitive control, if you need some information from that report for additional investigations, or you want to continue navigation from that report.

3. The scrollable section in the SYSRTD report is the link from the sysplex level to the single system. From these lines it is possible to "zoom" into any of the listed systems. Placing the cursor on the system-ID in one of the rows and pressing ENTER, leads to the SYSINFO report of that system, thus allowing further analysis based on the data shown there. Placing the cursor on a specific data column in one of the rows of the scrollable area leads to a specific report of that system that provides additional information related to the selected column.
4. Finally, when the single-system level is reached, navigation among those reports is possible as described in step 5.
5. In the workflow/Exceptions (WFEX) report, you can identify jobs and resources with low workflow values or jobs that have met exceptional conditions. For example, you can check the Reason field to identify the user or the possible cause of delay. Once you recognize a user or a resource with a potential problem, you can analyze the situation using cursor-sensitive control. "Using cursor-sensitive control" on page 11 describes how to invoke reports using this method.

If you are on the Delay report, check the delay value (for PROC, STOR, DEV, SUBS, OPER, ENQ) with the largest value associated with a job, use cursor-sensitive control to navigate to the Job Delay report for that type of delay to analyze the main reason for the delay.

In case of a delay due to devices (DEV) or enqueued resources (ENQ), you can use cursor-sensitive control to further investigate a problem by looking at the resource-oriented device report (DEVVR) and the resource-oriented enqueue report (ENQR). For storage problems involving paging or swapping delays, you can use the resource-oriented storage delay report (STORR). Use either the job entry subsystem (JES), hierarchical storage manager (HSM), or cross-system coupling facility (XCF) delays report for a delay associated with SUBS. For OPER delay, use cursor-sensitive control to see the appropriate Job Delay report.

## Using Monitor III reports

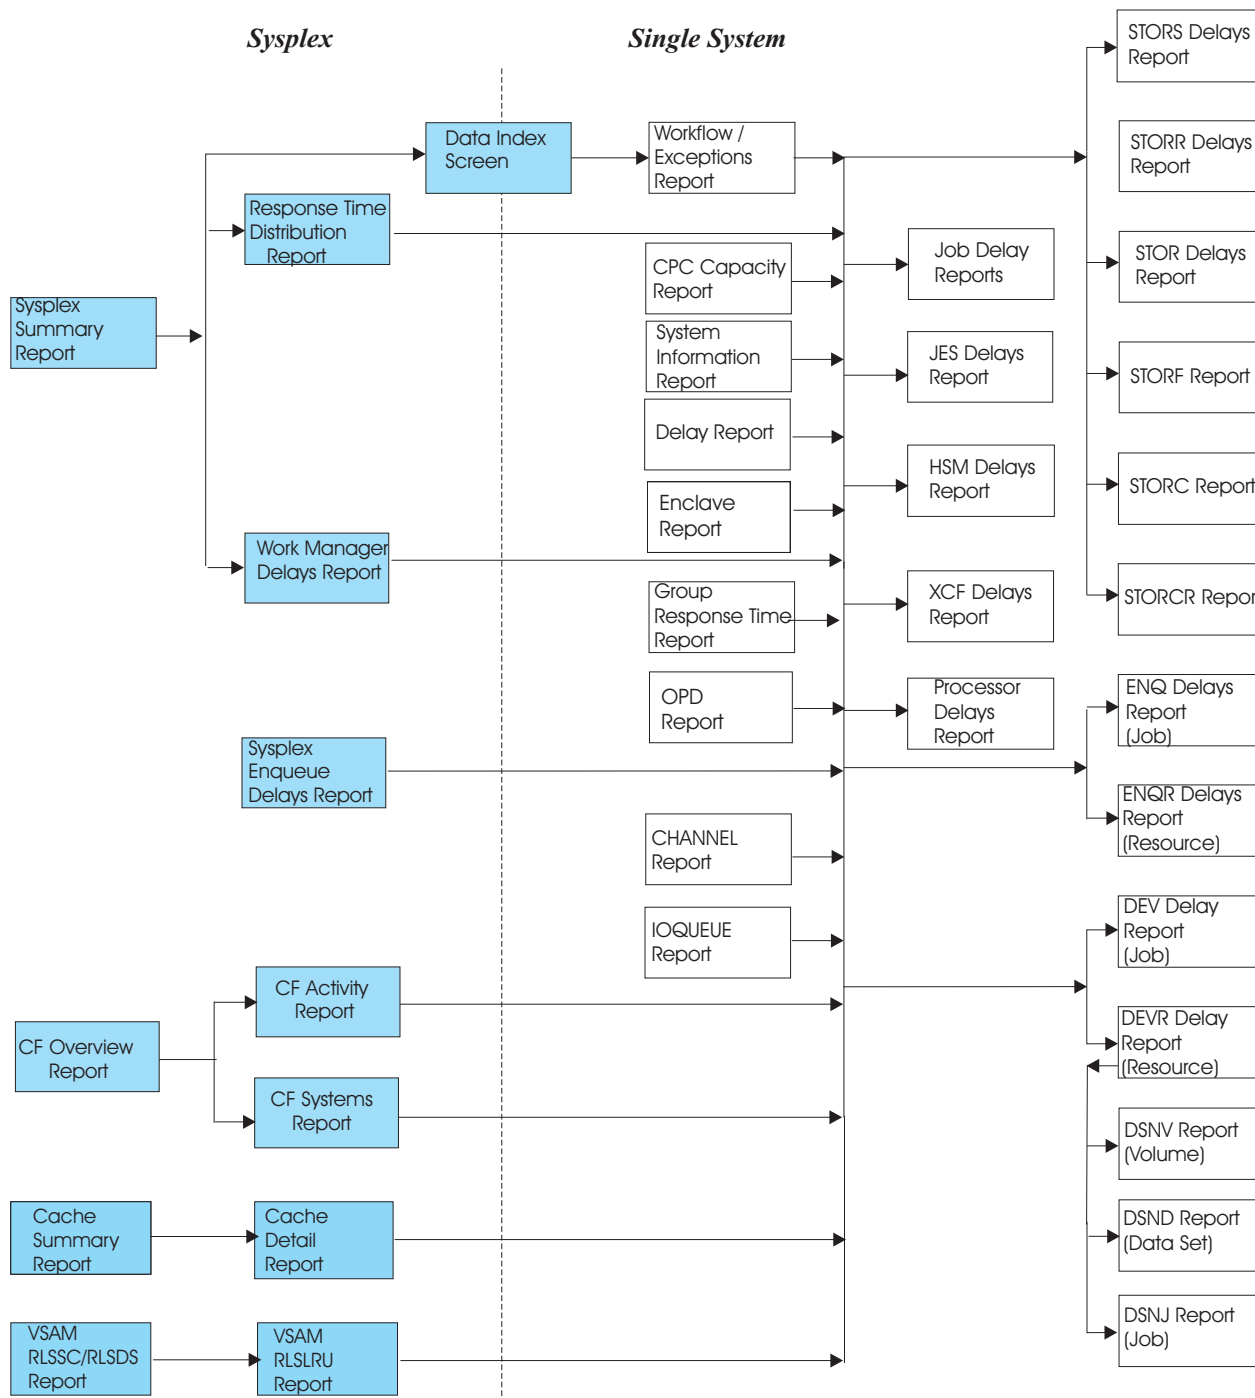


Figure 1. Suggested Sequence for Using Monitor III Reports

For a summary of common system storage consumption, use the Common Storage Summary report (STORC). To identify remaining storage, use the STORCR report.

For a summary of how the workflow and delay of the measured system affects performance, use the SYSINFO report.

Figure 1 shows a suggested sequence for using Monitor III reports to resolve potential problems.

## Getting information about data to be reported

For special purposes, where an overview of the available data may be helpful, you can also use the following path:

- You place the cursor on the sysplex field in the header line of a sysplex report and press ENTER.
- You call the DI command.

This leads to the Data Index screen.

The DI report shows all VSAM data sets used during data set recording, including data from other systems, or all preallocated data sets for the display session. From the DI report, you can decide what data in the data sets you want to display.

With the *Include data set names* option set to *NO*, this window gives an overview of all the data in the system.

Switching the *DDNAMES/DSNAMES* option to *YES* gives a Data Index screen with all data set names.

Based on the available data sets of the respective system, you can continue the analysis for that system.

## Using cursor-sensitive control

Cursor-sensitive control lets you place the cursor on a field in a tabular report, and press the ENTER key, to see another report containing additional information about the field where the cursor is positioned. You can easily navigate among the RMF reports without returning to the Primary Menu or entering specific commands.

For example, you can move from the ENQ Delays report to the ENQ Resource Delays report by using cursor-sensitive control on the **Major/Minor Names** field. Note that the result of using cursor sensitivity depends on the data. For example, if you use cursor-sensitive control on the **Primary Reason** field in the Delays report, the resulting report is that variation of the Job Delays report that is related to the main reason for the delay.

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

**Note:** If you press the RETURN (PF4) key, RMF displays the Primary menu and you lose all return paths.

If the path extends over reports that are built from different systems, the return path is lost.

If you issue any RMF command while using cursor-sensitive control, RMF will erase the return path.

Cursor-sensitive control is active on:

- most fields on all tabular reports except STORCR
- the *Jobname* field of the **Job Report Selection Menu**
- the *Report* field of the **Option Selection Menu**.

### Common Monitor III report measurements

Most values included in Monitor III session reports are similar in their calculation. The following definitions and general formulas are common to all RMF reports:

- Using (%) for address spaces
- Delay (%) for address spaces
- Workflow (%) for address spaces and resources
- Execution velocity

#### Using samples

**PROC** The number of address spaces found using one or more processors (which can be standard CPs (aka general purpose processors) or special purpose processors). An address space is considered using one or more processors when it has ready work (any ready SRB, interrupted ready task, asynchronous exit routine, or TCB is on the dispatching queue) that could be dispatched by the processor on which the Monitor III data gatherer is running.

**DEV** The number of address spaces found using one or more devices. An address space is considered using one or more devices when it issues an I/O request. However, because the channel subsystem accepts an I/O request whether the device, control unit or both are busy or not, the requests might or might not be delayed (queued) in the channel. Therefore, the using requestors for devices might also contain an unknown amount of delay. You must consider this delay when interpreting the workflow value.

#### Delay samples

**PROC** The number of address spaces found waiting for a processor (which can be general purpose or special purpose processors). An address space is considered waiting for a processor when the address space has at least one ready unit of work that is not dispatched. Primary source fields referenced in this calculation are the same as those listed under PROC for using samples.

**DEV** The number of address spaces found waiting for a measured device. An address space is considered to be waiting for a measured device when at least one I/O queue element in the I/O queue for the device identifies the address space as the issuer of the I/O request but the request is not active. I/O requests queued in the channel for devices are considered to be using the device, and therefore an unknown amount of delay is missing from the delayed requestor count for devices.

**ENQ** The number of address spaces found waiting for serially reusable resources.

**HSM** The number of address spaces found waiting for an HSM service.

**JES** The number of address spaces found waiting for a JES service.

**OPR** The number of address spaces found waiting for operator interventions.

**STR** The number of address spaces found waiting for storage operations.

**XCF** The number of address spaces found waiting for an XCF path.

#### Address space workflow (%)

The workflow of an address space represents how a job uses system resources and the speed at which the job moves through the system in relation to the maximum

average speed at which the job could move through the system. The speed at which the system performs the work of one job depends on the simultaneous work requested by other jobs.

A value from 0% to 100% indicates the workflow within the report interval. A low workflow value indicates that a job has few of the resources it needs and is contending with other jobs for system resources. A high workflow value indicates that a job has all the resources it needs to execute, and that it is moving through the system at a relatively high speed.

For example, a job that would take four minutes to execute if all the resources it needed were available, would have a workflow of 25% if it took sixteen minutes to execute.

The following formula defines the workflow of a *single* address space:

### Single Address Space

$$\text{Workflow (\%)} = \frac{\text{\# Using Samples}}{\text{\# Using Samples} + \text{\# Delay Samples}} * 100$$

**Note:** In calculating Workflow, Monitor III counts an address space as using a resource if at least one of its ready tasks is using the resource. Even if the address space has other ready tasks delayed for the same resource, Monitor III counts the address space as using the resource (single state case). For example, if a job has four ready tasks in its address space, and one task is using the processor while three tasks are simultaneously delayed for the processor, Monitor III considers this address space to have a using count of one and a delay count of one.

Also remember that a job can be using one resource and delayed for another at the same sample, or delayed for more than one resource at a time, or using more than one resource. The maximum per sample is two using (PROC and DEV) and eight delays (one for each resource).

### Example

A job was found to be delayed or productive 75 times. The job was found to be using the processor 5 times and a device 10 times. The job was also found delayed for the processor 15 times, for a device 20 times and for an enqueued resource 25 times. The Workflow (%) of the job would be:

$$\text{Workflow (\%)} = \frac{5 + 10}{(5 + 10) + (15 + 20 + 25)} * 100 = 20\%$$

The following formula defines the workflow of a *group* of address spaces:

### Group of Address Spaces

$$\text{Workflow (\%)} = \frac{\sum \text{Using Samples}}{\sum \text{Using Samples} + \sum \text{Delay Samples}} * 100$$

**Note:** The sums represent the values for all address spaces in the group.

### Resource workflow (%)

The workflow of resources indicates how efficiently users are being served. The speed with which each resource performs the work of all users is expressed as a value from 0% to 100%.

## Measurements

A low workflow value represents a large queue of work requests and a large number of delayed jobs, while a high workflow value represents little resource queuing contention and a small number of delayed jobs.

The following formula defines the workflow of a resource (DEV or PROC):

### Resource

$$\text{Workflow (\%)} = \frac{\# \text{ Using Samples}}{\# \text{ Using Samples} + \# \text{ Delay Samples}} * 100$$

### Address space using (%)

Jobs getting service from hardware resources (PROC or DEV) are *using* these resources. The use of a certain resource by an address space can vary from 0% to 100%, where 0% indicates no use of the resource during the report interval and 100% indicates that the address space was found using the resource in every sample during that period. If you use the default range of 100 seconds, 1% of using is equal to 1 second of using to the user.

The following formula defines the use of a resource by an address space during the report interval:

### Single Address Space

$$\text{Using (\%)} = \frac{\text{Using Samples}}{\# \text{ Samples}} * 100$$

**Note:** In calculating Using, Monitor III counts an address space as using a resource even if the address space is also delayed for the identical resource (single state case). For example, if a job has four ready tasks in its address space, and one task is using the processor while three tasks are simultaneously delayed for the processor, Monitor III considers this address space to have a Using count of one and a Delay count of one.

PROC and DEV using can add up to more than the overall using percentage, with the maximum being 200 %.

The using state of a group of address spaces for a certain resource during a report interval can also range from 0% to 100% and is calculated as follows:

### Group of Address Spaces

$$\text{Using (\%)} = \frac{\sum \text{Using Samples}}{\# \text{ Samples} * \text{Avg} \# \text{ Address Spaces}} * 100$$

### Address space 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. The delay of an address space for a specific resource or for all resources can vary from 0% to 100%. A delay of 0% indicates no delay during the report interval, while a delay of 100% represents a job that was found delayed at every sample during that period. Delay is a percent of Time during the period; with the default Range of 100 seconds, 1% delay is equal to one second of delay to the user.

The following formula defines the delay of an address space for a certain resource during a report interval:



### Single Address Space

$$\text{Delay (\%)} = \frac{\# \text{ Delay Samples}}{\# \text{ Samples}} * 100$$

**Note:** In calculating Delay, Monitor III counts an address space as delayed for a resource if at least one ready user (unit of work) is waiting for a device or processor. In the case of single state sampling, if a job has more than one ready tasks simultaneously delayed for the processor, Monitor III considers this address space to have a delay count of one.

The sum of individual delays can be more than overall delay, with a maximum of 600%.

The delay of a group of address spaces for a certain resource during a report interval can also range from 0% to 100% and is calculated as follows:

### Group of Address Spaces

$$\text{Delay (\%)} = \frac{\sum \text{ Delay Samples}}{\# \text{ Samples} * \text{ Avg } \# \text{ Address Spaces}} * 100$$

**Note:** This value needs to be checked carefully if the number of address spaces in the group is very small.

### Execution velocity

The execution velocity is a measure of how fast work is running compared to ideal conditions without delays.

The calculation of the execution velocity is:

### Execution Velocity

$$\text{Execution Velocity (\%)} = \frac{\# \text{ Using samples}}{\# \text{ Using Samples} + \# \text{ Delay Samples}} * 100$$

The values are taken from RCAETOTU and RCAETOTD which are described in the IWMWRCAA mapping (see *z/OS MVS Planning: Workload Management*).

## Monitor III MINTIME and range

The Monitor III data gatherer combines all samples gathered into a set of samples for a time interval called MINTIME. The value for MINTIME is specified as gatherer option. The recommended value is 100 seconds.

Reporting is performed based on this MINTIME interval and is defined by the Range value. Range can be set either on the Session Options dialog or directly in each report header line on the Report Options panel, or using the BREF/FREF command.

When choosing a range for your report interval, there are two things to consider:

- It must be a multiple of the MINTIME that the data was gathered for
- It can be defined in seconds or minutes:

**nnnnS** where nnnn represents a number from 0 to 9999

**nnnM** where nnn represents a number from 1 to 166

## MINTIME and range

**Note:** If you specify a value that is less than the MINTIME, the default value will be changed to equal the MINTIME.

### How the data gathered affects the data reported

If you request a report interval on a report heading that crosses two MINTIMES, data will be presented for both MINTIMES and Time and Range will be adjusted accordingly on the report header.

### Example

If the data gatherer runs with the recommended MINTIME of 100 seconds, data is gathered in the following intervals:

```
Time = 12:00:00   :01:40   :03:20   :05:00   :06:40
      |-----|-----|-----|-----|
```

The report that results from this data will have an initial time of 12:00:00 and a range of 100 seconds to match the data gathered.

If you revise the time to start at 12:02:00 and leave the range unchanged, the report heading changes to reflect a time of 12:01:40 and a range of 200. Here's why:

### Requested Report

Interval: 12:02:00 to 12:03:40

Range: 100 Sec

```
Time = 12:00:00   :01:40   :03:20   :05:00   :06:40
      |-----|-----|-----|-----|
                ↑         ↑
                |-----|
```

### Presented Report

Interval: 12:01:40 to 12:05:00

Range: 200 Sec

```
Time = 12:00:00   :01:40   :03:20   :05:00   :06:40
      |-----x-----x-----|
                ↑         ↑
                |-----|
```

Rather than present less data than you requested, RMF displays a report using the minimum number of MINTIMES that include the interval you requested. In the above example, this means the report interval must start at 12:01:40 and finish at 12:05:00.

To accurately reflect the data presented on the report, the header is adjusted accordingly. Time is changed to 12:01:40 (the start of the first MINTIME in the report interval) and Range is adjusted to 200 seconds (to include the last MINTIME in the report interval).

**Note:** The Range value that you specify on the Session Options panel is saved in your current option set and applies to all reports displayed when that option set is in effect. If you modify Range either directly in each report header line, or using the BREF/FREF command, the new range temporarily overrides the value on the Session Options panel, but is not saved in your current option set.

### Shortened intervals

The following events can cause a shortened report interval:

**A policy switch**

A report interval containing a policy switch can have data with different gatherer options.

**A system IPL****A change of the gatherer CYCLE time**

RMF cannot combine data that was collected using different gathering options, and so the reporting range will be adjusted to start where the change occurred.

**Example**

If the data gatherer runs with the recommended MINTIME of 100 seconds, and a mode switch occurred at 12:02:00, data will be gathered in the following intervals:

```
Time = 12:00:00   :01:40   :03:20   :05:00   :06:40
          |-----|-x
          |
          |         ↑ (Switch time)
          |-----|-----|-----|
          :02:00   :03:40   :05:20   :07:00
```

**Monitor III report options**

Most of the Monitor III reports can be tailored by using specific report options. You can either define these options by parameters together with the report command, or you can call the Report Options panel with the command ROPTIONS or RO. Depending on the specific report, you will get a panel where you can select and specify valid options. By pressing PF3, you leave the panel and activate the options which will stay valid until they will be changed explicitly. In addition, you can use the following commands in a Report Options panel:

**CANCEL**

You can use this command to exit the panel without making any changes.

**RESET**

To ensure the RMF default settings for option panels are in effect, enter RESET on the command line or the respective panel. RESET reestablishes RMF's default settings. Because there is no default value for jobname, the RESET command is not valid on the options panel of the Job Delays report.

**Monitor III sysplex support in different time zones**

This chapter describes how Monitor III is reporting a sysplex with systems running in different time zones.

Monitor III will always work with the local time of the system the reporter session has been started on. When requesting data from a system with a different local time, Monitor III will automatically adapt the begin and end time of the data request. Users do not have to care about different time zones, the Monitor III reporter does it internally.

**Example**

```
SYSA (local time is GMT+2)
SYSB (local time is GMT-1)
```

Monitor III reporter session started on SYSA

## MINTIME and range

- Sysplex report requested for data at 10.05 am (GMT 08.05 am)
  - SYSA returns data from 10.05 am local SYSA time
  - SYSB returns data from 07.05 am local SYSB time

Thus, both systems contribute data of the same point of time based on GMT.

- Single system report requested for remote system SYSB 10.05 am
  - SYSB returns data from 07.05 am local SYSB time

This matches 10.05 am local SYSA time. The reporter displays it with Time=10.05.00.

### Note:

1. The Monitor III Data Index screen displays in the columns Begin/End Date/Time still the local time from the system.
2. This processing is done only when running a Monitor III reporter session to display data gathered by an active Monitor III gatherer session, not when displaying data from VSAM data sets by preallocating DDNAMEs RMFDSnn before starting a reporter session.

---

## The reporter session

A Monitor III reporter session runs under ISPF and issues online reports about system performance. During a reporter session, you can obtain reports with current data from any system in the sysplex, with data collected earlier, or with preallocated data from any system. You can use the Data Index to choose what data you want to use. See the *z/OS RMF User's Guide* for information.

The tutorial, which is option T on the Monitor III Primary Menu, is an online introduction to Monitor III. It gives you an overview of an RMF reporter session and contains some examples on using RMF. You can also use the HELP (PF1) key for additional information on a panel, message, or report you are viewing.

## Starting and stopping a Monitor III reporter session

You can start the reporter session with the command:

```
RMF
```

This presents the RMF Primary Menu, and by selection 3, you get the Monitor III reporter session.

To end the RMF Monitor III data reporter, select X on the Primary Menu, or enter =X

on any command line.

### Messages during Monitor III start

There are two special cases that you might see a message at your terminal after calling Monitor III:

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

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

- IEC130I ADMPC DD STATEMENT MISSING

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

## The Primary Menu

After the Monitor III reporter session starts, RMF displays the Monitor III Primary Menu.

```

RMF Monitor III Primary Menu                                z/OS V2R2 RMF
Selection ==>
Enter selection number or command on selection line.

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

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

                    0 OPTIONS    T TUTORIAL    X EXIT

                    5650-ZOS Copyright IBM Corp. 1986, 2016.
                    Licensed Materials - Property of IBM

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

```

Figure 2. Monitor III Primary Menu

You can use the ISPF options to tailor the command and selection lines.

RMF reports and their fields are described later in this chapter. See the *z/OS RMF User's Guide* for information on RMF options.

## Selecting a report

You can select a report from the Primary Menu or from any other report panel in the following ways:

- Enter the report name or its abbreviation on any selection or command line.
- Enter the selection number of the report name on the selection line of the Primary Menu.
- Use the ISPF 'jump' function to enter the selection number on any selection or command line.

**Note:** While viewing tabular reports, you can use cursor-sensitive control to obtain additional detailed reports about several individual fields.

### Example

To invoke the Delay Report, enter the command DELAY on any selection or command line within Monitor III: Command ==> delay

Or enter a 1 on the Primary Menu: Selection ==> 1

and then a 4 on the Overview Report menu: Selection ==> 4

## Monitor III - Primary Menu

Or select the report from any other report panel by using the ISPF *jump function*:  
Command ==> =1.4

If you do not specify parameters for report commands (jobname, class, selection, resource), RMF defaults to the options already in effect for the session. If nothing was specified on the Job Delays report, RMF presents you with a report options panel so that you can specify a selection.

To change options for a report, enter the following on the command line of the report:

Command ==> ROPTIONS

RMF displays the Report Options panel for that report. On this panel, you can modify what is presented on RMF reports by changing the options. RMF saves any report options you change across sessions. When you are finished specifying the options, press END (PF3) to save your changes and return to the report.

**Note:** You can also use the option selection (OPTIONS) menu to access a Report Options panel.

## Monitor III report commands

Table 2 on page 21 lists all report commands with their parameters and abbreviations. The “How to request this report” section for each report shows an example of the command and parameters.

You can enter the commands on any command line.

The **Parameters** column in Table 2 on page 21 indicates what parameters, if any, you can specify on the respective commands:

### **cfname**

A coupling facility name

### **job\_class**

One of the following names of a job class:

ALL (A)  
ASCH (AS)  
BATCH (B)  
OMVS (O)  
STC (S)  
TSO (T)

### **Notes:**

1. This parameter is optional. If it is not specified, ALL is used by default.
2. In addition, ENC (or E) can be specified as class for the DELAY report.

### **dsname**

A data set name

### **jobname**

A job name

### **period**

A service or report class period

**resource**

A resource name

**service\_class**

A service class name

**s/r-class**

A service or report class name

**ssid**

A cache subsystem identifier

**sstype**

The name of a subsystem that schedules enclaves

**storage\_class**

A storage class name

**volser**

A serial number of a volume

**wlm**

The name of a workload group, a service class, or a report class

Table 2. Report Commands

Command	Parameters	Displays	Abbreviation
CACHDET	ssid	Cache detail report	CAD
CACHSUM		Cache summary report	CAS
CFACT	cfname	coupling facility activity report	CA
CFOVER	cfname	coupling facility overview report	CO
CFSYS	cfname	coupling facility system report	CS
CHANNEL		Channel path activity report	CHAN, CH
CPC		CPC capacity report	
DELAY	job_class, service_class	Delays report for all jobs or specified job groups	DEL, DLY, DL
DELAYJ	jobname	Job report variation for specified job reflecting primary delay reason	DLJ, DJ, DELJ, DLYJ, JOB, JO
DEV	job_class, service_class	Device delays report for all jobs or specified job groups	DD, DVD
DEVJ	jobname	Device delays variation of job report for specified jobname	DDJ, DVJ
DEVR	volser	Device delays report for all or specified resources	DR, DVR
DSINDEX		Data index information	DS, DI
DSND	dsname	Data set delays report for all or specified data sets	DSN
DSNJ	jobname	Data set delays - Job report for specified jobname	DSJ
DSNV	volser	Data set delays - Volume report for specified volume	DSV
ENCLAVE	sstype	Enclave activity report	ENCL
ENQ	job_class, service_class	Enqueue delays report for all jobs or specified job groups	ED
ENQJ	jobname	Enqueue delays variation of job report for specified jobname	EJ
ENQR	resource	Enqueue delays for all or specified resources	ER
GROUP	s/r-class, period	Group response time breakdown	GP, GRP, GD, RT, GRT

## Monitor III - Primary Menu

Table 2. Report Commands (continued)

Command	Parameters	Displays	Abbreviation
HSM	job_class, service_class	HSM delays report for all jobs or specified job groups	HD
HSMJ	jobname	HSM delays variation of job report for specified jobname	HJ
IOQUEUE		I/O queuing activity report	IOQ, IQ
JES	job_class, service_class	JES delays report for all jobs or specified job groups	JD
JESJ	jobname	JES delays variation of job report for specified jobname	JJ
JOB	jobname	Job report variation for specified job reflecting primary delay reason	JO, DELAYJ, DLYJ, DELJ, DLJ, DJ
LOCKSP	HELD   SPIN   <u>BOTH</u>	Spin Lock Report about held spin locks and/or address spaces spinning due to a request for a spin lock	LSP
LOCKSU	LOCAL   GLOBAL   <u>BOTH</u>	Suspend Lock Report about local and/or global suspend locks	LSU
MNTJ	jobname	Operator delays variation for mount request of job report for specified jobname	MTJ
MSGJ	jobname	Operator delays variation for message request of job report for specified jobname	MSJ
OPD		OMVS process data	
PCIE		PCIE activity	PCI
PROC	job_class, service_class	Processor delays report for all jobs or specified job groups	PD
PROJ	jobname	Processor delays variation of job report for specified job	PJ
PROCU	job_class, service_class	Processor usage of a job per processor type (standard or special purpose processors)	PU
QSCJ	jobname	Operator delays variation for quiesce command of job report for specified jobname	QJ
RLSDS	dsname	VSAM RLS activity by data set	RLD
RLSLRU		VSAM LRU overview	RLL
RLSSC	storage_class	VSAM RLS activity by storage class	RLS
SCM		SCM activity	
SPACED		Disk space report	SPD
SPACEG		Storage space report	SPG
STOR	job_class, service_class	Storage delays report for all jobs or specified job group	SD
STORC	job_class, service_class	Common storage report	SC
STORCR		Common storage remaining at end of job report	SCR
STORF	job_class, service_class	Detailed information on frame counts for all jobs or specified job group	SF
STORM	job_class, service_class	Detailed information about the use of memory objects within the system	SM
STORJ	jobname	Storage delays variation of job report for specified job	SJ
STORR		Storage space and paging activity report for all system volumes	SR
STORS	wlm	Summarized storage information by workload group, service or report class	SS
SYSENG		Sysplex enqueue delays report	ES



Table 2. Report Commands (continued)

Command	Parameters	Displays	Abbreviation
SYSINFO	wlm	System information, total and by user groups	SY, SYS, SI
SYSRTD	s/r-class, period	Response time distribution report	RTD
SYSUM	wlm	Sysplex summary	SUM
SYSWKM	s/r-class, period	Work manager delays report for subsystems	WKM
USAGE	job_class, service_class	Job usage report	USG
WFEX		Workflow/exceptions screen	WE, WF
XCF	job_class, service_class	Cross-system coupling facility delays report	XD
XCFJ	jobname	XCF delays variation of the job report for specified jobname	XJ
ZFSFS		zFS file system	ZFF
ZFSKN		zFS file system kernel	ZFK
ZFSOVW		zFS file system overview	ZFO

Table 3 contains commands for the examples of user-written reports that were delivered with RMF.

Table 3. User-Written Report Commands

Command	Parameters	Displays	Abbreviation
DEVN		Device activity	DA
DEVT		Device trend	DT
DSD		Detailed storage delays	
RG		Resource group data	
SYSTREND		System trend	ST

## Header for single-system reports

Figure 3 shows the common header for single-system Monitor III reports.

Command ==>	RMF V2R2 TITLE	Line 1 of 30
		Scroll ==> HALF
Samples: nnn	System: syst	Date: mm/dd/yy Time: hh.mm.ss Range: nnn Sec

Figure 3. Header of Monitor III Single-System Reports

All Monitor III single-system report headers contain the following information:

Table 4. Monitor III Heading Information

Heading	Definition
Report title	The type of measurement data
Date/Time	The starting date and time for the first set of samples included in the report.
Range	The length of time (in seconds) during which samples were gathered, starting with the time specified in the Time field.

## Header for sysplex reports

The sysplex report header differs from single-system reports in the following fields:

## Monitor III - Primary Menu

1. The sysplex reports show the sysplex name, whereas the single-system reports show the SMF system identification. In Figure 4, **SYSPLEXN** is the eight character sysplex name.
2. For some sysplex reports, for example, SYSSUM or SYSWKM, the RMF sample count is replaced by the number of WLM samples. This value is an average of the MINTIMEs from the different systems, that contributed to the report.
3. The number of systems participating in the report is shown and indicates whether the complete sysplex is contributing to the report or not. To find out which specific system is not included, check the HELP-Screen of the message *Not all systems included in the report*. Furthermore, the Data Index may be useful.

The sysplex report header is shown in Figure 4.

```
Command ==>>          RMF V2R2  TITLE - SYSPLEXN          Line 1 of 30
                               Scroll ==>> HALF

WLM Samples:   nnn   Systems: n Date: mm/dd/yy Time: hh.mm.ss Range: nnn   Sec
```

Figure 4. Header of Monitor III Sysplex Reports

## The Sysplex Report Selection Menu

```
RMF Sysplex Report Selection Menu
Selection ==>>

Enter selection number or command for desired report.

Sysplex Reports
  1 SYSSUM  Sysplex performance summary          (SUM)
  2 SYSRTD  Response time distribution           (RTD)
  3 SYSWKM  Work Manager delays                 (WKM)
  4 SYSENG  Sysplex-wide Enqueue delays        (ES)
  5 CFOVER  Coupling Facility overview         (CO)
  6 CFSYS   Coupling Facility systems          (CS)
  7 CFACT   Coupling Facility activity         (CA)
  8 CACHSUM Cache summary                      (CAS)
  9 CACHDET Cache detail                      (CAD)
 10 RLSSC   VSAM RLS activity by storage class (RLS)
 11 RLSDS   VSAM RLS activity by data set      (RLD)
 12 RLSLRU  VSAM LRU overview                 (RLL)
 13 ZFSOVW  zFS Overview                      (ZFO)
 14 ZFSFS   zFS File System                   (ZFF)
 15 ZFSKN   zFS Kernel                        (ZFK)

Data Index
  D DSINDEX Data index                       (DI)
```

Figure 5. Monitor III Sysplex Report Selection Menu

You can navigate to the *Sysplex Report Selection Menu* by selecting an **S** on the Primary Menu.

Use this menu to select one of the sysplex reports, or the Data Index. For more information about the Data Index, see “The Data Index” on page 28.

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

### The Overview Report Selection Menu

```

RMF Overview Report Selection Menu
Selection ==>
Enter selection number or command for desired report.

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

Detail Reports
  4 DELAY  Delays                  (DLY)
  4A USAGE Job Usage              (USG)
  5 GROUP  Group response time breakdown (RT)
  6 ENCLAVE Enclave resource consumption and delays (ENCL)
  7 OPD    OMVS process data
  10 SPACEG Storage space          (SPG)
  11 SPACED Disk space             (SPD)
  12 LOCKSP Spin locks             (LSP)
  13 LOCKSU Suspend locks         (LSU)
  
```

Figure 6. Monitor III Overview Report Selection Menu

You can navigate to the *Overview Report Selection Menu* by selecting a **1** on the Monitor III Primary Menu.

On the Overview Report Selection Menu, you can select among various basic and detail reports.

### The Job Report Selection Menu

To request the *Job Report Selection Menu*, select **2** from the Primary Menu or enter JOBS on any command line. Use this menu to choose the specific job and the type of delay which you want to analyze.

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

## Mon III - Job Report Selection Menu

```
RMF Job Report Selection Menu
Selection ==>

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

Jobname ==> _____

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

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

Figure 7. Monitor III Job Report Selection Menu

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

---

### The Resource Report Selection Menu

Use this menu to choose what resource you want to see delays or storage problems for.

To request the *Resource Report Selection Menu*, select 3 on the Primary Menu, or enter RESOURCE on any command line.

```

RMF Resource Report Selection Menu
Selection ==>

Enter selection number or command for desired report.

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

Figure 8. Monitor III Resource Report Selection Menu

The Storage report section of the menu allows you to choose one of the six types of storage report provided by RMF. There are two types of report: Storage reports and Common Storage reports.

## The Subsystem Report Selection Menu

The Subsystem report menu allows you to select HSM, JES, and XCF Delay reports.

To request the *Subsystem Report Selection Menu*, select **4** from the Primary Menu or enter SUBS on any command line.

```

RMF Subsystem Report Selection Menu
Selection ==>

Enter selection number or command for desired subsystem report.

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

Figure 9. Monitor III Subsystem Report Selection Menu

## The User Report Selection Menu

The User report menu allows you to select your user-written reports or those examples that are provided with Monitor III.

To request the *User Report Selection Menu*, select **U** from the Primary Menu or enter USER on any command line.

```
RMF User-written Report Selection Menu
Selection ==>

Enter selection number or command for desired report.

  1 MSI           Migration SYSINFO including Execution Velocity
  2 DSD           Detailed Storage Delays
  3 RG            Resource Group Data

Device Reports
DA DEVN          Device Activity
DT DEVT          Device Trend
                 Device   => _____

System Reports
ST SYSTREND      System and Workload Trend
                 Workload => _____
```

Figure 10. Monitor III User-written Report Selection Menu

## The Data Index

The Data Index (DI) shows you the data sets that are available throughout the sysplex.

The Data Index provides information about the data that is currently available for your reporter session. The data that it contains is either:

- Current data from all active Monitor III data gatherers in the sysplex. The current data represents all available data from every system that can be found in the sysplex.
- Previously stored data from a Monitor III gatherer session (so called preallocated data sets).

The Data Index displays the list of systems in alphabetical order by system ID of the RMF Monitor III data gatherer that recorded the data. You can also see if data is missing, or could not be retrieved due to one of the following reasons:

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

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

## How to request the Data Index

Select **S** on the Primary menu, and then **D** on the Sysplex Report menu, or you can enter the following command:

```
DI
```

## Contents of the Data Index

Read the following information about the Data Index:

- “If you are using active Monitor III gatherers in the sysplex” on page 29
- “If you are using preallocated data sets” on page 29

- “Condensed information on the Data Index” on page 30
- “Data Index — field descriptions” on page 31
- “Information and error messages on the Data Index” on page 32
- “Cursor-sensitive control” on page 33
- “Data Index options” on page 33

**If you are using active Monitor III gatherers in the sysplex**

For each active Monitor III gatherer in the sysplex, the following will be displayed on the Data Index:

- The RMF in-storage buffer; and
- If the Monitor III gatherer has been started with data set support, all data sets used by the Monitor III gatherer for recording data.

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

Figure 11 shows a sample Data Index that is using data from active Monitor III data gatherers in a sysplex.

```

RMF V2R2 Data Index - RMFPLEX1 Line 1 of 22
Command ==> Scroll ==> HALF
Samples: 118 System: MVS2 Date: 11/27/15 Time: 10.12.00 Range: 120 Sec
----Begin/End----
System --Date-- --Time-- -DDNAME- -----Data Set Name-----
MVS1 11/27/15 10.03.20
10.12.00 * * * In-storage buffer * * *
MVS1 SYS00002 RMF.MONITOR3.DATASET1.MVS3
* * * Data from system MVS3 * * *
MVS1 SYS00001 RMF.MONITOR3.DATASET2.MVS3
* * * Data from system MVS3 * * *

MVS2 11/27/15 09.11.00 SYS00002 RMF.MONITOR3.DATASET1.MVS2
09.14.00
MVS2 11/27/15 10.03.00 SYS00003 RMF.MONITOR3.DATASET2.MVS2
10.12.00 * * * Currently active * * *
MVS2 11/27/15 10.03.00
10.12.00 * * * In-storage buffer * * *

MVS3 11/27/15 09.11.00 SYS00002 RMF.MONITOR3.DATASET1.MVS3
09.14.00
MVS3 11/27/15 10.03.00 SYS00003 RMF.MONITOR3.DATASET2.MVS3
10.12.00 * * * Currently active * * *
MVS3 11/27/15 10.03.00
10.12.00 * * * In-storage buffer * * *
TEST
* * * No response * * *

```

Figure 11. Data Index

**If you are using preallocated data sets**

If a local session has had data sets preallocated to it before the RMF reporter is started, the Data Index will only display data from those data sets. This is independent of the active Monitor III gatherers in the sysplex. For information about how to preallocate data sets to the local session, refer to the *z/OS RMF User's Guide*.

## Mon III - Data Index

All rows will be shown in dark blue.

**Note:** It is possible to preallocate data sets from different systems, but only one sysplex can be represented by the data in those data sets. If the data represents more than one sysplex, the Data Index is displayed, and no other report can be generated. To resolve the problem, end the session, deallocate any data sets with a different sysplex ID, and start a new session.

Figure 12 shows a sample Data Index that is using data from preallocated data sets to a reporter session.

```
RMF V2R2 Data Index - RMFPLEX Line 1 of 18
Command ==> Scroll ==> HALF
Samples: 37 System: RMFB Date: 09/28/16 Time: 13.09.00 Range: 60 Sec
----Begin/End----
System --Date-- --Time-- -DDNAME- -----Data Set Name-----
RMFA 09/28/16 04.17.00 RMFDS01 RMF.RMFA.RMFDS01
          06.26.00
RMFA 09/28/16 04.20.00 RMFDS00 RMF.RMFA.RMFDS00
          06.28.00
RMFA 09/28/16 11.28.00 RMFDS02 RMF.RMFA.RMFDS02
          13.07.00
RMFB 09/28/16 04.16.00 RMFDS04 RMF.RMFB.RMFDS01
          06.24.00
RMFB 09/28/16 04.26.00 RMFDS05 RMF.RMFB.RMFDS00
          06.36.00
RMFB 09/28/16 11.58.00 RMFDS03 RMF.RMFB.RMFDS02
          13.10.00
RMFC 09/28/16 04.17.00 RMFDS07 RMF.RMFC.RMFDS01
          06.26.00
RMFC 09/28/16 04.27.00 RMFDS06 RMF.RMFC.RMFDS00
          06.37.00
RMFC                RMFDS07 RMF.RMFC.RMFDS02
          * * * Empty * * *
```

Figure 12. Data Index with Preallocated Data Sets - Detailed View

### Condensed information on the Data Index

The detailed version of the Data Index allows you to display all data sets that are available throughout the entire sysplex, or all data sets that are preallocated to one session. As this may be a long list, you can use the **DDNAMES/DSNAMES** option on the Report Options panel to compress the data set level information per system.

Figure 13 on page 31 and Figure 14 on page 31 show what the data displayed in Figure 11 on page 29 and Figure 12, respectively, look like if the **DDNAMES/DSNAMES** option is used to condense the information displayed.



```

RMF V2R2 Data Index - RMFPLEX                               Line 1 of 4
Command ==>                                               Scroll ==> HALF
Samples: 118      System: RMFE Date: 11/27/15 Time: 10.12.00 Range: 120 Sec
-----Begin-----           -----End-----
System --Date-- --Time--       --Date-- --Time--
MVS1  11/27/15 10.03.20        11/27/15 10.12.00
MVS2  11/27/15 09.11.00        11/27/15 10.12.00
MVS3  11/27/15 09.11.00        11/27/15 10.12.00
TEST                                     * * * No response * * *
    
```

Figure 13. Data Index - Condensed Version

```

RMF V2R2 Data Index - RMFPLEX                               Line 1 of 3
Command ==>                                               Scroll ==> HALF
Samples: 37      System: RMFB Date: 09/28/16 Time: 13.09.00 Range: 60 Sec
-----Begin-----           -----End-----
System --Date-- --Time--       --Date-- --Time--
RMFA  09/28/16 04.17.00        09/28/16 13.07.00
RMFB  09/28/16 04.16.00        09/28/16 13.10.00
RMFC  09/28/16 04.17.00        09/28/16 06.37.00
    
```

Figure 14. Data Index with Preallocated Data Sets - Condensed Version

The condensed version of the Data Index displays information about data that is available throughout the sysplex. It shows at a glance, for which time ranges data is available on each system, or if no data is available at all or could not be retrieved due to special conditions.

### Data Index — field descriptions

For a description of the report header area of the Data Index, refer to “Header for single-system reports” on page 23, and “Header for sysplex reports” on page 23, where the various header fields are described in more detail.

**Note:** The Data Index title line contains a sysplex ID field like in sysplex reports, but instead of a WLM Samples: field, it shows just the Samples: field as in single system reports. The sysplex ID field in the title line can be blank if you are using data from an old RMF gatherer or have preallocated data from a previous release of RMF.

Table 5. Fields in the Data Index

Field Heading	Meaning
System	The four-character SMF system identifier of the Monitor III gatherer that collected the data. If the identification of the system could not be determined, the field contains '????'.
Begin/End Date Time	These are the beginning and ending dates/times for the data in the usable and not-empty data sets or the in-storage buffers. If the beginning and ending dates are the same, RMF will only display the beginning date. If there is a problem with the data, the dates and times are left blank, and a message is shown.
DDNAME	For a Data Index using active gatherer's data, the DDNAME is the system generated DD name for the data set that has been dynamically allocated for the data gatherer's session. For a Data Index using preallocated data sets, this is the name that was specified in RMFDSxx on the ALLOCATE command issued before the reporter session was started. If there is a problem with the data, the field remains blank. Also, note that the field remains blank if this line is representing the gatherer's in-storage buffer.

## Mon III - Data Index

Table 5. Fields in the Data Index (continued)

Field Heading	Meaning
Data Set Name	This field has two lines. The first line contains the name of the VSAM data set containing the data. The second line is either blank or contains a comment concerning the status of the data. (See "Information and error messages on the Data Index.") Note that the first line is blank when the data represents a gatherer's in-storage buffer.
Begin Date Time	The begin date and time for which data is available on the respective system. This field is blank if there is no data to be represented for the system.
End Date Time	The end date and time for which data is available on the respective system. This field contains a comment about the status of the data if a problem was encountered.

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

### Information and error messages on the Data Index

The following messages can be shown in special cases:

**\*\*\* Currently active \*\*\***

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

**\*\*\* In-storage buffer\*\*\***

The local storage buffer entry of the Monitor III data gatherer

**\*\*\* Empty \*\*\***

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

**\*\*\* No data available \*\*\***

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

**\*\*\* No response \*\*\***

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

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

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

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

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

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

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

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

**\*\*\* Not Found \*\*\***

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

- \*\*\* **Invalid RECSIZE** \*\*\*  
Data set specified with an invalid record size (the data set is unusable)
- \*\*\* **Invalid CISIZE** \*\*\*  
Data set specified with an invalid control interval size (the data set is unusable)
- \*\*\* **Open Error RC=xx reason=xxx** \*\*\*  
Error in opening the data set (the data set is unusable)
- \*\*\* **Close Error RC=xx reason=xxx** \*\*\*  
Error in closing the data set (the data set is unusable)
- \*\*\* **VSAM error RC=xx reason=xxx** \*\*\*  
Error in reading the VSAM data set (the data set is unusable)
- \*\*\* **DYNALLOC RC=xx IRC=xxxx ERC=xxxx** \*\*\*  
Dynamic allocation error (the data set is unusable)
- \*\*\* **UNALLOC RC=xx IRC=xxxx ERC=xxxx** \*\*\*  
Data set unallocated (the data set is unusable)
- \*\*\* **Sample time exceeds current time** \*\*\*  
Data set with a sample time that is later than the current system time. The system time has probably been incorrectly set. (This message does not appear on the screen with preallocated data sets.)
- \*\*\* **Data from sysplex xxxxxxxx** \*\*\*  
For either preallocated data sets or gatherer data sets, a data set that is from a sysplex other than the one selected has been encountered. Only one sysplex can be represented by the data on the Data Index. No other reports can be shown as long as this error persists.
- \*\*\* **Data from system xxxx** \*\*\*  
The reporter cannot report data from gatherer data sets from another system. The gatherer marks the data sets as unusable if more than one system has written to a data set. The reporter cannot access the data in data sets that are marked unusable.  
  
The reporter also cannot report data from different sysplexes in one session.

### Cursor-sensitive control

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

### Data Index options




Figure 15. Data Index Options Panel

The Data Index has two options:

- The *DDNAMES/DSNAMES* option allows you to determine the amount of information that will be displayed. There are two valid values for the *DDNAMES/DSNAMES* option:
  - YES** This value gives you a more detailed Data Index. It contains, at a data-set level, the data that is used on other RMF reports. You can see if there are any specific problems with the available data. See Figure 11 on page 29 and Figure 12 on page 30 for example screens.
  - NO** This value gives you a condensed version of the Data Index. You can see what systems are available for reporting rather than the actual data that is available. See Figure 13 on page 31 and Figure 14 on page 31 for example screens.
- The *Sort Order* is available on the Data Index Options panel, so that it can be accessed with the *ROPTIONS* command from the Data Index screen. If more than one row with the same system identification exist, the usable data sets that are not empty are listed first, then the empty data sets, and finally the unusable data sets. The usable data sets that are not empty are sorted according to the end time of the stored data.

The entire Data Index is sorted by system ID. The *Sort Order* option allows you to change the sorting of the individual data entries for each system. The two valid values for this option are:

**ASCEND**

This value causes the entries in the Data Index to be sorted with the oldest data at the top of the individual system lists. See Figure 11 on page 29 for an example.

**DESCEND**

This value causes the entries in the Data Index to be sorted with the latest data first for each system. See Figure 16 for an example.

```

RMF V2R2 Data Index - RMFPLEX1                               Line 1 of 14
Command ==>                                                Scroll ==> HALF
Samples: 100      System: MVS1 Date: 09/28/16 Time: 10.31.40 Range: 100 Sec
-----Begin/End-----
System --Date-- --Time-- -DDNAME- -----Data Set Name-----
MVS1 09/28/16 09.55.00
           10.31.40                * * * In-storage buffer * * *
MVS1 09/28/16 05.53.20 SYS00003 SYS3.RMF.DS02
           10.11.40                * * * Currently active * * *
MVS1 09/28/16 01.31.40 SYS00002 SYS3.RMF.DS01
           05.53.20
MVS2 09/28/16 09.55.00
           10.31.40                * * * In-storage buffer * * *
MVS2 09/28/16 01.31.40 SYS00003 SYS3.RMF.DS01
           10.11.40                * * * Currently active * * *
TEST
                                           * * * No response * * *
    
```

Figure 16. Data Index - Sort Order Descend

**Note:** Since the condensed version of the Data Index (*DDNAMES/DSNAMES* = *NO*) has only one entry per system, changing the *Sort Order* option will not have any effect on the condensed version of the Data Index.

## CACHDET - Cache Detail Report

The CACHDET report provides detailed information about the activities of one cache subsystem.

### How to request this report

To request the CACHDET report, select **S** on the Primary Menu, and then select **9** on the Sysplex Report menu (shown in Figure 5 on page 24), or enter the following command:

```
CACHDET [subsystem_id]
```

### Contents of the report

Command ==>> _		RMF V2R2	Cache Detail	- SYSPLEX	Line 1 of 20 Scroll ==>> HALF							
Samples: 120	Systems: 5	Date: 09/28/16	Time: 16.11.30	Range: 120	Sec							
		CDate: 09/28/16	CTime: 16.11.25	CRange: 120	Sec							
Volume /Num	SSID	I/O %	I/O Rate	Hit %	Cache Read	Hit Rate	DFW	CFW	DASD Total	I/O Stage	Seq Rate	Async Rate
*ALL		100	42.9	99.5	30.8	11.9	0.0	0.2	0.2	0.6	1.8	
*NOCAC		0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
*CACHE		100	42.9	99.5	30.8	11.9	0.0	0.2	0.2	0.6	1.8	
SYSPXC	05CB 0043	27.9	12.0	100	10.4	1.6	0.0	0.0	0.0	0.0	0.0	
SYSSM3	05C3 0043	11.6	5.0	100	3.9	1.1	0.0	0.0	0.0	0.3	0.4	
SYSAXC	05CA 0043	11.2	4.8	100	3.2	1.6	0.0	0.0	0.0	0.0	0.0	
SYSSM6	05C6 0043	10.8	4.6	98.9	3.7	0.9	0.0	0.1	0.1	0.0	0.3	
SYSSMS	05C0 0043	9.7	4.2	97.6	0.4	3.7	0.0	0.1	0.1	0.0	0.0	
SYSSM8	05C8 0043	6.8	2.9	100	1.6	1.4	0.0	0.0	0.0	0.1	0.1	
SYSSM5	05C5 0043	5.4	2.3	100	2.3	0.1	0.0	0.0	0.0	0.0	0.2	
SYSSM2	05C2 0043	3.7	1.6	100	0.6	1.0	0.0	0.0	0.0	0.0	0.5	
SYSSMB	05DA 0043	3.6	1.6	100	1.4	0.1	0.0	0.0	0.0	0.1	0.0	
SYSSMA	05CD 0043	3.6	1.5	99.5	1.5	0.0	0.0	0.0	0.0	0.1	0.0	
SYSSM7	05C7 0043	3.3	1.4	98.2	1.2	0.2	0.0	0.0	0.0	0.0	0.1	
SYSOPE	05CC 0043	1.6	0.7	97.5	0.6	0.1	0.0	0.0	0.0	0.0	0.0	
SYSSM4	05C4 0043	0.6	0.3	96.9	0.0	0.2	0.0	0.0	0.0	0.0	0.0	
SYSSM9	05C9 0043	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
SYSSM1	05C1 0043	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
SYSSMC	05DB 0043	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
MVSWR1	05FF 0043	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	

Figure 17. CACHDET Report

There is no graphic version of this report available.

If you place the cursor on any field of the first two columns, a pop-up window appears showing details for the selected volume. Cursor-sensitive control of the third column leads you to a pop-up window with details for the selected SSID.

Mon III - CACHDET

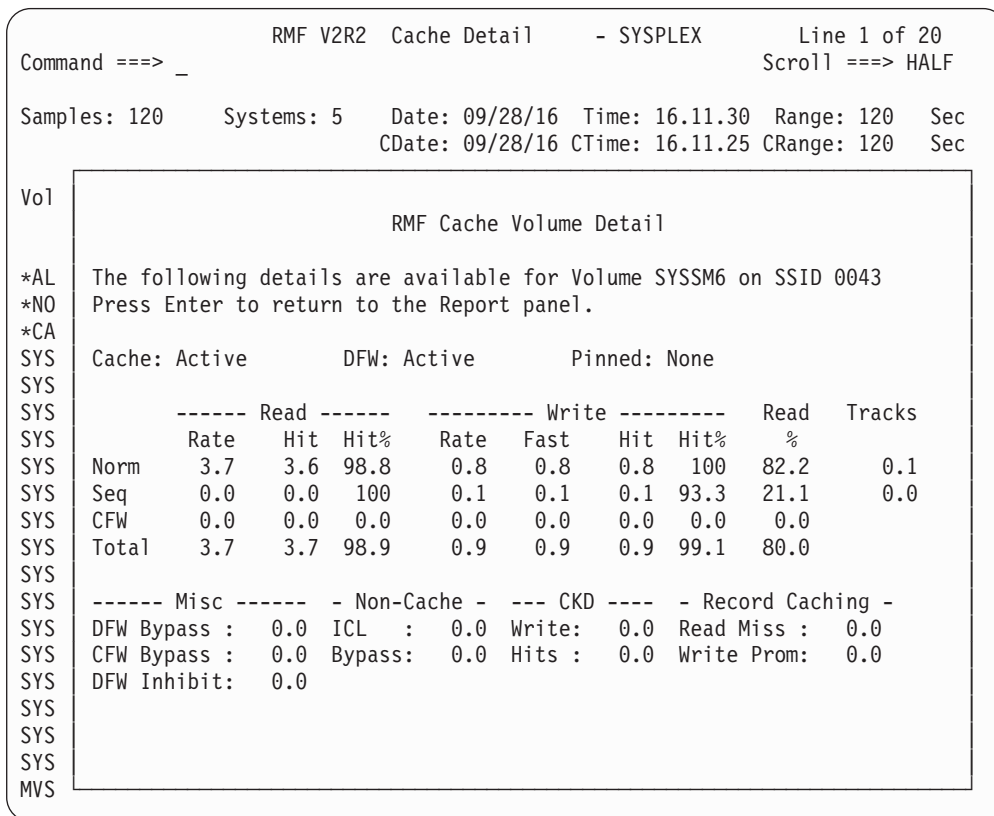


Figure 18. CACHDET Report - Volume Details

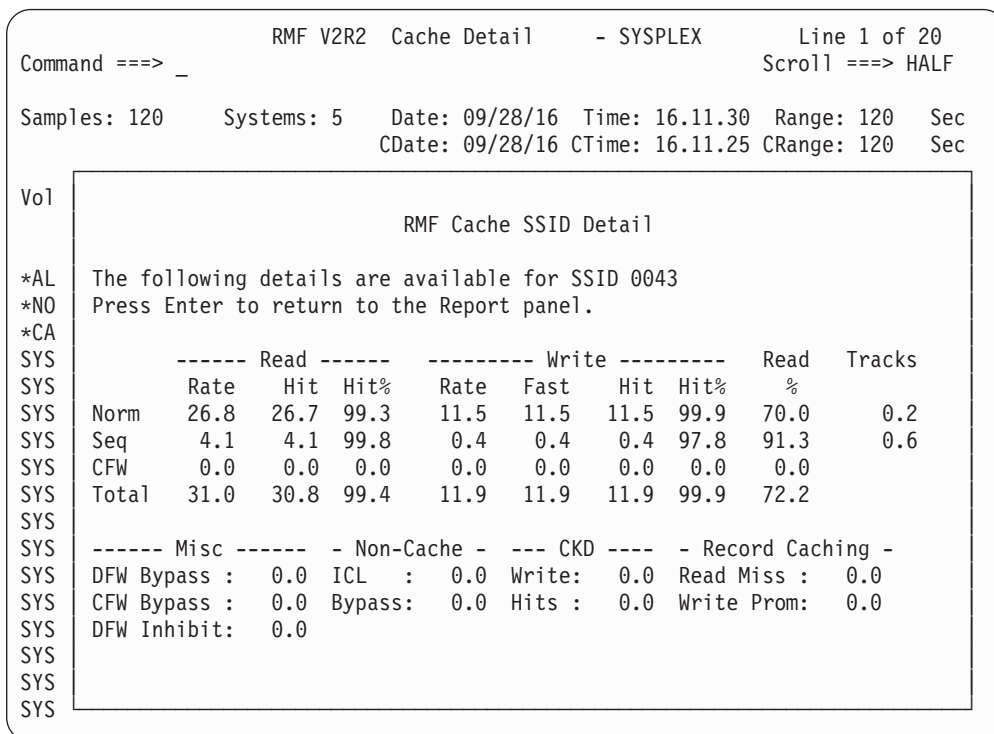


Figure 19. CACHDET Report - SSID Details

### Field descriptions

Table 6. Fields in the CACHDET Report

Field Heading	Meaning
	Device reserve activity can cause a data gatherer interface to wait until a RESERVE has been released. This in turn can cause the cache interval to be much longer than a regular RMF interval.  Therefore, CDate, CTime and CRange are used to show the actual point in time to which the cache interval start is related, and the actual cache interval length. All rates shown in the report are based on CRange, not on Range.
CDate	Date on which the cache interval started.
CTime	Time at which the cache interval started.
CRange	Cache interval time.
Volume	Volume serial number or one of the following: *ALL All devices belonging to the reported storage subsystem. *NOCAC All non-cached devices. *CACHE All cached devices.  These lines will be given only if the report shows data for one specific subsystem ID.
/Num	Device number.
SSID	Subsystem identifier; a number assigned during the installation of the subsystem that uniquely identifies the storage subsystem.
I/O %	Percentage of I/O requests to this device or category, compared to the total number of I/O requests sent to the subsystem.
I/O Rate	Rate of I/O requests.
Hit %	Percentage of I/Os that were processed within the cache (cache hits) based on the total number of I/Os.
Cache Hit Rate - I/O rate of all cache hits.	
Read	Rate of SEARCH/READ requests that completed without accessing the DASD.
DFW	Rate of DFW requests.
CFW	Rate of WRITE and READ-AFTER-WRITE requests that are processed in cache.
DASD I/O - I/O rate of all requests that accessed DASD.	
Total	I/O rate of all requests that accessed DASD. This is the sum of Stage rates (see below) and other request rates (inhibit cache load, DFW BYPASS, CFW BYPASS, DFW INHIBIT).
Stage	Rate of normal or sequential I/O requests that accessed DASD.
Seq Rate	Rate of tracks that have been staged due to cache misses for sequential I/O requests.
Async Rate	Rate of tracks that have been destaged asynchronously.

Table 7. Fields in the CACHDET Report - Volume and SSID Details

Field Heading	Meaning
Cache	Caching status of the device. <b>Active</b> Caching is active - requests to this device can be processed without DASD access. <b>Deact pending</b> Cache has been deactivated on request from host system or support facility, but transfer of modified data to DASD has failed. <b>Deactivated</b> Caching has been deactivated for this device.
DFW	Status of the DASD FAST WRITE option. <b>Active</b> DASD FAST WRITE requests can be processed for this device. <b>Deact pending</b> DASD FAST WRITE has been terminated on request by host system or support facility, but transfer of modified data to DASD is in progress or has failed. <b>Deactivated</b> DASD FAST WRITE requests are ignored for this device.

## Mon III - CACHDET

Table 7. Fields in the CACHDET Report - Volume and SSID Details (continued)

Field Heading	Meaning
Pinned	A device has failed, and data that has not yet been written to DASD is pinned in cache or NVS for later recovery. <b>None</b> No data is pinned for this device. <b>DFW not suspend</b> Pinned data exists, but DASD FAST WRITE has not been suspended. <b>DFW suspend</b> Pinned data exists, and DASD FAST WRITE has been suspended.
Cache I/O Request - The channel command DEFINE EXTENT specifies the way the cache will be used. There are three categories: Norm - Seq - CFW	
Norm	Cache will be managed by least-recently-used (LRU) algorithm for making cache space available.
Seq	Tracks following the track assigned in the current CCW chain are promoted. They will be transferred from DASD to cache in anticipation of a short-term requirement.
CFW	WRITE and READ-AFTER-WRITE requests are processed in cache. The data might not be written to DASD. Because CFW does not use the NVS, the application is responsible for restoring the data after a cache or subsystem failure.
Total	This is either the sum of I/O requests, the total I/O rate, or the average hit ratio for the three categories previously described.
Read - Cache I/O requests that searched or read data from DASD. This is the number of channel operations that had at least one SEARCH or READ command but no WRITE commands. It is counted for cache devices only.	
Rate	Rate of SEARCH/READ requests.
Hit	Rate of SEARCH/READ requests that completed without accessing the DASD.
Hit%	Percentage of SEARCH/READ requests that completed without accessing the DASD.
Write - Cache I/O requests that wrote data to DASD. This is the number of channel commands that had at least one WRITE command. It is counted for cache devices only.	
Rate	Rate of WRITE requests.
Fast	Rate of DASD/CACHE FAST WRITE requests.
Hit	Rate of DASD/CACHE FAST WRITE requests that completed without accessing the DASD (fast write hit).
Hit%	Percentage of DASD/CACHE FAST WRITE requests that completed without accessing the DASD (fast write hit).
Read %	Percentage of READ requests based on the sum of all READ and WRITE requests (excluding ICL and BYPASS).
Tracks	Rate of tracks transferred from DASD to cache.
Misc - Miscellaneous cache activity rates.	
DFW Bypass	Rate of DASD FAST WRITE requests that would have resulted in a DFW hit; however, NVS was overutilized causing writes to be sent directly to DASD. This values is also known as DFW RETRY.
CFW Bypass	Rate of operations that did not transfer a track from DASD into cache because no free segments were available. In this case, there is no destaging in favor of I/O requests with the CACHE FAST WRITE attribute. The I/O goes directly to the DASD.
DFW Inhibit	If DASD FAST WRITE is active, this is the rate of WRITE requests that inhibit DASD FAST WRITE.  If DASD FAST WRITE is inactive, this is the rate of WRITE requests that directly accessed the DASD, even with DASD FAST WRITE turned on.
Non-Cache - I/O READ requests that switched off cache processing.	
ICL	Inhibit cache load. Rate of I/O requests that inhibited load of data into cache although the data was not found in the cache.
Bypass	Rate of I/O requests that explicitly bypassed the cache, irrespective of whether the data is in the cache or not.
CKD - CKD (Count-key-data) is a format used to store data on DASD. The counts shown in this section are contained in the total WRITE count.	
Write	Rate of WRITE I/O requests in CKD format.



Table 7. Fields in the CACHDET Report - Volume and SSID Details (continued)

Field Heading	Meaning
Hits	Rate of I/O requests in CKD format that could be resolved in the cache.
Record Caching	Record caching is done dynamically upon a decision made by DCME or the microcode. It may improve overall cache performance if caching of whole tracks would waste cache storage. The decision is based on the number of I/Os, the hit ratio, and the locality of reference of a certain entity of data.
Read Miss	Rate of instances in which a record requested for READ was not found in the cache, and access to DASD was required.
Write Prom	Rate of instances in which a record requested for WRITE was found in the cache, and access to DASD was not required.

## Report options

```

RMF Cache Report Options                               Line 1 of 2
Command ==>>                                         Scroll ==>> HALF

Change or verify parameters. To exit press END.
Changes will apply to the Cache Detail report.

SSID          ==>> 0046 ALL or one of the available subsystem IDs below

                Available Subsystem IDs
0040   0041   0044   0046   0047   0048   004A   004B   004C   004D
0050   0051   0054   0056   0060   006A   006B   007A   007B   008A
    
```

Figure 20. CACHDET Report Options

In the Report Options panel, you can select whether you want to get the CACHDET report with one or with all available subsystem IDs.

## CACHSUM - Cache Summary Report

The Cache Summary report (CACHSUM) provides an overview about the activities in the cache subsystem for all SSIDs. You might take this as starting point when analyzing I/O performance to get a first impression about the I/O processing.

If you feel that further analysis is required, you may continue with the Cache Detail report (see “CACHDET - Cache Detail Report” on page 35).

### How to request this report

To request the CACHSUM report, select **S** on the Primary Menu, and then select **8** on the Sysplex Report menu (shown in Figure 5 on page 24), or enter the following command:

```
CACHSUM
```

## Contents of the report

```

RMF V2R2 Cache Summary - SYSPLEX Line 1 of 21
Command ==> _ Scroll ==> HALF
Samples: 120 Systems: 5 Date: 09/28/16 Time: 08.25.00 Range: 120 Sec
CDate: 09/28/16 CTime: 08.24.55 CRange: 120 Sec

SSID CUID Type-Mod Size I/O Hit Hit -- Miss --- Read Seq Async Off
Rate % Rate Total Stage % Rate Rate Rate

0010 071D 9396-001 1024M 99.8 98.8 98.6 1.2 1.2 99.2 1.6 0.1 0.0
0011 0520 9396-001 1024M 5.4 100 5.4 0.0 0.0 100 0.0 0.0 0.0
0030 0269 3990-006 256M 90.6 99.8 90.4 0.2 0.2 100 0.0 0.0 0.0
0040 0401 9393-002 1536M 142.0 100 141.9 0.0 0.0 97.7 0.0 0.3 0.0
0041 0460 9393-002 1536M 1.1 100 1.1 0.0 0.0 68.2 0.0 0.2 0.0
0042 05AA 9393-002 1536M 0.9 100 0.9 0.0 0.0 70.4 0.0 0.1 0.0
0043 05C6 9393-002 1536M 42.9 99.5 42.7 0.2 0.2 72.2 0.6 1.8 0.0
0044 0200 3990-006 64M 129.7 99.5 129.0 0.7 0.7 98.2 0.2 0.0 0.0
0060 0627 9393-002 1536M 9.2 29.3 2.7 0.0 0.0 40.3 0.0 0.4 0.0
0061 0654 9393-002 1536M 4.3 99.8 4.3 0.0 0.0 64.6 0.3 0.4 0.0
0062 06B0 9393-002 1536M 7.4 30.9 2.3 0.1 0.1 17.5 0.0 0.6 0.0
0063 06CE 9393-002 1536M 0.3 100 0.3 0.0 0.0 88.9 0.0 0.0 0.0
    
```

Figure 21. CACHSUM Report

There is no graphic version of this report available.

If you place the cursor on any field of the first three columns, a pop-up window appears showing details for the selected SSID. Cursor-sensitive control of the other columns leads you to the CACHDET report.

```

RMF V2R2 Cache Summary - SYSPLEX Line 1 of 12
Command ==> _ Scroll ==> HALF
Samples: 120 Systems: 5 Date: 09/28/16 Time: 16.11.30 Range: 120 Sec
CDate: 09/28/16 CTime: 16.11.25 CRange: 120 Sec

SSI
RMF Cache SSID Information

001 The following details are available for SSID 0043
001 Press Enter to return to the Report panel.
003
004 CUID : 05C6 Cache : Active NVS : Active
004 Type-Mod: 9393-002 Config: 1536M Config: 8192K
004 Avail : 1536M Pinned: 0
004 Offl : 0
004 Pinned: 0
006
006 ----- Read ----- Write ----- Read Tracks
006 Rate Hit Hit% Rate Fast Hit Hit% %
006 Norm 26.8 26.7 99.3 11.5 11.5 11.5 99.9 70.0 0.2
006 Seq 4.1 4.1 99.8 0.4 0.4 0.4 97.8 91.3 0.6
006 CFW 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0
006 Total 31.0 30.8 99.4 11.9 11.9 11.9 99.9 72.2
006
006
006
    
```

Figure 22. CACHSUM Report - SSID Details

## Field description

Table 8. Fields in the CACHSUM Report

Field Heading	Meaning
	Device reserve activity can cause a data gatherer interface to wait until a RESERVE has been released. This in turn can cause the cache interval to be much longer than a regular RMF interval.  Therefore, CDate, CTime and CRange are used to show the actual point in time to which the cache interval start is related, and the actual cache interval length. All rates shown in the report are based on CRange, not on Range. <b>Note:</b> The reported storage capacities for cache and non-volatile storage (NVS) represents only the Cluster Processor Complex in the storage server, that controls the subsystem. Since a typical storage server has two clusters, you must double the reported capacities to get the actual sizes.
CDate	Date on which the cache interval started.
CTime	Time at which the cache interval started.
CRange	Cache interval time.
SSID	Subsystem identifier: a number assigned during the installation of the subsystem that uniquely identifies the storage subsystem.
CUID	Physical control unit number of the caching subsystem. This is equal to the lowest device number, or the device that has been turned online first, respectively.
Type-Mod	Hardware type and model.
Size	Amount of physical storage that is configured in this storage subsystem (in giga- or megabytes).
I/O Rate	Rate of I/O requests.
Hit %	Percentage of I/Os that where processed within the cache (cache hits) based on the total number of I/Os.
Hit Rate	I/O rate of all cache hits. This is the sum of: <b>READ</b> Rate of SEARCH/READ requests that completed without accessing the DASD <b>DFW</b> Rate of DASD FAST WRITE requests <b>CFW</b> Rate of WRITE and READ-AFTER-WRITE requests that are processed in cache
Miss Total	I/O rate of all requests that accessed DASD.  This is the sum of Stage rates (see below) and other request rates (inhibit cache load, BYPASS, CFW BYPASS, DFW INHIBIT).
Miss Stage	Rate of normal or sequential I/O requests that accessed DASD.
Read %	Percentage of READ requests based on all READ and WRITE requests.
Seq Rate	Rate of tracks that have been staged due to cache misses for sequential I/O requests.
Async Rate	Rate of tracks that have been destaged asynchronously.
Off Rate	Rate of I/O requests to non-cached devices.

Table 9. Fields in the CACHSUM Report - SSID Details

Field Heading	Meaning
	<b>Note:</b> The reported storage capacities for cache and non-volatile storage (NVS) represents only the Cluster Processor Complex in the storage server, that controls the subsystem. Since a typical storage server has two clusters, you must double the reported capacities to get the actual sizes.
CUID	Physical control unit number of the caching subsystem. This is equal to the lowest device number, or the device that has been turned online first, respectively.
Type-Mod	Hardware type and model.

## Mon III - CACHSUM

Table 9. Fields in the CACHSUM Report - SSID Details (continued)

Field Heading	Meaning
Cache	<p>Caching status of the subsystem.</p> <p><b>Active</b> Caching is active (online and usable).</p> <p><b>Deact failed</b> A command requesting deactivation of cache has been received, but destaging to DASD has failed.</p> <p><b>Deact in process</b> A command requesting deactivation of cache has been received, and destaging from cache to DASD is still in progress.</p> <p><b>Deactivated</b> Cache has been deactivated by request from host system or support facility.</p> <p><b>Error termination</b> An internal error stopped cache (cache is offline).</p> <p><b>Maintenance</b> Cache has been disabled for maintenance.</p> <p><b>Pending active</b> Caching is pending active, that is, cache is in the process of being brought online.</p>
Cache Config	Amount of physical storage that is installed in this storage subsystem.
Cache Avail	Amount of storage that is available for caching. This is the total cache size minus the amount used by the subsystem for the cache directory, minus the amount pinned and offline storage.
Cache Offl	Amount of storage that is offline because of a host or subsystem error.
Cache Pinned	Amount of storage that is unavailable because a DASD failure is preventing the subsystem from destaging data. The data is pinned in cache.
NVS	<p>Overall status of the non-volatile storage (NVS).</p> <p><b>Active</b> NVS is online and usable.</p> <p><b>Deact failed</b> A command requesting deactivation of NVS has been received but transfer from NVS to DASD has failed.</p> <p><b>Deact in process</b> A command requesting deactivation of NVS has been received, and destaging to DASD is still in progress.</p> <p><b>Deactivated</b> NVS has been deactivated by request from host system or support facility.</p> <p><b>DFW inhibited</b> DASD FAST WRITE is inhibited because the battery is defective.</p> <p><b>Error termination</b> An internal error caused termination of NVS.</p> <p><b>Maintenance</b> NVS has been disabled for maintenance by the support facility.</p>
NVS Config	Amount of NVS that is installed in this storage subsystem.
NVS Pinned	Amount of NVS that is unavailable because a DASD failure is preventing the subsystem from destaging the data. The data is pinned in NVS.

Please refer to Table 7 on page 37 for a description of the other fields in the pop-up window of the CACHSUM report.

## CFACT - Coupling Facility Activity Report

The Coupling Facility Activity report (CFACT) gives you information about the activities in each structure.

You can use this report for analyzing in detail each structure in your coupling facilities. You see the type of a structure and the activities from each system that uses a structure. If you want to get more details, you will receive them through cursor-sensitive control. A pop-up window shows the allocation details and the name of the address space which is currently using the structure. If you experience performance problems for one or several structures in your coupling facilities, you should investigate into the appropriate applications.

## How to request this report

To request the Coupling Facility Activity report, select **S** from the Primary Menu and then select **7** on the Sysplex Report menu (shown in Figure 5 on page 24), or enter the following command:

```
CFACT [cfname]
```

In addition, you can navigate to this report through cursor-sensitive control from the CFOVER report or CFSYS report.

## Special considerations

Data gathering for this report is enabled by default, using the gathering option CFDETAIL. With CFDETAIL, a large amount of data is being gathered that enables you to get many details about the usage of each structure in the coupling facility. This data gathering is done only on one member of the sysplex. This is called *sysplex master gathering* and has been implemented to reduce performance overhead on non-master members and to reduce the amount of data in SSHs and SMF records. The RMF Sysplex Data Server determines internally which member of the sysplex will be the master. This can be controlled externally by the operator or system administrator specifying the Monitor III MASTER/NOMASTER data gatherer option.

If you run the Monitor III reporter with preallocated VSAM data sets, you should ensure that you have allocated all data sets belonging to the sysplex to be able for reporting everything that has been gathered.

## Contents of the report

Command ==>		RMF V2R2	CF Activity	- TRXPLEX	Line 1 of 78 Scroll ==> CSR					
Samples: 30	Systems: 2	Date: 09/28/16	Time: 08.41.30	Range: 30	Sec					
CF: X7CFP87	Type	ST	System	CF Util %	--- Sync Rate	--- Avg Serv	----- Async Rate	Avg Serv	Chng %	Del %
ISGLOCK	LOCK	A	*ALL	0.0	0.5	3	0.0	0	0.0	0.0
	LOCK		R70		0.5	3	0.0	0	0.0	0.0
	LOCK		R71		0.0	0	0.0	0	0.0	0.0
ISTGENERIC	LIST	AP	*ALL	0.0	27.1	14	0.0	0	0.0	0.0
	LIST		R70		27.1	14	0.0	0	0.0	0.0
	LIST		R71		0.0	0	0.0	0	0.0	0.0
SYSZWLM_BC772827	CACHE	AP	*ALL	0.0	0.4	8	0.1	126	0.0	0.0
	CACHE		R70		0.4	8	0.1	126	0.0	0.0
	CACHE		R71		0.0	0	0.0	0	0.0	0.0
SYSZWLM_WORKUNIT	CACHE	AP	*ALL	0.0	0.0	0	0.0	0	0.0	0.0
	CACHE		R70		0.0	0	0.0	0	0.0	0.0
	CACHE		R71		0.0	0	0.0	0	0.0	0.0
SYSZWLM_7A862827	CACHE	AS	*ALL	0.0	0.0	0	0.0	0	0.0	0.0
	CACHE		R70		0.0	0	0.0	0	0.0	0.0
	CACHE		R71		0.0	0	0.0	0	0.0	0.0
THRLSTSCMKP1_1	LIST	A	*ALL	14.6	1106	6	9.0	121	0.0	0.0
	LIST		R70		1106	6	9.0	121	0.0	0.0
	LIST		R71		0	0	0.0	0	0.0	0.0

Figure 23. CFACT Report

There is no graphic version of this report available.

If you place the cursor on any of the lines with coupling facility structure values, a pop-up window appears (one out of Figure 24 through Figure 26 on page 45, depending on the structure type), showing details for this structure. The pop-up window from Figure 27 on page 45 appears instead of the one shown in Figure 25, if you select a List Structure which also uses SCM storage.

Some of these detail values are available only, if the Monitor III gatherer is running with the gathering option **CFDETAIL**, which is the default.

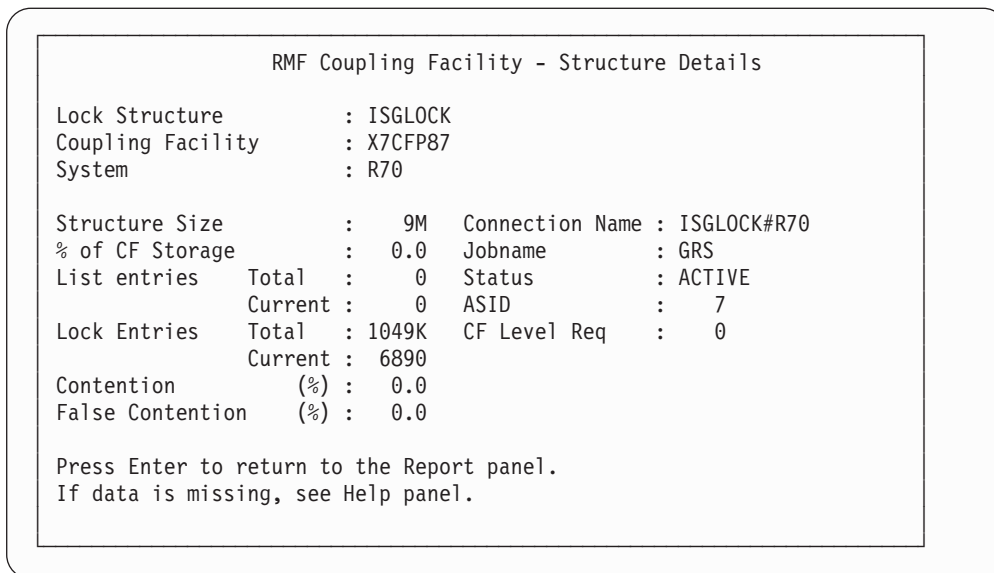


Figure 24. CFACT Report - Details for a Lock Structure

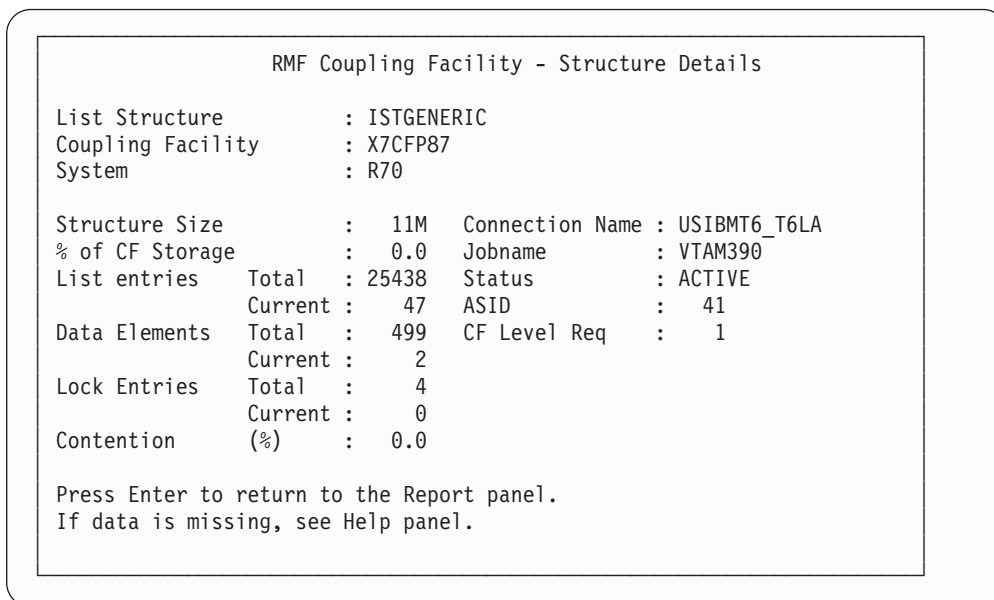


Figure 25. CFACT Report - Details for a List Structure

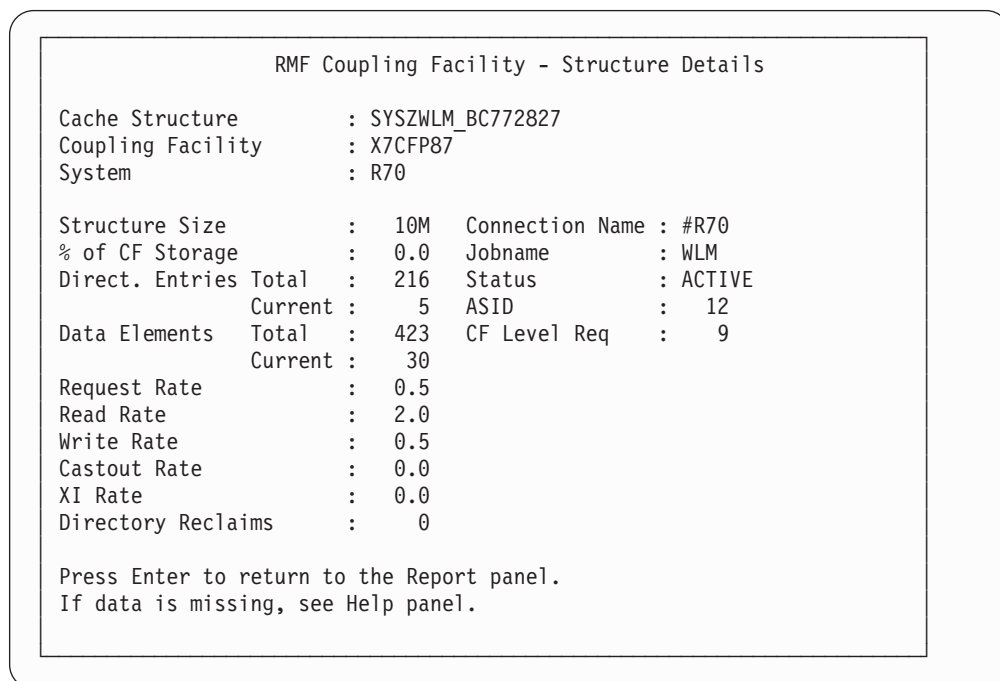


Figure 26. CFACT Report - Details for a Cache Structure

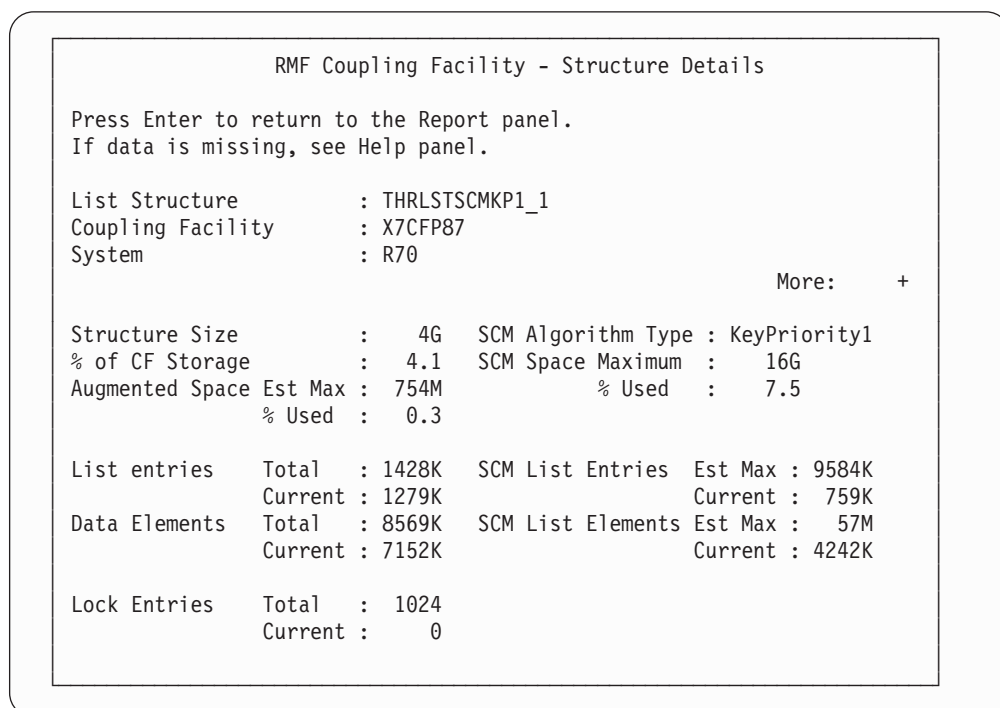


Figure 27. CFACT Report - Details for a List Structure using SCM storage (1)

Pressing PF8 displays more information as shown in Figure 28 on page 46.

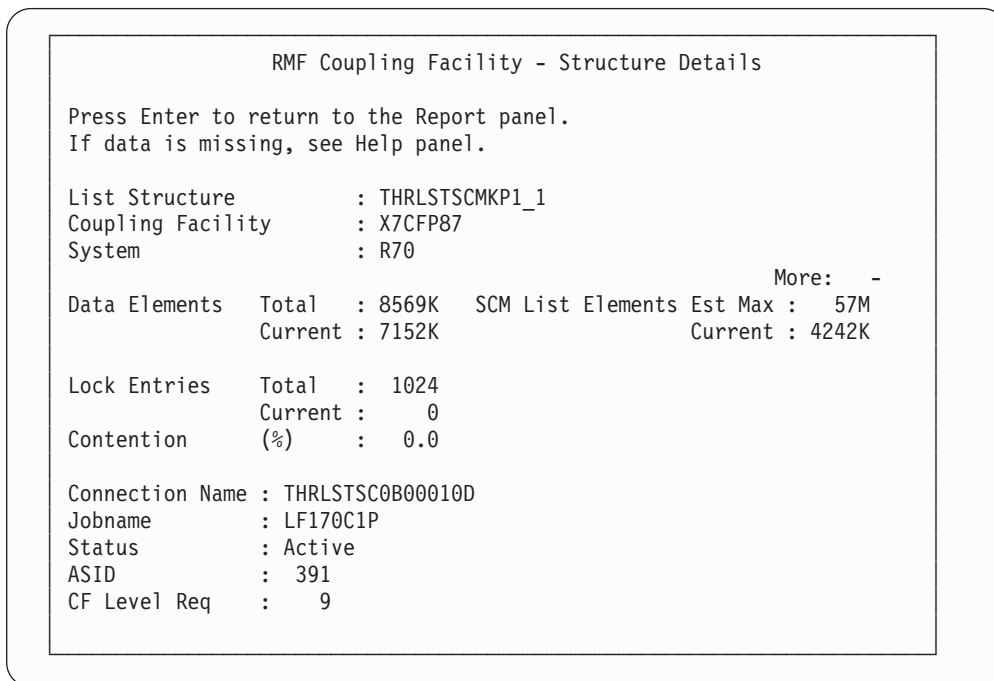


Figure 28. CFACT Report - Details for a List Structure using SCM storage (2)

### Field descriptions

Table 10. Fields in the CFACT Report

Field Heading	Meaning
<b>Note:</b> Each rate is reported as '<0.1' if the value is greater than 0 but would be rounded to 0.	
CF	Coupling facility name.
Structure Name	Name given to the structure by the coupling facility policy specification in the Function Couple Data Set. It has up to 16 characters and is unique within a sysplex.
Type	Type indicates whether the structure is a list (LIST), lock (LOCK) or cache (CACHE) structure. The structures being reported are grouped by structure type.
ST	<p>Status - can be one of following states in the SYSPLEX (*ALL) view data line for a structure:</p> <p><b>A</b> Active - structure is allocated and connected to at least one system during the entire MINTIME.</p> <p><b>AP</b> Active/primary - structure has been active as primary structure during MINTIME (rebuild-old).</p> <p><b>AS</b> Active/secondary - structure has been active as secondary structure during MINTIME (rebuild-new).</p> <p><b>I</b> Inactive - structure got disconnected from all systems during MINTIME.</p> <p><b>N</b> New - structure became allocated and connected to at least one system during MINTIME.</p> <p><b>PA</b> Active/primary - Structure is the rebuild-old (primary) structure in an asynchronous duplexing rebuild process.</p> <p><b>SA</b> Active/secondary - Structure is the rebuild-new (secondary) structure in an asynchronous duplexing rebuild process.</p> <p><b>Note:</b> There is no structure activity data reported for an inactive structure even if it was active earlier in the MINTIME. The same applies for structures that became active during MINTIME. Therefore, all values for these structures are reported as blank.</p>
System	<p>System name for the system connected to the structure (from IEASYSxx Parmlib member, SYSNAME parameter).</p> <p>In the first data line for a structure, the name is '*ALL' to indicate that this line shows the SYSPLEX view of the data rather than a single system view.</p>



Table 10. Fields in the CFACT Report (continued)

Field Heading	Meaning
CF Util %	The percentage of CF processor time used by the structure. The sum of the values in this column is less than 100%, because not all CF processor time is attributable to structures.  'N/A' is shown in this field if the CF level is lower than 15.
Sync Rate	Number of hardware operations per second that started and completed synchronously to the coupling facility on behalf of connectors to the structure.
Sync Avg Serv	Average time in microseconds required to satisfy a synchronous coupling facility request for this structure.
Async Rate	Number of hardware operations per second that started and completed asynchronously to the coupling facility on behalf of connectors to the structure.
Async Avg Serv	Average time in microseconds required to satisfy an asynchronous coupling facility request for this structure. This value also includes operations that started synchronously but completed asynchronously.
Async Chng %	Percentage of asynchronous requests for this structure that changed from synchronous to asynchronous because the requests could not be serviced as synchronous request. This field reports only those requests which were changed due to a subchannel busy condition and can be used as an indicator of a shortage of subchannel resources.  Request conversions caused by heuristic sync/async algorithms used to optimize the coupling efficiency of workloads using the CF are not included.
Async Del %	Percentage of asynchronous hardware operations for this structure being delayed by either subchannel contention or dump serialization.  This value can exceed 100% if there are several delays for one request during the MINTIME.
<p><b>Note:</b></p> <ol style="list-style-type: none"> <li>The availability of the data in the pop-up panel depends on the <b>CFDETAIL</b> option of the Monitor III gatherer session. If this option is not active, all values marked as follows have to be used carefully:   <b>DET</b> Value will not be reported   <b>MON I</b> Value is a snapshot value taken at the end of the previous Monitor I gathering interval.                       In some cases, the values can be blank, for example, if SMF data gathering for the coupling facility (SMF record type 74-4) is not active, or if a structure has no connection to some members in the sysplex.</li> <li>Fields marked with <sup>1)</sup> are only available for a single system, but not for the sysplex view from the *ALL summary line. You get the single systems view by selecting option Detail ==&gt; Yes from the <i>RMF Coupling Facility Report Options</i> (Figure 29 on page 49).</li> <li>Each rate is reported as &lt;0.1 if the value is greater than 0 but would be rounded to 0.</li> <li>Fields marked with <sup>2)</sup> only apply to List Structures exploiting SCM storage extension.</li> <li>Each rate is reported as '&lt;0.1' if the value is greater than 0 but would be rounded to 0.</li> </ol>	
Structure Size (MON I)	Amount of storage in bytes that is currently allocated for this structure.
% of CF Storage (MON I)	The percentage of the total coupling facility storage allocated to this structure.
List Entries (MON I)	Number of list entries in use in a LIST or LOCK structure. <b>Total</b> Total number. <b>Current</b> Number of list entries in use.
Data Elements (MON I)	Number of data elements in use in a LIST or CACHE structure. <b>Total</b> Total number. <b>Current</b> Number of list data elements in use.
SCM Algorithm Type <sup>2)</sup> (MON I)	Type of algorithm used by the coupling facility to control the movement of structure objects between coupling facility real storage and storage class memory.

## Mon III - CFACT

Table 10. Fields in the CFACT Report (continued)

Field Heading	Meaning
Augmented Space <sup>2)</sup> (MON I)	<b>Est Max</b> Estimated maximum amount of augmented space in bytes that can be assigned for this structure. <b>% Used</b> Percentage of maximum augmented space that is in use by the structure.
SCM Space <sup>2)</sup> (MON I)	<b>Maximum</b> Maximum amount of storage class memory in bytes that this structure can use. <b>% Used</b> Percentage of maximum storage class memory that is in use by the structure.
SCM List Entries <sup>2)</sup> (MON I)	<b>Est Max</b> Estimated maximum number of list entries that can reside in storage class memory for the structure. <b>Current</b> Number of existing structure list entries that reside in storage class memory.
SCM List Elements <sup>2)</sup> (MON I)	<b>Est Max</b> Estimated maximum number of list elements that can reside in storage class memory for the structure. <b>Current</b> Number of existing structure list elements that reside in storage class memory.
Lock Entries (MON I)	Number of lock table entries in use in a serialized LIST or a LOCK structure. <b>Total</b> Total number. <b>Current</b> Number of lock table entries in use. <b>Note:</b> This is an approximate number, since it is based on sampling.
Contention (%)	For serialized LIST structures and for LOCK structures only: percentage of all external requests issued by connectors delayed due to contention on a lock.
False Contention (%)	For LOCK structures only: percentage of all external requests issued by connectors that experience "hash contention".  This occurs because a hashing algorithm is used to map a lock request to a lock table entry. When more than one lock request maps to a lock table entry, there is the potential for contention delay. You may need to increase the size of the lock table. <b>Note:</b> It is possible for an application to have unusual lock reference patterns that cause storage contention regardless of the size of the lock structure.
Connection Name <sup>1)</sup> (DET)	Name of the last connection from the selected system.
Jobname <sup>1)</sup> (DET)	Name of the job that made the last connect from the selected system.
Status <sup>1)</sup> (DET)	The status of the last connection from the selected system. <b>Active</b> Connection established. <b>FailPers</b> Failed Persistent: Connection with CONDISP=KEEP has failed and all of the event exit responses have been received with RELEASECONN=NO. <b>Failing</b> Connection terminated abnormally and not all of the event exit responses have been received. <b>Disc</b> Disconnecting: Connection disconnected and not all of the event exit responses have been received. <b>NotKnown</b> None of the above.
ASID <sup>1)</sup> (DET)	ASID of the job that made the last connect from the selected system.
CF Level Req <sup>1)</sup> (DET)	The CFCC Microcode Level requested by the last connect from the selected system.
The following values apply to CACHE structures only.	
Directory Entries (MON I)	The number of directory entries in a CACHE structure. <b>Total</b> Total number. <b>Current</b> Number of currently filled directory entries.

Table 10. Fields in the CFACT Report (continued)

Field Heading	Meaning
Request Rate	Number of external requests to this structure on behalf of connectors per second.
Read Rate (DET)	Number of occurrences the coupling facility returned data on a read request by any connector (read hit) per second.  Directory-only caches will always have a zero value reported since there is no data to be returned.
Write Rate (DET)	Number of occurrences per second data has been written to the cache structure.  Directory-only caches will always have a zero value reported since there are no data writes possible.
Castout Rate (DET)	Number of CASTOUT processings per second.  Castout is the process of writing changed cache data to permanent storage. This rate is of interest for store-in cache structures (for example, DB2® global buffer pool structures) in determining the volume of changed data being removed from the structure.
XI Rate (DET)	Number of times per second a data item residing in a local buffer pool was marked invalid by the coupling facility.  XI values are seen for directory, store-in and store-thru caches. This rate reflects the amount of data sharing among the users of the cache and the amount of write or update activity against the data bases.
Directory Reclaims (DET)	Number of cache directory reclaims happened during the RMF MINTIME.  Directory reclaims occur when the total number of used unique entities exceeds the total number of directories. Whenever this shortage of directory entries occurs, the coupling facility will reclaim in-use directory entries associated with unchanged data. All users of that data must be notified that their copy of the data is invalid. As a consequence, it may happen that this data must be re-read from DASD and registered to the coupling facility again.  Directory reclaim activity can be avoided by increasing the directory entries for a particular structure.

## Report options

```

RMF Coupling Facility Report Options  Line 1 of 1
Command ==>                               Scroll ==> HALF

Change or verify parameters. To exit press END.
Changes will apply to the CFACT, the CFOVER and the CFSYS report.

Name  ==> ALL      ALL or one of the available coupling facilities below
Type  ==> ALL      Structure type (LIST, LOCK, CACHE or ALL) in CFACT report
Detail ==> YES     Show single system data (YES or NO) in CFACT report

                                     Available Coupling Facilities
CF5B      CF6B
    
```

Figure 29. Coupling Facility Report Options Panel

**Name** Either **ALL** or the name of one of the coupling facilities being available in the sysplex as shown in the field **Available Coupling Facilities**.

The value for Name that you specify on this panel affects all Coupling Facility reports.

**Type** To select a specific structure type in the CFACT report, you can request LIST, LOCK, CACHE, or ALL.

**Detail** With this option, you can select the level of detail in the CFACT report:

- YES The report contains data for the sysplex and all single systems.
- NO The report contains data for the sysplex only.

**Available Coupling Facilities**

The list of all coupling facilities which are currently connected to the sysplex.

## CFOVER - Coupling Facility Overview Report

The Coupling Facility Overview report (CFOVER) gives you information about all coupling facilities which are connected to the sysplex.

You might start the investigation of the performance of the coupling facilities in your sysplex with the CFOVER report. You get an overview about all coupling facilities showing the utilization of the processors and the storage. If you experience high values for these resources, this might indicate contention in the coupling facilities which could lead to internal queues causing performance problems.

In addition, you can evaluate the request rates to analyze whether the usage of the coupling facilities is well balanced. This, of course, will not be the case if you have one coupling facility for production and the other one just as stand-by.

### How to request this report

To request the Coupling Facility Overview report, select **S** from the Primary Menu and then select **5** on the Sysplex Report menu (shown in Figure 5 on page 24), or enter the following command:

```
CFOVER [cfname]
```

### Contents of the report




Figure 30. CFOVER Report

There is no graphic version of this report available.

### Field descriptions

Table 11. Fields in the CFOVER Report

Field Heading	Meaning
CF Policy	Name of the current coupling facility policy.
Activated at	Date and time the current coupling facility policy was activated.

Table 11. Fields in the CFOVER Report (continued)

Field Heading	Meaning
Coupling Facility	<p>The following information about the coupling facility is provided:</p> <p><b>Name</b> coupling facility name</p> <p><b>Type</b> coupling facility processor type</p> <p><b>Mod</b> coupling facility processor model</p> <p><b>Lvl</b> coupling facility microcode level</p> <p><b>Dyn</b> Dynamic CF dispatching status (ON, OFF, or THIN). THIN indicates that coupling thin interrupts are enabled for the coupling facility (only for CFLEVEL 19 or higher).  <b>Note:</b> Dynamic CF dispatching is provided by PR/SM and available to all CF partitions with shared engines. It allows the installation to limit the impact of polling for CFs with low activity rates. The amount of CP resource used by the CF is reduced. There is, however, a performance trade-off when working with dynamic dispatching: though the CPU resource consumed by the CF is reduced, the responsiveness of the CF partition is also reduced. In Parallel Sysplex environments with a CFLEVEL 19 or higher, it is recommended to enable coupling thin interrupts for shared-engine coupling facilities.</p>
Processor	<p>The following information about the processors within the coupling facility is provided:</p> <p><b>Util%</b> Percentage of processor utilization by the coupling facility.</p> <p>In case of a stand-alone coupling facility, the utilization of the individual CPs should be approximately the same. In a PR/SM environment where this CP is shared with other partitions, the utilization is the logical utilization of the CP (that is, only the utilization by the coupling facility).</p> <p>If the utilization is high, you can take the following actions:</p> <ul style="list-style-type: none"> <li>• In a PR/SM environment, you can dedicate the CP to the integrated coupling facility or assign additional CPs to the partition.</li> <li>• Move structures to a coupling facility with lower utilization.</li> <li>• Consider additional or larger coupling facilities.</li> </ul> <p><b>Def</b> Number of logical processors defined for the coupling facility.</p> <p><b>Shr</b> Number of shared processors defined for the coupling facility.</p> <p><b>Wgt</b> Average weight of shared logical processors. This value is not displayed if no shared processors are assigned to this CF.</p> <p><b>Eff</b> Number of effective available logical processors in a shared environment. This value is only useful in a CFCC environment. CFCC measures the time of real command execution as well as the time waiting for work. The reported value shows the ratio of the LPAR dispatch time (CFCC execute and wait time) to the RMF MINTIME length.</p> <p>For example, if a CFCC CEC contains 6 logical processors, and the measured CF LPAR has two logical processors and is limited at 5%, the number of effective logical processors is 0.3.</p>
Req Rate	<p>The sum of all requests (internal and external) that utilize the subchannel. Specifically:</p> <ul style="list-style-type: none"> <li>• External requests to send/receive data on behalf of a structure. The sum of synchronous and asynchronous requests completed against any structure within this coupling facility per second. This includes requests that changed from synchronous to asynchronous.</li> <li>• Internal requests that utilize the subchannels (but are not aggregated by the structure).</li> </ul> <p>The value is reported as '&lt; 0.1' if the rate is greater than 0 but would be rounded to 0.</p>
Storage Size	The total amount of coupling facility storage in bytes, including both allocated and available space.
Storage Avail	The amount of coupling facility space in bytes that is not allocated to any structure, not allocated as dump space, and not allocated as augmented space.
SCM Size	The total amount of coupling facility storage class memory in bytes which may be concurrently used as structure extensions.
SCM Avail	The total amount of available storage class memory in bytes.

## CFSYS - Coupling Facility Systems Report

The Coupling Facility Systems (CF Systems) report (CFSYS) gives you information about the distribution of coupling facility requests among the systems and about the activities in the subchannels and paths attached to the coupling facilities in the sysplex.

Using the CFSYS report, for each coupling facility, you see their activity and all connected systems. High activity values are indicators for contention and possible bottlenecks in the configuration. The pop-up panel with the details provides information about the configuration and you see the path IDs for all channels which are connecting each coupling facility with a system. You can use the CHANNEL report to get the link utilization for each path.

### How to request this report

To request the Coupling Facility Systems report, select **S** from the Primary Menu and then select **6** on the Sysplex Report menu (shown in Figure 5 on page 24), or enter the following command:

```
CFSYS [cfname]
```

### Contents of the report

Command ==>		RMF V2R2	CF Systems	- TRXPLEX				Line 1 of 12 Scroll ==> CSR			
Samples:	60	Systems:	2	Date:	09/28/16	Time:	15.55.00	Range:	30	Sec	
CF Name	System	Subchannel Delay %	Busy %	-- Paths -- Avail Delay %	-- Sync --- Rate	Avg Serv	----- Rate	Async Avg Serv	----- Chng %	----- Del %	
X7CFH89	R70	0.0	0.0	4	0.0	0.8	20	53.0	94	0.0	0.0
	R71	0.0	0.0	4	0.0	0.7	17	54.2	105	0.0	0.0
X7CFP87	R70	0.0	0.1	4	0.0	1170	22	135.7	51	0.0	0.0
	R71	0.0	4.4	4	0.0	246K	5	193.4	37	0.0	0.0
X7CFP89	R70	0.0	0.0	4	0.0	451.3	5	3.5	92	0.0	0.0
	R71	0.0	0.1	4	0.0	229.5	22	224.2	55	0.0	0.0

Figure 31. CFSYS Report

There is no graphic version of this report available.

If you place the cursor on any of the lines with coupling facility systems values, a pop-up panel appears showing details for the subchannels and paths.

```

RMF V2R2 CF Systems - TRXPLEX Line 1 of 12
Command ==> Scroll ==> CSR
Samples: 60 Systems: 2 Date: 09/28/16 Time: 15.55.00 Range: 30 Sec
CF
RMF Coupling Facility - Subchannels and Paths
Press Enter to return to the Report panel.
X7C Details for System : R70
Coupling Facility : X7CFP89
X7C Subchannels Generated : 28
In Use : 28
X7C Max : 128
Channel Path Details:
X7C ID Type Operation Mode Deg Distance CHID AID Port --IOP IDs--
More: +
C4 CIB 1x IFB HCA3-0 LR N <1 0708 000D 01 02
C5 CIB 1x IFB HCA3-0 LR N <1 0709 000D 01 02
C6 CIB 1x IFB HCA3-0 LR N <1 070A 001D 03 08
C7 CIB 1x IFB HCA3-0 LR N <1 070B 001D 03 08
EE ICP
F0 CFP 2Gbit N 1.4 01F0 10
F1 CFP 1Gbit N 1.3 01F1 10
F1=Help F2=SplitScr F3=End F6=RMFHelp F7=Backward
F8=Forward F9=SwapScr F12=Return
F1=HELP F2=SPLIT F3=END F4=RETURN F5=RFIND F6=TOGGLE
F7=UP F8=DOWN F9=SWAP F10=BREF F11=FREF F12=RETRIEVE

```

Figure 32. CFSYS Report - Subchannels and Paths

### Field descriptions

Table 12. Fields in the CFSYS Report

Field Heading	Meaning
<b>Note:</b> Each rate is reported as '<0.1' if the value is greater than 0 but rounded to 0.	
CF Name	Coupling facility name.
System	Name of the system attached to the coupling facility (from IEASYSxx Parmlib member, SYSNAME parameter).
Subchannel Delay %	The percentage of all coupling facility requests z/OS had to delay because it found all coupling facility subchannels busy.  If this percentage is high, you should first ensure that sufficient subchannels are defined (see MAX field below).  If there are sufficient subchannels and this percentage is still high, it indicates either a coupling facility path constraint or internal coupling facility contention.
Subchannel Busy %	Percentage of the coupling facility subchannel utilization. This value is calculated from the sum of synchronous and asynchronous coupling facility request times related to the MINTIME and to the number of subchannels.
Paths Avail	Number of physical paths (coupling facility channels) available to transfer coupling facility requests between this system and the coupling facility.

## Mon III - CFSYS

Table 12. Fields in the CFSYS Report (continued)

Field Heading	Meaning
Paths Delay %	<p>Percentage of all coupling facility requests that were rejected because all paths to the coupling facility were busy. This value can exceed 100% if requests encounter a <i>path busy</i> condition more than once.</p> <p>A high percentage results in elongated service times which is a reduction of the capacity of the sending processor. If coupling facility channels are being shared among PR/SM partitions, the contention could be coming from a remote partition.</p> <p><b>Identifying path contention:</b> There can be path contention even when this count is low. In fact, in a non-PR/SM environment where the subchannels are properly configured, Subchannel Busy, not Path Busy, is the indicator for path contention. If Path Busy is low but Subchannel Busy is high, it means z/OS is delaying the coupling facility requests and in effect gating the workload before it reaches the physical paths. Before concluding you have a capacity problem, however, be sure to check that the correct number of subchannels is defined in the I/O generation (see Subchannel Max).</p> <p><b>PR/SM environment only:</b> If coupling facility channels are being shared among PR/SM partitions, Path Busy behaves differently. Potentially, you have many subchannels mapped to only a few coupling facility command buffers. You could have a case where the subchannels were properly configured (or even under-configured), Subchannel Busy is low, but Path Busy is high. This means the contention is due to activity from a remote partition.</p> <p><b>Possible actions:</b> Dedicate the coupling facility links on the sending processor or add additional links.</p>
Sync Rate	Number of hardware operations per second that started and completed synchronously to the coupling facility on behalf of connectors from this system.
Sync Avg Serv	Average time in microseconds required to satisfy a synchronous coupling facility request on behalf of connectors from this system.
Async Rate	Number of hardware operations per second that started and completed asynchronously to the coupling facility on behalf of connectors from this system.
Async Avg Serv	Average time in microseconds required to satisfy an asynchronous coupling facility request on behalf of connectors from this system. This value also includes operations that started synchronously but completed asynchronously.
Async Chng %	<p>Percentage of asynchronous requests for this structure that changed from synchronous to asynchronous because the requests could not be serviced as synchronous request. This field reports only those requests which were changed due to a subchannel busy condition and can be used as an indicator of a shortage of subchannel resources.</p> <p>Request conversions caused by heuristic sync/async algorithms used to optimize the coupling efficiency of workloads using the CF are not included.</p>
Async Del %	Percentage of asynchronous hardware operations on behalf of connectors from this system being delayed by either subchannel contention or dump serialization.

Table 13. Fields in the CFSYS Report - Subchannels and Paths

Field Heading	Meaning
<i>Subchannels and Paths</i>	



Table 13. Fields in the CFSYS Report - Subchannels and Paths (continued)

Field Heading	Meaning
Subchannels	<p>Subchannel configuration data.</p> <p><b>Generated</b> Number of subchannels generated by the I/O configuration. This could be more than the number z/OS can optimally use for coupling facility requests.</p> <p><b>In Use</b> Number of subchannels z/OS is currently using for coupling facility requests.</p> <p><b>Max</b> Maximum number of coupling facility subchannels z/OS can optimally use for coupling facility requests.</p> <p>The limit is calculated by z/OS to be the number of physical paths to the coupling facility times the number of command buffer sets per path. It represents the number of parallel requests the coupling facility configuration can handle.</p> <p>If this number is less than the subchannels generated by the I/O configuration, you should reduce the number of coupling facility subchannels in the I/O to match this number. Over-specifying subchannels causes unnecessary storage usage and can cause a high number of rejected coupling facility requests due to path busy.</p>
<b>Channel Path Details</b>	
<b>Note:</b> If the hardware cannot provide values for a measurement, the field remains blank.	
ID	The hexadecimal identifier of a channel path (CHPID) that is connecting a system with the coupling facility. The physical path utilization for these coupling facility links is shown in the CHANNEL report.
Type	Channel path type.
Operation Mode	<p>Channel path operation mode. It describes the data rate bandwidth, protocol, and adapter type of the channel path.</p> <p>A data rate of, for example, 1GBIT denotes a rate of 1.0625 gigabit per second.</p> <p>A bandwidth of, for example, 12X denotes a twelve-fold bandwidth.</p> <p>Protocols:</p> <ul style="list-style-type: none"> <li>• IFB – InFiniBand</li> <li>• IFB3 – InFiniBand 3</li> <li>• GEN3 – PCIe third generation protocol</li> </ul> <p>Adapter types:</p> <ul style="list-style-type: none"> <li>• HCA2-O – Host Channel Adapter2-optical</li> <li>• HCA2-O LR – Host Channel Adapter2-optical long reach</li> <li>• HCA3-O – Host Channel Adapter3-optical</li> <li>• HCA3-O LR – Host Channel Adapter3-optical long reach</li> <li>• PCIE-O SR – Peripheral Component Interconnect Express short reach</li> </ul> <p>Unknown operation mode:</p> <ul style="list-style-type: none"> <li>• Unknown</li> </ul>
Deg	Character Y in this column indicates that the channel path is operating at reduced capacity (degraded) or not operating at all.
Distance	<p>Estimated distance in kilometers. The value is calculated as follows:</p> <p style="text-align: center;">Average round-trip path time in microseconds ----- 10 microseconds / kilometer</p> <p>A value of zero means that the time was not measured.</p>
PCHID	Physical channel identifier.
AID	The hexadecimal coupling adapter ID associated with the CHPID.
PORT	The hexadecimal port number associated with the CHPID.
IOP IDS	The hexadecimal identifiers of I/O processors (System Assist Processors) to which the channel path is accessible.

## CHANNEL - Channel Path Activity Report

The Channel Path Activity report (CHANNEL) gives you information about channel path activity for all channel paths in the system. The report contains data for every channel path that is online during data gathering.

For all channels that are managed by Dynamic Channel Path Management (DCM), additional information is available. DCM allows an installation to identify channels which they wish to be managed dynamically. These channels are not assigned permanently to a specific control unit, but belong to a pool of channels. Based on workload requirements in the system, these channels will be assigned dynamically by DCM. On top of the report, there is a consolidated data section for managed channels displaying the total number of channel paths for each type and the average activity data. The character **M** as suffix of the acronym for the channel path type is an indicator that the channel is managed by DCM.

You can use channel path activity information together with I/O device activity and I/O queuing activity information to identify performance bottlenecks associated with channel paths.

To find out which logical control unit is using the channel, look in the I/O Queuing Activity report. From there you can go to check device response times. For example, if a channel path to a device shows excessive use, you could define additional paths to the device or introduce a different job mix to produce better performance.

### How to request this report

To request the Channel Path Activity report, select **3** from the Primary Menu and then select **12** on the Resource Report Selection Menu (shown in Figure 8 on page 27), or enter the following command:

```
CHANNEL
```

### Special considerations of report output

You can obtain the report whether or not a Monitor I session measuring channel path activity is active.

Data for total utilization and partition utilization is gathered independently. Because the internal interval used to gather this data is a few seconds, the total utilization and the sum of the partition's utilization sharing that channel might differ if a short RMF interval is specified. If the interval is too small or the appropriate data cannot be gathered, dashes (---) will be reported instead of data. Please refer to the information APAR II05151 for a list of channel types for which channel utilization data is not gathered.

## Contents of the report

```

RMF V2R2 Channel Path Activity
Command ==>>>
Line 1 of 69
Scroll ==>> HALF

Samples: 60      System: CB88 Date: 09/28/16 Time: 08.00.00 Range: 60 Sec

Channel Path      Utilization(%)  Read(B/s) Write(B/s)  FICON OPS  zHPF OPS
ID No  G  Type  S  Part Tot  Bus  Part Tot  Part Tot  Rate Actv  Rate Actv

   4    *CNCSM      0.1 0.3
   4    *FC_SM      0.0 0.0 0.0    0  0    0  0    0  0    0  0
  12      OSD  Y      0.0 0.0 0.0    2K 19K    0  0
  14      OSD  Y      0.0 0.0 0.0    5K 531K 511K 514K
  16      OSD  Y      0.4 1.3 0.0   511K  5M    3M  5M
  20      CTC_S Y      0.0 0.0
  27      CNC_S Y      0.0 0.0
  2B      CNC_S Y      1.3 5.2
  2C      CNC_S Y      0.2 0.5
  30     5 FC_S  Y      0.0 32.5 8.9   205 52M  205 235K  186  1    0  0
  31     5 FC_S  Y      0.0 33.3 8.5   429 50M  330 249K  185  2    0  0
  37     4 FC_S  Y      0.0 0.5 0.1    0 619K    0 42K   24  1    0  0
  38     4 FC_S  Y      0.0 0.5 0.1    0 613K    0 73K   30  2    0  0
  39     4 FC_S  Y      0.0 0.1 0.0   374 23K    0 31K    8  1    0  0
  3A     4 FC_S  Y      0.0 0.1 0.0   365 21K    0 32K    7  1    0  0
  3E     4 FC_S  Y      0.0 0.0 0.0    0 10K    0  1K    3  1    0  0
  7C      CNCSM Y      0.3 0.8
  7D      CNCSM Y      0.1 0.1
  81     3 FC_S  Y      1.1 14.2 3.2  801K 18M  147K  1M   738  2  132  1
  82     5 FC_S  Y      0.1 0.4 0.2   37K 870K  28K  86K    7  1   36  1
  83     5 FC_S  Y      0.0 0.4 0.2   36K 887K  27K  83K    8  1   36  1
  84     4 FC_S  Y      0.0 0.0 0.0    25 101    0  0    0  1    0  0
  85     3 FC_S  Y      0.4 6.8 0.7   62K  2M   61K  1M   420  1  157  1
  8C     3 FC_S  Y      0.6 10.8 1.4  344K  6M   61K 801K  720  2    0  0
  A6     5 FC_SM Y      0.0 0.0 0.0    0  0    0  0    0  0    0  0
  B6     5 FC_SM Y      0.0 0.0 0.0    0  0    0  0    0  0    0  0
  E0      IQD  Y      0 315K
  E1      IQD  Y      0  0
  E2      IQD  Y      0  0
  E3      IQD  Y      0  0
    
```

Figure 33. CHANNEL Report

The graphic form of the Channel Path Activity report shows the percentage of total utilization for each channel.

### Field descriptions

Table 14. Fields in the CHANNEL Report

Field Heading	Meaning
Channel Path ID	Hexadecimal channel path identifier (CHPID).
Channel Path No	For each channel type which is managed by DCM, a summary line is shown with the average values for all channels in this group. These summary lines are characterized by an '*' preceding the channel path type, and the number of channels of the group is displayed in column No.

## Mon III - CHANNEL

Table 14. Fields in the CHANNEL Report (continued)

Field Heading	Meaning
Channel Path G	<p>Generation.</p> <p>The generation is used to differentiate between channels of the same channel type, when one has significant differences from the other. Newer generations with significant differences (for example, the channel throughput) are indicated by a number (1, 2, ...).</p> <p>For example, for a FICON channel, a number 1 indicates that the channel has an auto-negotiated throughput of 1Gbit/sec, or a number 4 indicates a throughput of 2Gbit/sec on a FICON Express4 card or a FICON Express2 card.</p>
Channel Path Type	<p>Type of channel path.</p> <p>You may issue the console command <code>D M=CHP(xx)</code> to see an explanation of the channel path type.</p> <p>If RMF encounters an error while processing the type, this field is blank. RMF continues to measure channel path activity. Check the operator console for messages.</p>
Channel Path S	<p>The indication of whether a channel path is defined as shared between one or more logical partitions. Y indicates that the channel path is shared. A blank indicates it is not shared.</p>
<p><b>Note:</b></p> <ol style="list-style-type: none"> <li>1. On a machine running in LPAR mode, but with only one LPAR defined, the <i>Part</i> columns for the <i>Read</i>, <i>Write</i> and <i>Utilization</i> fields display a zero value for channels of type FC (FICON).</li> <li>2. When Channel Path Measurement Facility (CPMF) is not available, for example, on z/OS systems running as z/VM guests, RMF uses sampled data from SRM so that the reported channel utilization is only an approximate value. With increasing channel speed, the channel utilization value becomes more and more inaccurate. Therefore, in such cases, RMF does not provide accurate values of FICON channel utilization.</li> </ol> <p>Beginning with z990 processors, the channel data from SRM is no longer available. As a result, the channel utilization data on a z/OS system running as z/VM guest, is reported as '-----'</p>	
Utilization (%) Part	<p>The channel path utilization percentage for an individual partition. RMF uses the values provided by CPMF.</p> $\text{Part Utilization (\%)} = \frac{\text{Channel Path Busy Time}}{\text{Channel Path Elapsed Time}} * 100$ <p>For channels like FICON, OSA Express, or OSA Direct Express, which are running in extended CPMF mode, the calculation is as follows:</p> $\text{Part Utilization (\%)} = \frac{\text{LPAR \# of Channel Work Units}}{\text{Max \# of Channel Work Units} * \text{Channel Path Elapsed Time}} * 100$ <p>For some channels like OSAEGbE, FICON EXPRESS/EXPRESS2, this value reflects the microprocessor utilization.</p> <p>For hipersockets, this value is not available.</p>

Table 14. Fields in the CHANNEL Report (continued)

Field Heading	Meaning
Utilization (%) Tot	<p>The channel path utilization percentage for the CPC during an interval.</p> <p>For processors earlier than z990 and shared channels in LPAR mode, where CPMF is not available, the calculation is:</p> $\text{Total Utilization (\%)} = \frac{\text{\# SRM Observations of Channel Path Busy}}{\text{\# SRM samples}} * 100$ <p>For unshared channels in LPAR mode, the value for total utilization is the same as partition utilization.</p> <p>For channels like for example FICON, OSA Express, or OSA Direct Express, which are running in extended CPMF mode, the calculation is as follows:</p> $\text{Total Utilization (\%)} = \frac{\text{Total \# of Channel Work Units}}{\text{Max \# of Channel Work Units} * \text{Channel Path Elapsed Time}} * 100$ <p>For some channels like OSAEGbE, FICON EXPRESS/EXPRESS2, this value reflects the microprocessor utilization.</p> <p>For hipersockets, this value is not available.</p>
Utilization (%) Bus	<p>Percentage of bus cycles, the bus has been found busy for this channel in relation to the theoretical limit.</p> <p>For OSAEGbE, the value reflects the PCI bus utilization.</p> <p>For hipersockets, this value is not available.</p>
Read(B/s)	<p><b>Part</b> Data transfer rates from the control unit to the channel for this partition.</p> <p><b>Total</b> Data transfer rates from the control unit to the channel for the CPC.</p> <p>For hipersockets, this value is not available.</p>
Write(B/s)	<p><b>Part</b> Data transfer rates from the channel to the control unit for this partition.</p> <p><b>Total</b> Data transfer rates from the channel to the control unit for the CPC.</p>
FICON OPS	<p><b>Rate</b> Number of native FICON operations per second.</p> <p><b>Actv</b> The average number of native FICON operations that are concurrently active during the report interval.</p>
zHPF OPS	<p><b>Rate</b> Number of zHPF (High Performance FICON) operations per second.</p> <p><b>Actv</b> The average number of zHPF operations that are concurrently active during the report interval.</p>

### Monitor III Utility fields

You can use the Monitor III Utility to customize the CHANNEL report in a way that the following additional values are shown:

Table 15. Additional Fields in the CHANNEL Report

Field Name	Meaning
CHACFDFR	Number of deferred native FICON operations per second that could not be initiated by the channel due to a lack of available resources.
CHACXDFR	Number of deferred zHPF operations per second that could not be initiated by the channel due to lack of available resources.
CHACNET1	Physical-network identifier (PNET ID) of first channel path port.
CHACNET2	Physical-network identifier (PNET ID) of second channel path port.
The following fields are only available for HiperSockets:	

Table 15. Additional Fields in the CHANNEL Report (continued)

Field Name	Meaning
CHACTMVC	Total message sent rate
CHACTSVC	Total message sent size
CHACTFVC	Total receive fail rate
CHACPMVC	Rate of messages sent by this LPAR
CHACPSVC	Average size of messages sent by this LPAR
CHACPFVC	Rate of messages received by this partition that failed due to an unavailable buffer. The value could indicate that more receive buffers are required.
CHACSFVC	Rate of messages sent by this partition that failed.

## CPC - CPC Capacity Report

The CPC Capacity (CPC) report provides the capability to monitor values that are relevant for software pricing as well as partition related processor activities.

Prior to z/OS, software products were typically priced based on the computing capacity of the central processor complex (CPC) on which the software is running. With z/OS running on a zSeries hardware, charging can be based on the capacity defined for workloads. WLM LPAR CPU management offers the support that allows pricing based on partition capacity. It will ensure that the average CPU consumption of a partition does not exceed a defined capacity value (in millions of unweighted CPU service units per hour - MSU/h) over a defined period of time. WLM allows the actual workload to rise above the defined MSU limit, but takes care that the four-hours average stays below. This is done by dynamically turning capping on and off. Prerequisites are uncapped partitions with shared CPs.

With this CPC capacity report, you can compare the defined capacity limits against the actual MSU consumption for all partitions of the CPC. In addition, the report contains MSU information related to the last four hours, for the partition RMF is running in, which clearly shows if the defined capacity limit is appropriate to the workload running in this partition or if WLM has to cap this partition's workload because the defined limit was set too low.

### How to request this report

To request the CPC Capacity report, select **1** from the Primary Menu and then **3** on the Overview Report Selection Menu (shown in Figure 6 on page 25), or enter the following command:

```
CPC
```

### Contents of the report

The *CPC Capacity* report provides:

- **Header information** which offers MSU related values with the scope of that partition which requested the report as well as the processor type, model, and capacity.
- **Partition data** which displays the values for all partitions belonging to the CPC. If multithreading is enabled for a processor type (LOADxx PROCVIEW CORE parameter is in effect), processor data is reported at core granularity.

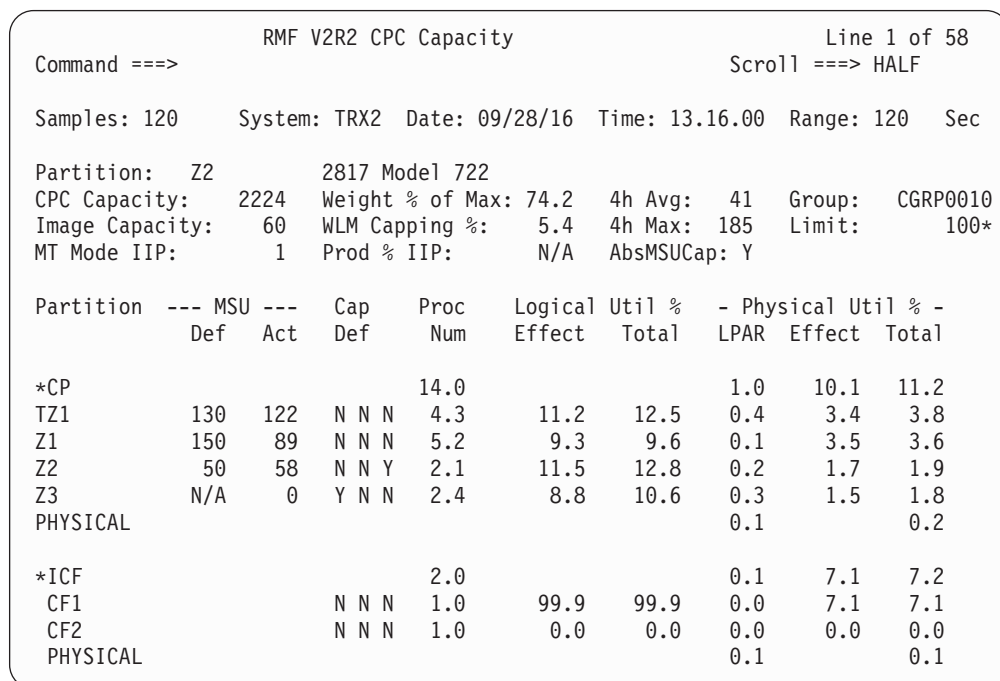


Figure 34. CPC Capacity report

### Field descriptions

Table 16. Fields in the CPC Capacity Report

Field Heading	Meaning
All MSU-related values are measured in MSU/h (millions of service units per hour).	
<b>Values for the partition which requested the report</b>	
Partition	Partition name.
Processor/Model	Processor family and model number of the measured system.
CPC Capacity	Effective processor capacity available to the central processor complex (CPC), measured in MSU/h.
Image Capacity	Processor capacity available to the z/OS image (partition) which requested the report, measured in MSU/h. The field is calculated as minimum of the following capacities: <ul style="list-style-type: none"> <li>the capacity based on the partition's logical CP configuration (both online and offline)</li> <li>the defined capacity limit of the partition, if available (image softcap)</li> <li>the capacity limit of the related WLM capacity group, if the partition belongs to a capacity group.</li> <li>the absolute physical hardware capping limit.</li> <li>the capacity based on the hardware group capping limit.</li> </ul>
MT Mode IIP	The multithreading mode for processor type zIIP designates the number of active threads for each online logical zIIP core. If this value is greater than 1, multithreading becomes effective for zIIP cores.  N/A is shown if the LOADxx PROCVIEW CPU parameter is in effect or no IIP is currently installed or online.
Prod % IIP	The multithreading IIP core productivity represents the percentage of the maximum IIP core capacity that was used while the IIP cores were dispatched to physical hardware.  When this value equals 100% in multithreading mode, all threads on all IIP cores that were configured ONLINE for the complete MINTIME are being used. If the LOADxx PROCVIEW CPU parameter is in effect or no IIP is currently installed or online, no core productivity is calculated and N/A is reported.

Table 16. Fields in the CPC Capacity Report (continued)

Field Heading	Meaning
Weight % of Max	Average weighting factor in relation to the maximum defined weighting factor for this partition.  With 'Initial Capping ON', which the operator can set on the Hardware Management Console, this value is not available and therefore, this field shows '****' in this case.
WLM Capping %	Percentage of time when WLM capped the partition because the four-hours average MSU value exceeds the defined capacity limit.
4h Avg	Average value of consumed MSU/h during the last four hours.
4h Max	Maximum value of consumed MSUs during the last 4 hours (retrieved from 48 sample intervals of five minutes). This value can be greater than the defined capacity.
AbsMSUCap	Absolute MSU capping is active for the partition: Y or N.
Group	Name of the partition's capacity group. If the partition does not belong to a capacity group, N/A is displayed.
Limit	Capacity limit (in MSUs) defined for the partition's capacity group.  An '*' following the limit value indicates that this partition started to be a member of this capacity group less than four hours ago. This partition will have a different view of unused group capacity and, therefore, might cap differently than existing group members.
<b>Values for all configured partitions are grouped by general and special purpose processor types. The term logical processor refers to a logical core if the LOADxx PROCVIEW CORE parameter is in effect.</b>	
Partition	Partition name. <b>Notes:</b> <ol style="list-style-type: none"> <li>Partitions identified by the name PHYSICAL are not configured partitions. Data reported in these lines includes the time during which a physical CPU was busy, but the time could not be attributed to a specific logical partition.</li> <li>The summary lines (for example, *CP or *ICF) show the total percentages for the indicated processor type.</li> <li>Starting with z9 processors, IFLs (Integrated Facility for Linux) and zAAPs are reported separately and no longer as ICFs (Internal Coupling Facility).</li> </ol>
MSU	Millions of unweighted CPU service units per hour:  <b>Def</b> Defined MSU capacity limit for the partition.  <b>Act</b> Actual MSU consumption.  These values are only provided for general purpose processors.
Cap Def	The hardware capping option of the partition. Each Cap Def value is a three position character string denoting which hardware capping mechanisms have or have not been applied in the logical partition controls of the Hardware Management Console (HMC) for the partition. The values in the first, second and third position of the string are either Y (Yes) or N (No) and have the following meaning:  The first character (Y or N) indicates whether "Initial Capping ON" has been set.  The second character (Y or N) indicates whether an absolute physical hardware capping limit (maximal number of CPUs) has been defined.  The third character (Y or N) indicates whether an absolute hardware group capping limit (maximal number of CPUs) has been defined.  An asterisk (*) to the right of a value indicates that the capping status changed during the report interval.
Proc Num	The number of logical processors which were online during the report interval.
Average Processor Utilization Percentages. <ul style="list-style-type: none"> <li>The average utilization of logical processors is based on the total online time of all logical processors assigned to the partition.</li> <li>The average utilization of physical processors is based on the total interval time of all physical processors.</li> </ul>	
Logical Util % - Effect	The average partition effective dispatch time percentage.  $\frac{\text{Effective Dispatch Time}}{\sum \text{Online Times}} * 100$



Table 16. Fields in the CPC Capacity Report (continued)

Field Heading	Meaning
Logical Util % - Total	The average partition total dispatch time percentage. $\frac{\text{Total Dispatch Time}}{\sum \text{Online Times}} * 100$
Physical Util % - LPAR	The average LPAR management time percentage. $\frac{\text{Total Dispatch Time} - \text{Effective Dispatch Time}}{\# \text{ Physical Processors} * \text{Range Time}} * 100$ <p>The calculation for the PHYSICAL partition is:</p> $\frac{\text{Time PHYSICAL}}{\# \text{ Physical Processors} * \text{Range Time}} * 100$ <p>Time PHYSICAL is the time that could not be attributed to a specific logical partition but was used by PR/SM to control the physical processor (LPAR management time).</p>
Physical Util % - Effect	The effective utilization of the physical processor resource by the partition. $\frac{\text{Effective Dispatch Time}}{\# \text{ Physical Processors} * \text{Range Time}} * 100$
Physical Util % - Total	The total utilization of the physical processor resource by the partition. $\frac{\text{Total Dispatch Time}}{\# \text{ Physical Processors} * \text{Range Time}} * 100$ <p>The Total Dispatch Time for the PHYSICAL partition includes the time during which a physical CPU was busy, but the time could not be attributed to a specific logical partition. This time includes the time PR/SM was controlling the physical processor (LPAR management time), as well as any other time the processor was busy for any reason such as managing coupling facility traffic.</p>

## DELAY - Delay Report

The Delay report allows you to determine which system resources are causing delays for jobs or job groups, and to what extent the jobs are delayed.

The report gives you information about job delay for every type of delay that RMF monitors. This includes processor delay (PRC), device delay (DEV), storage delay (STR), subsystem delay (SUB), operator delay (OPR), and enqueue delay (ENQ). RMF provides a detail report for each of these delays except OPR. Operator delay includes message, mount, and quiesce requests. SUB is divided into an HSM, JES, and XCF detail report. The names of the detail reports correspond to the names that appear in the Delay report.

### How to request this report

To request the Delay report, select **1** from the Primary Menu, and then select **4** on the Overview Report menu (shown in Figure 6 on page 25) or enter the following command using the format:

```
DELAY [job_class,service_class]
```

For example, to get a Delay report for TSO service class TSOPRIME, enter:

```
DELAY T, TSOPRIME
```

## Contents of the report

RMF V2R2 Delay Report											Line 1 of 58											
Command ==>											Scroll ==> HALF											
Samples: 120											System: MVS1			Date: 09/28/16			Time: 12.00.00			Range: 120 Sec		
Name	Service	CX Class	Cr	WFL %	USG %	DLY %	IDL %	UKN %	---- PRC	% Delayed for	---- STR	---- SUB	---- OPR	---- ENQ	Primary Reason							
*SYSTEM				49	1	1	62	36	0	0	0	0	0	0								
*TSO				56	1	1	95	2	0	0	0	1	0	0								
*BATCH				39	2	4	0	94	1	0	0	4	0	0								
*STC				40	0	1	51	48	0	0	0	0	0	0								
*ASCH					0	0	0	0	0	0	0	0	0	0								
*OMVS					0	0	0	100	0	0	0	0	0	0								
*ENCLAVE					0	0	0	0	0	N/A	0	N/A	N/A	N/A								
JES2	S	SYSSTC		0	0	1	0	99	0	1	0	0	0	0	SCLSP4							
BMAI	T	PRDTSO	S	16	9	66	13	4	0	1	0	65	0	0	HSM							
HSM	S	STCCMD	C	30	26	62	0	23	0	1	0	0	50	11	Mount							
HIRW2	B	BATCHMED		35	6	14	0	1	2	1	0	12	0	0	HSM							
TCPNET	SO	SYSSTC	SC	60	3	2	0	97	2	0	0	0	0	0	NET							
*MASTER*	S	SYSTEM		67	3	2	0	95	0	2	0	0	0	0	M00202							

Figure 35. DELAY Report

The graphic form of this report shows the percent of time that each user spent delayed for the above resources.

### Field descriptions

Table 17. Fields in the DELAY Report

Field Heading	Meaning
Name	Name of the job, job group or enclave.  The enclave names, starting with the letters ENC, and belonging to class E, are created dynamically by RMF. You cannot use the names to track a particular enclave through different time ranges. However, the enclave token is used when combining multiple set-of-samples, so that data are combined only for the same individual enclaves, thus providing consistent data. N/A is shown if the value does not apply to enclaves.
CX	Abbreviation for the address space types as follows: <b>S</b> Started task <b>T</b> TSO <b>B</b> Batch <b>A</b> ASCH <b>O</b> OMVS <b>?</b> Data is missing or not valid.  Or it can indicate an enclave: <b>E</b> Enclave  For summary entries, this field is blank.  An <b>O</b> as second character indicates that an OMVS process exists for this address space.
Service Class	The name of the service class that a specified job has been running in.  If a job changes its service class during the report interval, RMF displays eight asterisks (*****) instead of the service class name. If the service class is not available, RMF displays eight dashes (-----).

Table 17. Fields in the DELAY Report (continued)

Field Heading	Meaning
Cr	An indication whether WLM managed the address space as <i>storage critical</i> and/or <i>CPU critical</i> during the report interval. <b>C</b> CPU critical <b>S</b> Storage critical <b>SC</b> Both storage and CPU critical
WFL %	The workflow percentage of the job or job group. "Address space workflow (%)" on page 12 shows the formula used to calculate this value.
USG %	The using percentage for the job or job group. "Address space using (%)" on page 14 shows the formula used to calculate this value.
DLY %	The delay percentage for the job or job group. See "Address space delay (%)" on page 14 for more information.
IDL %	<p>The idling percentage for a job or job group. Jobs in terminal wait, timer wait, or waiting for job selection by JES are in an <i>idling</i> state if they are not using the processor or devices and are not delayed for any monitored reason.</p> <p>Jobs classified as in terminal wait meet all of the following conditions:</p> <ul style="list-style-type: none"> <li>• They are not found using any monitored resource</li> <li>• They are not found delayed for any monitored reason</li> <li>• They are swapped out</li> <li>• They are in terminal wait</li> <li>• They are waiting for a user ready indication before being swapped in.</li> </ul> <p>Jobs classified as in timer wait meet all of the following conditions:</p> <ul style="list-style-type: none"> <li>• They are not using or delayed for a monitored resource.</li> <li>• Their address space is waiting for a timer.</li> </ul> <p>The idling percentage of an address space can vary from 0 to 100%, where 0% indicates that the user is not idling during the report interval, and 100% represents a job that is idle at every sample.</p> <p>The idling percentage for an address space during a refresh period is calculated as follows:</p> $IDL \% = \frac{\# \text{ Idle Samples}}{\# \text{ Samples}} * 100$ <p><b>Idle samples</b>  The number of samples that show the job in an idle state.</p> <p>The idling percentage for a group of address spaces during a range period is calculated as follows:</p> $IDL \% = \frac{\sum \text{ Idle Samples}}{\# \text{ Samples} * \text{Avg} \# \text{ Address Spaces}} * 100$ <p><b>Note:</b> The value reported might include some delay for a non-monitored resource.</p>
UKN %	<p>RMF considers jobs that are not delayed for a monitored resource, not using a monitored resource, or not in an idling state to be in an unknown state.</p> <p>Examples of address spaces in an unknown state are:</p> <ul style="list-style-type: none"> <li>• Idle address spaces that use a non-monitored mechanism for determining when they are active. Most system tasks (STC) show up as unknown when they are idle.</li> <li>• Address spaces waiting for devices other than DASD or tape.</li> </ul> <p>The unknown state percentage for an address space can vary from 0 to 100%, where 0% indicates that the state was always known during the report interval and 100% represents a job in an unknown state throughout the report interval.</p>

## Mon III - DELAY

Table 17. Fields in the DELAY Report (continued)

Field Heading	Meaning
% Delayed for	<p>The percentage that each defined resource contributes to the overall delay of the job or job group.</p> <p>The overall delay value DLY % may exceed the sum of the reported resource delay values, because there are other resources which contribute to the overall delay, such as WLM capping delay.</p> <p>If the percentages add up to more than DLY %, there is an overlap of delay states.</p> <p>The defined resources that can delay the job or job group are as follows:</p> <p><b>PRC</b> The job or job group has ready work on the dispatching queue, but it is not being dispatched.</p> <p><b>DEV</b> The job or job group is delayed for a DASD or tape.</p> <p><b>STR</b> The job or job group is waiting for a COMM, LOCL, SWAP, XMEM, HIPR or VIO page, or is on the out/ready queue. See the Storage Delay report.</p> <p><b>SUB</b> The job or job group is delayed for a JES, HSM, or XCF subsystem request.</p> <p><b>OPR</b> The job or job group is delayed by a message or a mount request or a quiesce. Quiesce means that the operator has quiesced the address space. A quiesced address space can show unexpected data:</p> <ul style="list-style-type: none"> <li>• A swappable address space will be swapped out, thus it can be OUTR and show storage delays.</li> <li>• A non-swappable address space will get lowest priority, thus it can show CPU delay, paging delay, or other delays, and even some USG % from time to time depending on the load on the system.</li> </ul> <p><b>ENQ</b> The job or job group is waiting to use an enqueued (reserved) resource.</p>

Table 17. Fields in the DELAY Report (continued)

Field Heading	Meaning
Primary Reason	<p>Reported only for a specific job, this field provides additional information about the primary reason for the delay. The contents depend on the resource having the largest % Delayed for value.</p> <p>If the resource with the maximum delay is:</p> <p><b>PRC</b> This field contains the name of the job that used the processor most frequently while the reported job was delayed.</p> <p><b>STR</b> This field identifies the cause of the largest percentage of delay:</p> <p><b>COMM</b> common storage paging (includes shared pages)</p> <p><b>LOCL</b> local storage paging (includes shared pages)</p> <p><b>VIO</b> virtual I/O paging</p> <p><b>SWAP</b> swap-in delay</p> <p><b>OUTR</b> swapped out and ready</p> <p><b>XMEM</b> cross memory address space</p> <p><b>HIPR</b> standard hiperspace paging delays</p> <p><b>DEV</b> This field contains the volume serial number of the device that the reported job was most frequently delayed for.</p> <p><b>SUB</b> This field contains either JES, HSM, or XCF depending on which subsystem is causing the most delay.</p> <p><b>OPR</b> This field contains <i>Message</i> if most of the delay was due to a message or <i>Mount</i> if most of the delay was due to a mount request.</p> <p>The field can contain <i>QUIESCE</i> if the operator quiesced the address space. A quiesced address space can show unexpected data:</p> <ul style="list-style-type: none"> <li>• A swappable address space will be swapped out, thus it can be OUTR and show storage delays.</li> <li>• A non-swappable address space will get lowest priority, thus it can show CPU delay, paging delay, or other delays, and even some USG % from time to time depending on the load on the system.</li> </ul> <p>Cursor-sensitive control on this field gives you the Quiesce delay variation of the Job Delay report.</p> <p><b>ENQ</b> This field contains the major name of the resource most responsible for the delay.</p> <p><b>*ENCLAVE</b> One or more enclaves are holding the processor.</p> <p><b>RG-Cap</b> The job is delayed due to WLM resource capping. That means that</p> <ul style="list-style-type: none"> <li>• either the resource group for which the job is running, has used up its CPU service specified in the WLM policy,</li> <li>• or the work for which the job is running is overachieving its goal. So this work may be capped in order to divert its resources to run discretionary work (see also section 'Using Discretionary Goals' in <i>z/OS MVS Planning: Workload Management</i>).</li> </ul>

### Monitor III Utility fields

You can use the Monitor III Utility to customize the DELAY report. In addition to the delays previously described, you can use the Utility to have the following delay percentages shown.

Table 18. Additional Fields in the DELAY Report

Field Heading	Meaning
% Delayed for	<ul style="list-style-type: none"> <li>• JES delay percentage</li> <li>• HSM delay percentage</li> <li>• XCF delay percentage</li> <li>• Operator mount delay percentage</li> <li>• Operator message delay percentage</li> <li>• Operator quiesce delay percentage</li> <li>• WLM resource group capping delay percentage</li> </ul>

### Cursor-sensitive control on the Delay report

To see all delays for a particular class or summary line (\*SYSTEM, \*TSO, \*BATCH, \*STC, \*ASCH or \*OMVS), use cursor-sensitive control on any name starting with an asterisk (\*) under the name column or on any value in the CX or Service Class columns, to display a subset of the Delay report for that group.

When you use cursor-sensitive control on the \*ENCLAVE summary line, you are shown a subset of individual enclave names.

Using cursor-sensitive control on an enclave name displays a pop-up panel that shows information you extracted from the WLM Enclave Classification Data (ECD) control block. You can use this information to identify the transactions that are processed in the enclave. See “Enclave Classification Attributes” on page 93 for an example.

To see all jobs using or delayed for processor, use cursor-sensitive control on any indicator under USG % to display either the Processor Delays or the Device Delays report, depending on which is contributing more to the delay.

To investigate which jobs or resources are contributing to a delay, use cursor-sensitive control on any indicator under DLY % or % Delayed for to display the related resource report or job delay report.

### Report options

```

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

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

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

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

                Available Service classes
GPMSSERVE  OE          OMVSKERN  STCDEF    TSODEF    SYSTEM    SYSSTC
    
```

Figure 36. DELAY Report Options Panel

The parameters that you specify on this panel (except Summary and Criterion) affect all job-oriented detail delay reports.

**Class** The class of jobs for which you want delay and common storage data reported. For Class, you can request:

- T or TSO
- B or BATCH
- S or STC for started task
- A or ALL for all jobs in the system
- AS or ASCH for ASCH address spaces
- O or OMVS

Your selection for Class applies to all delay and common storage reports and is saved across sessions in the current option set.

**Service Class**

The service class for which you want data reported. For Service Class, you can specify any of the available service classes listed under Available Service Classes.

If the service class you want is not listed, it was not active during the current report interval. If you specify the service class, it will appear on the report when it is available.

Your selection applies to all delay and common storage reports and is saved across sessions in the current option set.

**Summary**

Summary allows you to specify whether you want summary lines for the DELAY report.

To produce one summary line for all jobs in the system and one summary line for each class (TSO, BATCH, STC, ASCH or OMVS), enter ALL for Class and YES for Summary.

To only produce a summary line for one class, group or service class, enter the name for Class and YES for Summary.

Your selection for Summary applies only to the DELAY report and is saved across sessions in the current option set.

**Criterion**

The value (from 0% to 100%) that RMF compares to each job's computed delay value in deciding whether to include the job in the DELAY report.

RMF displays all jobs whose delay values meet or exceed the Criterion.

The value that you specify for Criterion applies only to the DELAY report and is saved across sessions in the current option set.

**Jobs** A YES for JOBS displays the name of all the active jobs in the Class, Group or Service class you specified and any jobname that you previously selected or excluded.

You can use this list to view active jobs in the system and to select and exclude jobs from your report.

**Available Service classes**

The list of available service classes includes all of the service classes that belong to the Class you specified and that had any activity during the current report interval.

Press the END key to make these values active for the session.

**Job Selection/Exclusion Option panel**

If you select YES for Jobs on the Delay Report Options panel, RMF displays a Job Selection/Exclusion panel shown in Figure 37 on page 70.

```

RMF Delay Report Options: DELAY                               Line 1 of 77
Command ==>>>                                               Scroll ==>> HALF

Select (S), exclude (X), or fill-in jobs for report. Press END.

Sel  Jobname      Sel  Jobname      Sel  Jobname      Sel  Jobname      Sel  Jobname
-----
S  *ALL           *MASTER*      520252           EFIBERC          ALISONW
  ALLOCAS        ALPERTA       ALTER2          AMSAQTS          AMYH
  ANDREA         ANDREW        ANN             ARTHUR           ARTI
  ASTER2        AUXCFTH       AULT           BARBARA          BARBIE
  BCOVEN        BEENA        BERNIEP        BERRZA           BETHC
  
```

Figure 37. DELAY Report Job Selection/Exclusion Panel

The Job Selection/Exclusion panel allows you to select or exclude specific jobs from your delay reports.

The panel lists:

- Active jobs in the class and group specified on the Delay Report Options panel.
- All jobs that you previously selected or excluded, selection codes appear to the left of jobs previously selected or excluded.

To select a job for your delay reports, type an s to its left, under SEL; to exclude a job, type an x to its left. (You can select \*ALL for all jobs in the specified class and group and then exclude specific jobs. Similarly, you can exclude \*ALL and then select specific jobs.)

To select several jobs with similar names, use an asterisk (\*) as a "wild card" character under Jobname. For example: to request a report for all jobs starting with A, specify 's' under Sel, 'a\*' under Jobname and ensure that there is an 'x' beside \*ALL.

You can also specify multiple wild card entries, for example, to list all jobs starting with A and all jobs starting with BK, specify:

```

Sel  Jobname      Sel  Jobname
S  A*_____    S  BK*_____
X  *ALL
  
```

To select or exclude a job that is not listed, enter the job name in the top row and the appropriate selection code to its left.

All the jobs might not fit on this panel. Use PF8 and PF7 to scroll through the remaining job names.

## DEV - Device Delays Report

The Device Delays report (DEV) shows jobs delayed by contention for devices. RMF lists the jobs included by descending delay percentages; that is, the job experiencing the most significant delay appears first.

### How to request this report

To request the Device Delays report, select 3 from the Primary Menu and then select 2 on the Resource Report Selection Menu (shown in Figure 8 on page 27), or enter the following command:



DEV [job\_class,service\_class]

For example, to get a Device Delays report for TSO service class TSOPRIME, enter:  
DEV T, TSOPRIME

## Contents of the report

```

RMF V2R2 Device Delays
Command ==>
Line 1 of 57
Scroll ==> HALF

Samples: 100 System: MVS1 Date: 09/28/16 Time: 10.03.20 Range: 100 Sec

Jobname Service DLY USG CON ----- Main Delay Volume(s) -----
C Class % % % % VOLSER % VOLSER % VOLSER % VOLSER
MARYPATM B NRPRIME 70 51 54 70 TSOL11 1 DUMP00
MICHAELL B NRPRIME 39 15 14 39 BPXLK1
MCPDUMP S SYSSTC 36 18 20 36 D24PK2
CHARLESR B NRPRIME 33 13 13 28 BPXLK1 3 HSML02 2 BPXSSK
DFHSM S SYSSTC 30 83 35 10 HSML17 5 SMS026 4 HSMOCD 4 HSMBCD
SHUMA3 T TSOPRIME 18 52 53 13 D83ID0 5 HSML02
DAVEP T TSOPRIME 16 9 10 4 HSM009 3 HSM005 2 HSML06 1 SMS013
CATALOG S SYSTEM 9 15 21 2 CLR007 1 HSM036 1 HSM018 1 HSM011
DB2MDBM1 S SYSSTC 9 7 5 7 DB2MS2 1 DB2MD0 1 DB2MS0
GINNI T TSOPRIME 8 10 9 3 HSML17 2 CLR010 1 HSM032 1 NATPK1
TREVORJ T TSOPRIME 6 10 11 2 HSM022 1 HSM001 1 RESPK1 1 HSM024
RHANSON T TSOPRIME 6 9 8 4 HSML17 1 RESPK1 1 NATPK1
KOCH T TSOPRIME 6 3 3 2 HSML17 1 CLR010 1 HSM018 1 HSM043
RSTSHYSO B NRPRIME 5 8 7 5 HSML17
BEENA T TSOPRIME 5 6 6 4 HSM036 1 HSM020
CRISMAN T TSOPRIME 5 6 3 4 HSML17 1 SMS005
JACKF T TSOPRIME 5 4 1 2 HSML17 2 TS0063 1 HSM004
    
```

Figure 38. DEV Report

The graphic form of this report shows each user's device delay percentage and device using percentage.

## Field descriptions

Table 19. Fields in the DEV Report

Field Heading	Meaning
Jobname	Name of a job that is delayed by device volumes. The Device Delay report does not summarize data by job groups; all jobs within a job group are reported individually.
C	A one-character abbreviation for the job class as follows: <b>S</b> Started task <b>T</b> TSO <b>B</b> Batch <b>A</b> ASCH <b>O</b> OMVS
Service Class	The name of the service class that a specified job has been running in.
DLY %	Delay the waiting job (address space) is experiencing because of contention for devices during the report interval, expressed as a percentage. <b>Note:</b> This DLY% value is also found in the DEV field on the job delay report.
USG%	The percentage of time when the job is transferring data between DASD or tape and central storage. (Not just the volumes listed under the VOLSER columns on the report.) <b>Note:</b> To find all the using volumes for a jobname you must scan an entire resource-oriented device delay (DEV) report, using the FIND command.

Table 19. Fields in the DEV Report (continued)

Field Heading	Meaning
CON %	<p>The percentage of time during the report interval when devices used by the address space were connected to channel path(s) that actually transferred data between the devices and central storage. This value measures connect time of the DEV volumes as well as I/O requests to any device on a block multiplex channel for which the measurement facility is active. RMF obtains connect time at each sample.</p> <p><b>Note:</b></p> <ol style="list-style-type: none"> <li>1. When comparing the CON % and the USG% fields in this report, you must be aware that CON % is a measured multi-state value, while USG% is a sampled single state value. Thus, CON % includes time while the job was using more than one device at the same time, while USG % does not. The value in the CON % field might include more devices than the USG% field. The USG % field may include a considerable amount of delay.</li> <li>2. Some of the connect time from the previous range period might be included in the CON % value, while some of the connect time in the current report interval might be absent. This discrepancy is noticeable on devices that have very long channel programs, such as paging devices.</li> </ol>
Main Delay Volume(s)	<p>Up to four DEV volumes contributing most to the delay of the job. The DEV volume having the largest delay percentage appears first.</p> <p><b>VOLSER</b> The serial number of a DASD or tape contributing to the job delay.</p> <p>% The percentage of delay caused because the job was waiting to use the named volume.</p>

### Report options

```

RMF Delay Report Options: DEV                               Line 1 of 4
Command ==>>>                                           Scroll ==>> HALF

Change or verify parameters. To exit press END.
Changes will apply to DELAY, DEV, ENQ, HSM, JES, PROC, PROCU, STOR, STORC,
STORF, STORM, and XCF.

Class           ==>> ALL           Classes: ALL TSO BATCH STC ASCH OMVS
Service class  ==>> *ALL          *ALL or one of available service classes below

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

Available Service classes
ZOSUN         OMVSKERN  PRDTSO      STCCMD      SYSTEM      SYSSTC
    
```

Figure 39. DEV Report Options Panel

The DEV Report Options panel is similar to the Delay Report Options panel, but does not contain Summary or Criterion. See “Report options” on page 68 (the Delay Report Options panel) for a description of the fields.

The parameters that you specify on this panel affect all job-oriented detail delay reports.

## DEVN - Device Activity Report

The Device Activity (DEVN) report gives information about all or a subset of online devices. The report is based on the Device Resource Utilization (DEVR) report, but only shows the average number of jobs using or being delayed for the devices and not every job, as shown on the DEVR report.

The report provides the capability to select a subset of all available devices and to sort the displayed devices. You can rearrange the displayed list of devices by any

activity category you want and focus on devices with common characteristics (for example, same volser number or device number prefix, or devices having the same type or are connected to the same control unit type). This is done using cursor-sensitive control (see “Cursor-sensitive control” on page 74).

## How to request this report

To request this report, select **U** from the Primary menu, then **DEVN** from the User Selection menu, or you can enter the following command:

```
DEVN
```

## Contents of the report

```

RMF V2R2 Device Activity                               Line 1 of 118
Command ==>                                           Scroll ==> HALF

Samples: 100      System: AQTS  Date: 09/28/16  Time: 14.23.20  Range: 100  Sec
Devices reported:  ALL
Report is sorted by: Jobs, DEL

-- Device Identification -- -- Activity -- ACT CON DSC - Pending - - Jobs -
VolSer Num  Type  CU    S  Rate RspT IosQ %  %  %  %  Rsn. %  USG  DEL
HSM013 006C 33903 3990-3 S  9.1 .092 .018 68  4  62  2 DB  1  0.0 0.8
CLR010 0051 33903 3990-3 S 80.7 .011 .005 47 24  1 22 DB 11  0.2 0.7
HSML17 0703 33903 3990-3 S 52.2 .015 .000 76 22 54  0  0.2 0.6
HSM015 006E 33903 3990-3 S 11.1 .024 .001 26  3 20  3  0.0 0.3
TS0060 0056 33903 3990-3 S  8.9 .034 .001 30  9 18  3 DB  2  0.1 0.2
D22SHR 0B70 3380 3880-3 S 13.5 .014 .000 18  2 16  0  0.0 0.1
HSML10 0043 33902 3990-3 S 68.4 .008 .000 53 43  3  7  0.4 0.1
CLR014 0149 33903 3990-3 S 37.2 .005 .000 19 11  7  1  0.1 0.1
TS0024 0842 33902 3990-3 S  5.8 .015 .000  9  1  8  0  0.0 0.1
HLPVOL 02E2 3380D 3880-3 S  2.8 .024 .000  7  1  5  1 DB  1  0.0 0.1
HSM011 006A 33903 3990-3 S  2.3 .033 .000  7  1  6  0  0.0 0.1
HSML04 005B 33903 3990-3 S 13.9 .006 .001  7  3  2  2  0.0 0.1
MIG015 01E3 3380K 3990-3 S  5.5 .017 .000 10  5  5  0  0.0 0.1
SPOL16 0844 33902 3990-3 S  9.4 .007 .001  6  2  1  3 DB  3  0.0 0.1
SYSLBX 01AE 33902 3990-3 S 53.7 .002 .000  9  7  0  2 DB  1  0.1 0.0
    
```

Figure 40. DEVN Report

The DEVN report has two parts.

- The top part provides information about the selection criteria and the sort criteria for the displayed devices.
- The bottom part is based on information from the DEVR report. It is similarly organized as the Postprocessor DASD Activity report (see “DEVICE - Device Activity report” on page 367).

The first four columns show the device identification (volser, device number, device type and control unit type). These columns can be used with cursor-sensitive control to change the scope of displayed devices.

The columns on the right side of the report display the device utilization information. These columns can be used to sort the report.

The graphic form of the report shows for each device the response time in milliseconds broken down in IOS queue time and service time.

### Field descriptions

All fields in the DEVN report are the same as in the DEVR report (see Table 22 on page 76) except for the following:

Table 20. Fields in the DEVN Report

Field Heading	Meaning
Devices reported:	The criteria selected for the devices being reported.  The devices being reported can be selected using cursor-sensitive control from the fields listed under Device Identification.
Report is sorted by:	The sort criteria for the devices being reported.  The sort criteria can be selected using cursor-sensitive control from any of the fields listed under the columns between Activity and Jobs.
Activity IosQ	The average number of seconds an I/O request must wait on an IOS queue before a SSCH instruction can be issued. A delay occurs when a previous request to the same sub-channel is in progress.  The value is calculated as: $\text{IosQ} = \frac{\text{IOS Queue Count} / \# \text{ Samples}}{\text{Device Activity Rate}}$  This field is not shown on the DEVR report but is available in the ISPF table of the DEVR report.
Pending Reasons	Only the highest delay reason and percentage is listed.
Jobs	<b>USG</b> The average number of jobs using the device during the report interval. <b>DEL</b> The average number of jobs being delayed for the device during the report interval.

### Cursor-sensitive control

Cursor-sensitive control on the DEVN report is extended (compared to other Monitor III reports) by new capabilities. In addition to navigation control as in other reports, you can

- Recreate the report with a different scope of selected devices
- Get a different sort order of the displayed devices

Therefore, cursor-sensitive control does not maintain the return path. Pressing PF3 on a subsequent report will always return you to the Primary menu.

In addition, the latest selection criteria and sort order is saved throughout the session and will be used on the next invocation of the report. The initial display is always shown according to the jobs being delayed for the device in descending order, and the initial selection criteria display all online devices.

The following table provides an overview about selection and sort using cursor-sensitive control.

Table 21. DEVN Report - Cursor-sensitive Control for Select and Sort

Cursor-sensitive Column	Cursor Position within Column	Result
VolSer	1, 2	Device Activity Trend report for the selected volume

Table 21. DEVN Report - Cursor-sensitive Control for Select and Sort (continued)

<i>Cursor-sensitive Column</i>	<i>Cursor Position within Column</i>	<i>Result</i>
VolSer	3 - 6	Device Activity report for devices starting with the same volser prefix. Example: Cursor position is 3 below volser TSO060: The result is a DEVN report for all TSOxxx devices.
Num	1-3	Device Activity report for devices starting with the same Num prefix. Example: Cursor position is 3 below num 006E: The result is a DEVN report for all devices with an address of 006x.
Type	any	Device Activity report for all online devices with the same device type.
CU	any	Device Activity report for all online devices with the same CU type.
S	--	No cursor-sensitive control.
Rsn %	--	No cursor-sensitive control.
All other	any	Device Activity report sorted in descending order by the selected column.

**Note:** Selecting the same column (VolSer, Num, Type, or CU) a second time displays the Device Activity report for ALL online devices again.

---

## DEVR - Device Resource Delays Report

The Device Resource Delays report (DEVR) shows the devices (volumes) and the jobs using or being delayed by them (as indicated on the Device Delays report).

On the DEVR report, the type of delay is listed under Pend Reasons as:

**DB** Device busy delay  
**CMR** Initial command response time

### How to request this report

To request the DEVR report, select **3** from the Primary menu, and then select **3** on the Resource Report Selection Menu (shown in Figure 8 on page 27) or enter the following command:

```
DEVR [volser]
```

## Contents of the report

```

RMF V2R2 Device Resource Delays Line 1 of 374
Command ==> Scroll ==> HALF
Samples: 100 System: MVS1 Date: 09/28/16 Time: 10.03.20 Range: 100 Sec

Volume S/ Act Resp ACT CON DSC PND %, DEV/CU Service USG DLY
/Num PAV Rate Time % % % Reasons Type Jobname C Class % %
160444 S 17.5 1.2 2 2 0 PND 0 33909 GPMSRVPM S GPMSERVE 0 1
8803 1.2H 2105 RMF S SYSSTC 2 0
SYSUSR S 0.3 0.4 0 0 0 PND 0 33903 *MASTER* S SYSTEM 1 0
D440 2* 2105 CATALOG S SYSTEM 1 0
MVSTGT S 0.8 0.9 0 0 0 PND 0 33909 GPMSRVPM S GPMSERVE 1 0
7790 2105
SYSXCP S 1.4 0.7 1 1 0 PND 0 33903 XCFAS S SYSTEM 1 0
D32A 2105
    
```

Figure 41. DEVR Report

The first block of columns in the report contain information related to each volume.

The remaining columns contain information related to each job. RMF sorts the volumes in descending order according to the average number of delayed users (1 user delayed 100% is equivalent to 100 users each delayed 1% of the time), and the waiting jobs by descending delay percentages. If RMF is unable to obtain valid hardware data, it prints dashes (---) for the hardware measurements, while percentages normally appear.

The fields DLY DB % and DLY CU % contain data about I/O request delays caused by contention at the control unit and device level.

The graphic DEVR report shows the average number of active users for each device that were delayed, connected, disconnected or pending.

### Field descriptions

Table 22. Fields in the DEVR Report

Field Heading	Meaning
Volume/Num	The name of an online volume and the device number where the volume is mounted.
S/PAV	<p><b>S</b> An S in the first line of this column indicates that the device was generated during system generation as a shared device.</p> <p><b>PAV</b> PAV count — A value in the second line of this column gives the number of parallel access volumes (base and alias) which were available at the end of the reporting.</p> <p>If the number has been changed during the report interval, it is followed by an '*'. If the device is a HyperPAV base device, the number is followed by an 'H', for example, 1.2H. The value is the average number of HyperPAV volumes (base and alias) for that range.</p> $\text{Average \# of HPAV devices} = \frac{\text{Accumulated \# of HPAV devices}}{\text{Number of Samples}}$

Table 22. Fields in the DEVR Report (continued)

Field Heading	Meaning
Act Rate	<p>The rate per second that I/O instructions (SSCH, RSCH, and HSCH) to a device completed successfully.</p> <p>The calculation is:</p> $\text{Act Rate} = \frac{\# \text{ I/O Instructions}}{\text{Range Time}}$
Resp Time	<p>The average response time (in milliseconds) that the device required to complete an I/O request.</p> <p>The calculation is:</p> $\text{Resp Time} = \frac{\text{Active Time}}{\# \text{ I/O Instructions}} + \text{IOS Queue Time}$
ACT %	<p>The percentage of time during the report interval when the device was active. To derive this value, RMF computes the accumulated percent active time as follows:</p> $\text{ACT \%} = \text{PEND \%} + \text{CON \%} + \text{DSC \%}$ <p><b>PEND %</b> Percentage of time all I/O requests wait in the logical control unit queue (CU-HDR) before there is an available path. Pending time includes the time spent waiting for a channel, control unit, or head of string, or for the actual device (if it is a shared device that is reserved by another processor).</p> <p><b>CON %</b> Percentage of time the device was connected to a channel path to actually transfer data between the device and storage.</p> <p><b>DSC %</b> Percentage of time the device has an active channel program and is disconnected (not transferring data). Disconnect time includes seek time, normal rotation delay time, and extra rotation delay because the channel was busy when the device needed to reconnect.</p>
CON %	<p>The percent connect time. See the description under % ACT. RMF calculates the value as follows:</p> $\text{CON \%} = \frac{\text{Accumulated Connect Time}}{\text{Range Time}} * 100$
DSC %	<p>The percent disconnect time. See the description under %ACT. RMF calculates the value as follows:</p> $\text{DSC \%} = \frac{\text{Accumulated Disconnect Time}}{\text{Range Time}} * 100$
<p><b>Note:</b></p> <ol style="list-style-type: none"> <li>When comparing the ACT %, CON %, DSC %, or PND % fields with the USG % field in this report, you must be aware that ACT %, PND %, CON %, and DSC % are measured multi-state values, while USG % is a sampled single state value. If a single I/O request is very long (such as a long-running channel program), the PND %, CON %, and DSC % values might be too low because of timer overflow errors.</li> <li>The channel updates the data fields used to calculate CON %, DSC %, and PND % when the I/O operation completes. Therefore, some of the time from the previous report interval might be included in these values, while some of the time in the current report interval might be absent from these values. This discrepancy is noticeable on devices that have very long channel programs, such as paging devices.</li> </ol>	

## Mon III - DEVR

Table 22. Fields in the DEVR Report (continued)

Field Heading	Meaning
PND % Reasons	<p>The first entry is always the pending percentage (PND). See the description under % ACT. RMF calculates the value as follows:</p> $\text{PND \%} = \frac{\text{Accumulated Pending Time}}{\text{Range Time}} * 100$ <p>DLY DB % and DLY CU % are included in pending time.</p> <p>Below PND % are the pend reasons that contribute to the total pending percentage. A value appears only when there is a non-zero delay percentage. Pend Reasons can be one of the following:</p> <p><b>DB</b> Device busy delay, which is the percentage of time during the report interval when the channel subsystem measured I/O request delay because the device was busy. Device busy might mean that the volume is in use by another system, the device is reserved by another system, a head of string busy condition caused the contention, or some combination of these conditions has occurred.</p> $\text{DLY DB\%} = \frac{\text{Accumulated DB Delay Time}}{\text{Range Time}} * 100$ <p><b>CMR</b> Command response time delay, which is the percentage of time during the report interval when the first command of an I/O instruction of the channel program is sent to the device, until the device indicates it has accepted the command.</p> $\text{DLY CMR\%} = \frac{\text{Accumulated CMR Delay Time}}{\text{Range Time}} * 100$ <p><b>Note:</b> If either hardware data or volume-related percentages are not available, this field is blank.</p>
DEV/CU Type	The top number represents the device type. The bottom number represents the control unit model.
Jobname	Name of a job using or being delayed by the DEV volume. The DEVR delay report does not summarize data by job groups; all jobs within a job group are reported individually. RMF lists all jobs for each device, by descending delay percentages.
C	A one-character abbreviation for the job class as follows: <b>S</b> Started task <b>T</b> TSO <b>B</b> Batch <b>A</b> ASCH <b>O</b> OMVS
Service Class	The name of the service class that a specified job has been running in.
USG %	The percentage of time when the job has had a request accepted by the channel for the specified Volume, but the request is not yet complete.
DLY %	Delay the waiting job (address space) is experiencing because of contention for a specific volume during the report interval, expressed as a percentage.

### Monitor III Utility fields

Table 23 shows additional fields for the Device Resource Delay report.

You can use the Monitor III Utility to customize the DEVR report.

Table 23. Additional Fields in the DEVR Report

Field Heading	Meaning
Percentage of pending time	The percentage of time during the report interval when the device was pending.
IOS queue time	The average number of milliseconds an I/O request must wait on an IOS queue before a SSCH instruction can be issued. Delay occurs when a previous request to the same subchannel is in progress.
Percentage of device busy delay	The percentage of time during the report interval when the channel subsystem measured I/O request delay because the device was busy. Device busy might mean that the volume is in use by another



Table 23. Additional Fields in the DEVR Report (continued)

Field Heading	Meaning
Percentage of control unit busy delay time	The percentage of time during the report interval when there is I/O request delay because the control unit was busy.
Percentage of switch port busy delay time	The percentage of time during the report interval when there is I/O request delay because the switch port was busy.

### Report options

You can use the DEVR Report Options panel to select the volume to be included in the DEVR report, or all volumes, from a list of available volumes.

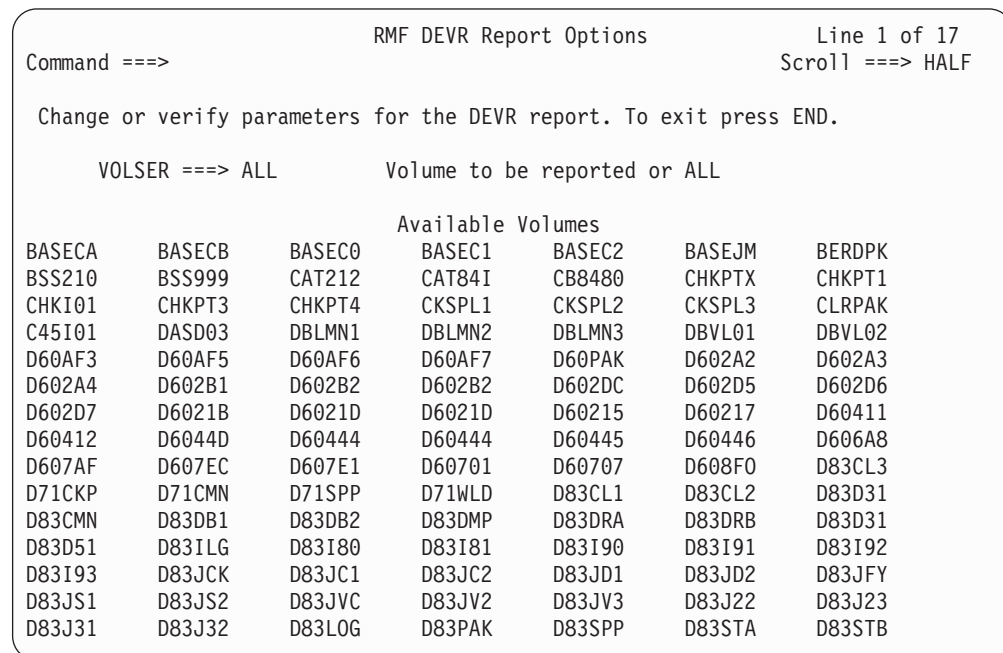


Figure 42. DEVR Report Options Panel

### VOLSER

The volume serial number of the device that you want information about.

Enter

- ALL for information about all devices that have jobs using it or being delayed by it in the system.
- A name with an asterisk (\*) as a "wild card" character. For example: to request a report for all volumes starting with D8, specify 'D8\*' for VOLSER.

**Note:** You cannot use the wild card when calling the report, that is, when you use the command `DEVR volser`. Here, *volser* must be a complete volume serial number, an asterisk will be interpreted as part of the volser.

- One of the volumes listed under Available Volumes.
- The volume serial number of a device that will be in the system at a later time.

If the volume that you specify is not currently available, it will appear on the report when it is available.

Your selection is saved across sessions in the current option set.

**Available Volumes**

The list of the online volumes in the system.

If the volume you want is not listed, it was not online during the current report interval. If you specify the volume, it will appear on the report when it is online.

**DEVT - Device Activity Trend Report**

The Device Activity Trend (DEVT) report shows the device activity for a selected volume for the last 20 reporting ranges. The report is based on the Device Activity (DEVN) report and can be used for a selected device as follows:

- To identify times of peak device utilization
- To analyze the device utilization over time
- To analyze device delay situations
- As a device summary report

**How to request this report**

To request this report, select **U** from the Primary menu, and then **DEVT** together with a volser from the User Selection menu.

**Note:** The report can also be requested from the Device Activity (DEVN) report using cursor-sensitive control. If the report is selected from the User Selection menu and no volser is specified or an invalid volser is specified, the DEVN report is displayed.

**Contents of the report**

RMF V2R2 CLR010 Activity Trend											Line 1 of 20			
Command ==>											Scroll ==> HALF			
Samples: 100		System: AQTS		Date: 09/28/16		Time: 11.58.20		Range: 100		Sec				
VolSer: CLR010		Number: 0051		Type and CU-Type: 33903		3990-3								
Latest: 09/28/16		at 11.58.20		Range/Line: 100		Sec								
Earliest: 09/28/16		at 11.26.40		Total Range: 2000		Sec		00.33.20						
Time	-----	Activity	-----	ACT	CON	DSC	- Pending	-	---	Jobs	---	WFL		
	S	Rate	RspT	IosQ	%	%	%	%	Rsn. %	USG	DEL	TOT	%	
11.58.20	S	154	.011	.006	79	41	1	37	DB	23	0.4	1.3	1.7	24
11.56.40	S	138	.018	.012	88	34	3	51	DB	36	0.4	2.2	2.6	15
11.55.00	S	159	.016	.010	87	40	5	42	DB	26	0.4	2.1	2.5	17
11.53.20	S	146	.011	.006	75	37	3	35	DB	20	0.4	1.3	1.7	22
11.51.40	S	125	.014	.008	69	33	2	34	DB	22	0.3	1.4	1.7	18
11.50.00	S	124	.016	.009	80	32	2	46	DB	34	0.4	1.7	2.1	17
11.48.20	S	127	.015	.008	79	34	3	42	DB	30	0.3	1.6	1.9	18
11.46.40	S	127	.021	.014	87	31	3	53	DB	41	0.3	2.5	2.8	12
11.45.00	S	135	.021	.015	89	33	3	53	DB	41	0.3	2.6	2.9	11
11.43.20	S	107	.026	.018	86	29	4	53	DB	40	0.3	2.5	2.8	11
11.41.40	S	97.9	.031	.022	88	25	2	61	DB	46	0.2	2.8	3.0	8

Figure 43. DEVT Report

The DEVT report has two parts.

- The top part provides information about the selected device, its volser, device number, device and control unit type, and information about the reported range.

- The bottom part is based on the DEVN report.  
Each row is preceded by a time stamp to identify the start time of the reporting range. The device activity columns are exactly the same as those shown on the DEVN report.  
At the right, a column showing the device workflow percentage is added. The workflow column is calculated from the average number of users using or being delayed for the device from the DEVN ISPF table. Please keep in mind that the value is not as precise as workflow values shown on the Workflow/Exception report.

## Field descriptions

Table 24. Fields in the DEVT Report

Field Heading	Meaning
VolSer:	The name of an online volume.
Number:	The device number where the volume is mounted.
Type and CU-Type:	The device type and the control unit model.
Latest:	Begin date and time of the last reported range on the report.
Range/Line:	Reported range per displayed line on the report.
Earliest:	Begin date and time of the first reported range on the report.
Total Range:	Total reported range on the report, expressed in seconds and HH.MM.SS.
Time	The start time of the reported range.

You find the description of all other fields in the report either in Table 22 on page 76 or in Table 20 on page 74.

### Cursor-sensitive control

Placing the cursor on the time stamp for a selected row will recreate the report starting at the selected time period. The return path is not maintained, which means that pressing PF3 will return you to the Primary Menu.

Using cursor-sensitive control from any other column will invoke the Device Resource Utilization (DEVR) report for the selected reporting range and the selected device. In this case, pressing PF3 on the DEVR report will return you to the DEVT report.

---

## DSD - Detailed Storage Delays Report

Figure 44 on page 82 shows a modified version of the Storage Delays report that replaces **Working Set Central** and **Expanded** with three columns: **VIO**, **XMEM** and **HIPR**. On the Storage Delays report, this information is combined and shown in the **OTHR** column.

### How to request this report

To request the DSD report, select **U** on the Primary menu, and then select **2** on the User menu, or enter the following command:

```
DSD
```

### Contents of the report

DLY % , or delay percentage, is the percentage of time during the report interval that the job is experiencing a delay because of contention for storage. If DLY % is greater than 10%, it could indicate a problem.

% Delayed for breaks down the number under DLY % into the various types of storage delays affecting each job.

The COMM and LOCL fields include shared storage paging.

```

RMF V2R2 Storage Delays
Command ==>
Line 1 of 206
Scroll ==> HALF

Samples: 100 System: MVS1 Date: 09/28/16 Time: 10:31:40 Range: 100 Sec

Jobname C Service DLY % Delayed for
Class % COMM LOCL SWAP OUTR VIO XMEM HIPR
*MASTER* S STC_HIGH 0 0 0 0 0 0 0 0 0
PCAUTH S STC_HIGH 0 0 0 0 0 0 0 0 0
RASP S STC_LOW 0 0 0 0 0 0 0 0 0
TRACE S STC_LOW 0 0 0 0 0 0 0 0 0
XCFAS S STC_HIGH 0 0 0 0 0 0 0 0 0
GRS S STC_HIGH 0 0 0 0 0 0 0 0 0
SMXC S STC_HIGH 0 0 0 0 0 0 0 0 0
SYSBMAS S STC_HIGH 0 0 0 0 0 0 0 0 0
DUMPSRV S STC_HIGH 0 0 0 0 0 0 0 0 0
CONSOLE S STC_HIGH 0 0 0 0 0 0 0 0 0
ALLOCAS S STC_HIGH 0 0 0 0 0 0 0 0 0
TLCS S STC_LOW 0 0 0 0 0 0 0 0 0
GPDB S STC_LOW 0 0 0 0 0 0 0 0 0
NETVIEW1 S STC_HIGH 0 0 0 0 0 0 0 0 0
TSO S STC_HIGH 0 0 0 0 0 0 0 0 0
APFTABLE S STC_HIGH 0 0 0 0 0 0 0 0 0
    
```

Figure 44. Modified STOR Report Showing all Storage Delays in Detail

### Field descriptions

The fields in the DSD report are identical to the fields in the STOR report (see Table 64 on page 155) except for the XMEM and HIPR fields.

Table 25. Fields in the DSD Report

Field Heading	Meaning
Delayed for XMEM	This column contains the paging delays from cross memory address spaces.
Delayed for HIPR	This column contains the paging delays from standard hiperspaces (including waits during scroll wait), but not ESO hiperspaces.

There are no report options to be specified for the DSD report.

## DSND - Data Set Delays Report

The DSND report presents information about the utilization (using and delay) of one data set or a group of data sets. For each selected data set, information is given about

- The volume the data set resides on
- All jobs that are using this data set or that are waiting for this data set.

RMF sorts the data sets by descending overall delay percentages.

You can use this report as base for further analysis:

- To investigate the performance of a volume and list all jobs that are delayed because of it, use cursor-sensitive control on any indicator under **Volume** to display the related DSNV report.
- To see performance information for a specific job, use cursor-sensitive control on any job listed under **Jobname** to display the related DSNJ report.
- To view all data sets which RMF found active in the report interval, or to change the list of data sets to be reported on, enter the command ROPTIONS to display the DSND Report Options panel.

## How to request this report

To request the Data Set Delays report, select **3** from the Primary Menu, and then select **3A** on the Resource Report Selection Menu (shown in Figure 8 on page 27). or enter the following command:

```
DSND [dsname]
```

In addition, you can navigate to this report through cursor-sensitive control from the DSNJ report or DSNV report.

## Special considerations

The Device Resource Delays report (DEVR) provides USG and DLY values for jobs that are using devices or are waiting for them. This data is gathered in a multistate fashion, this means that there may be several wait records for the same job for the same device. The reporter changes to "pseudo multistate", this can result in one USG counter and one DLY counter in parallel within a cycle, but does not take multiple wait records into account.

Data gathering for the Data Set Delays reports (DSND, DSNJ, and DSNV) is different. Here, several wait records referring to the same device are not treated as being the same and counted only once because they may refer to different data set names, and have to be counted individually.

As a result, the sum of the USG and DLY percentage values in these reports can be different to the USG and DLY percentage values in the DEVR report. Therefore, the three reports contain the headings DUSG% and DDLY% instead of USG% and DLY% to indicate a potential difference to the related values in the DEVR report.

## Contents of the report

```

RMF V2R2 Data Set Delays
Command ==>>>
Line 1 of 6
Scroll ==>> HALF

Samples: 100 System: MVS1 Date: 09/28/16 Time: 10.03.20 Range: 100 Sec

Input Data Set Name: BDA.CTT*

----- Data Set Name ----- Volume Jobname ASID DUSG% DDLY%
BDA.CTT.MSPCT.SP41XCTT.CTTGUIDE.BOOK EDSS99 BOECL2 003C 20 15
                                           BGGEETEO 0201 15 10
BDA.CTT.MSPCT.SP41XCTT.GUIDE EDSS99 BOECL2 003C 13 10
                                           BGGEETEO 0201 10 8
BDA.CTTX.TEST DATA94 BSHR 0022 3 16
BDA.CTT0.INFORM.SEQ DATA67 BSHR 0074 3 8
    
```

Figure 45. DSND Report

There is no graphic version of this report available.

### Field descriptions

Table 26. Fields in the DSND Report

Field Heading	Meaning
Input Data Set Name	Name of the data set or group of data sets to be reported on.  This is the name which has been specified as command parameter or has been selected via report options.
Data Set Name	Name of a data set which was utilized during the report interval and selected for reporting.  RMF lists the data set names by descending overall delay percentages.
Volume	Name of the volume on which the data set resides.
Jobname	Name of a job using or being delayed by the data set.  RMF lists all jobs for each data set by descending delay percentages.
ASID	Hexadecimal address space identifier (ASID) of the job using the data set or waiting for its availability.
DUSG%	Percentage of time when the job has had an I/O request accepted by the channel for the volume on which the data set resides, but the request is not yet complete. <b>Note:</b> See "Special considerations" on page 83.
DDLY%	Percentage of time when the job was waiting to use the data set because of contention for the volume where the data set resides.

### Report options

The DSND Report Options panel displays a list of all data set names which have been found active during the report interval. You can select a data set name by placing an **S** in front of the displayed data set name. The selected name is re-displayed in the header field "Selected Data Set Name". This field is an input field which can be used also to enter a data set name or a group of data set names using a "wild card", for example: BDA.CTT\*, directly. The wild card example **BDA.CTT\*** lets RMF select all data set names which start with the character sequence **BDA.CTT**. If someone specifies only "\*", RMF reports on all data sets which are being utilized in the report interval.

The current selection is displayed on top of the data set names list.

```

RMF DSND Report Options                               Line 1 of 12
Command ==>>>                                       Scroll ==>> HALF

Select (S) or fill-in a data set name or a group of data set names
for the DSND report. To exit press END.

Selected Data Set Name: BDA.CTT*

Sel  Data Set Name
-    BDA.CTT.MSPCT.SP41XCTT.CTTGUIDE.BOOK
-    BDA.CTT.MSPCT.SP41XCTT.GUIDE
-    BDA.CTT0.INFORM.SEQ
-    BDA.CTTX.TEST
-    BHEW.ERBMFDTS.LST90514
-    BSHR.FIX.LINKLIB
-    RMF.R430.NLS.OLDENG.ERBCOPS3.SEQ
-    RMF.R430.NLS.OLDENG.ERBFMTS3.SEQ
-    RMF.R430X10.LPALIB
-    RMF.R430X10.SRMFCLS
-    RMF.R430X10.SRMFJPN
-    RMF.R530.FPFS

```

Figure 46. DSND Report Options Panel

#### Selected Data Set Name

The currently selected name of a data set or group of data sets to be reported on.

This field is an input field and can be overwritten according to the rules for z/OS data set names. It is possible also to use an '\*' as "wild card" as last character of the data set name. By using a wild card, all data sets starting with the character sequence before the '\*' are reported on no matter which characters follow.

**Sel** An S can be placed in front of the data set name to be selected. This results in replacing the data set name in the header field "Selected Data Set Name".

#### Data Set Name

The name of a data set which was found active during the report interval. The data set names are sorted in alphabetical order.

#### Note:

1. The *RESET* command is not supported.
2. Only one data set name can be selected.
3. If a data set name is selected and the data set name in the input field is changed at the same time, the selected data set name is used.
4. If the data set name is blanked out, it is possible to leave the panel, but the fields in the report will be empty.

---

## DSNJ - Data Set Delays - Job Report

The DSNJ report presents information about data set utilization for a specific job:

- The EXCP rate and the percentage of time when data transfer for this job took place.
- A list of all data sets being utilized by the job.

You can use this report as base for further analysis:

- To investigate the performance of a volume and list all jobs that are delayed because of it, use cursor-sensitive control on any indicator under **Volume** to display the related DSNV report.
- To see performance information for a specific data set, use cursor-sensitive control on any data set listed under **Data Set Name** to display the related DSND report.

## How to request this report

To request the Data Set Delays - Job report, select **2** from the Primary Menu, and then select **1A** on the Job Report menu (shown in Figure 7 on page 26). or enter the following command:

```
DSNJ [jobname]
```

In addition, you can navigate to this report through cursor-sensitive control from the DSND report or DSNV report.

## Contents of the report

RMF V2R2 Data Set Delays - Job						Line 1 of 5
Command ==>						Scroll ==> HALF
Samples: 100	System: MVS1	Date: 09/28/16	Time: 10.03.20	Range: 100	Sec	
Jobname: BOECL2	EXCP Rate: 123.5	Connect: 41%				
ASID	----- Data Set Name -----	Volume	Num	DUSG%	DDL%	
003C	BDA.CTT.MSPCT.SP41XCTT.CTTGUIDE.BOOK	EDSS99	0312	20	15	
	Bshr.FIX.LINKLIB	DATA68	0257	2	10	
	BDA.CTT.MSPCT.SP41XCTT.GUIDE	EDSS99	0312	13	10	
	RMF.R430.NLS.OLDENG.ERBCOPS3.SEQ	EDSS09	0312	3	8	
	-- N/A --	-----	----	1	7	
	BHEW.ERBMFDTs.LST90514	DATA38	0122	1	6	

Figure 47. DSNJ Report

There is no graphic version of this report available.

## Field descriptions

Table 27. Fields in the DSNJ Report

Field Heading	Meaning
Jobname	Name of the job for which reporting was requested.
EXCP Rate	Number of EXCP requests per second for the job being reported on.
Connect	Percentage of time during the report interval when devices used by the job were connected to channel path(s) that actually transferred data between the devices and central storage.
ASID	Address space identifier (ASID) of the job being reported on.
Data Set Name	Name of the data set being utilized by the current job.  RMF lists all data sets by descending delay percentages. <b>Note:</b> The using and delay information for all I/Os for which the data set name information is not available is accumulated in a single slot. In this case, -- N/A -- is provided instead of a data set name. If these I/Os are directed to different volumes, dashes are shown in columns for <b>Volume</b> and <b>Num</b> . This happens if only those I/O instructions have been detected for which no data set information is provided by the SMS subsystem, like <ul style="list-style-type: none"> <li>• I/Os to system data sets (like paging or spooling)</li> <li>• I/Os to any data set which was opened prior to SMS subsystem initialization</li> <li>• I/Os like SENSE or RELEASE</li> <li>• System I/Os not done by an access method</li> </ul>



Table 27. Fields in the DSNJ Report (continued)

Field Heading	Meaning
Volume	Name of the volume on which the data set resides which was utilized during the current report interval.
Num	Device number where the volume is mounted.
DUSG%	Percentage of time when the job has had an I/O request accepted by the channel for the volume on which the data set resides, but the request is not yet complete. <b>Note:</b> See "Special considerations" on page 83.
DDLY%	Percentage of time when the job was waiting to use the data set because of contention for the volume where the data set resides.

## DSNV - Data Set Delays - Volume Report

The DSNV report presents information about the utilization of data sets that reside on a specific DASD volume.

The first part of the report provides a general overview on important activity and delay data for the volume. The second part of the report displays a list of all data sets on this volume that were found active during the reporting interval.

RMF sorts the data sets by descending overall delay percentages.

You can use this report as base for further analysis:

- To investigate the performance of a specific job that is using data sets on this volume or is waiting for them, use cursor-sensitive control on any indicator under **Jobname** to display the related DSNJ report.

### How to request this report

To request the Data Set Delays - Volume report, select **3** from the Primary Menu, and then select **3B** on the Resource Report Selection Menu (shown in Figure 8 on page 27). or enter the following command:

```
DSNV [volser]
```

In addition, you can navigate to this report through cursor-sensitive control from the DEVR report, the DSND report, or the DSNJ report.

## Contents of the report

```

RMF V2R2 Data Set Delays - Volume                               Line 1 of 10
Command ==>>>                                               Scroll ==>> HALF
Samples: 100      System: MVS1 Date: 09/28/16 Time: 10.03.20 Range: 100 Sec
----- Volume EDSS99 Device Data -----
Number:  0B4A      Active:    84%      Pending:  22%      Average Users
Device:  3380A    Connect:   10%      Delay DB:  22%      Delayed
Shared:  Yes      Disconnect: 52%      Delay CM:  0%      0.4
PAV:     1.6H

----- Data Set Name ----- Jobname  ASID  DUSG%  DDLY%
BDA.CTT.MSPCT.SP41XCTT.CTTGUIDE.BOOK  BOECL2  003C   20    15
                                           BGGEETE0 0201   15    10
BDA.CTT.MSPCT.SP41XCTT.GUIDE          BOECL2  003C   13    10
                                           BGGEETE0 0201   10     8
RMF.R430X10.LPALIB                     BWS0    0058    1    12
-- N/A --                               *MASTER* 0001    0    10
RMF.R430.NLS.OLDENG.ERBCOPS3.SEQ       BOECL2  003C    3     8
    
```

Figure 48. DSNV Report

There is no graphic version of this report available.

### Field descriptions

Table 28. Fields in the DSNV Report

Field Heading	Meaning
Device Data Section	This sections contains identical information as provided in the Device Delays variation of the Job Delay report (see "Device Delay variation" on page 118).
Data Set Name	Name of a data set which was utilized during the report interval and resides on the selected volume.  RMF lists the data set names by descending overall delay percentages.
Jobname	Name of a job using or being delayed by the data set.  RMF lists all jobs for each data set by descending delay percentages. <b>Note:</b> The using and delay information for all I/Os for which the data set name information is not available is accumulated, and -- N/A -- is provided instead of a data set name. This happens if only those I/O instructions have been detected for which no data set information is provided by the SMS subsystem, like <ul style="list-style-type: none"> <li>• I/Os to system data sets (like paging or spooling)</li> <li>• I/Os to any data set which was opened prior to SMS subsystem initialization</li> <li>• I/Os like SENSE or RELEASE</li> <li>• System I/Os not done by an access method</li> </ul>
ASID	Address space identifier (ASID) of the job using the data set or waiting for it.
DUSG%	Percentage of time when the job has had an I/O request accepted by the channel for the volume on which the data set resides, but the request is not yet complete. <b>Note:</b> See "Special considerations" on page 83.
DDLY%	Percentage of time when the job was waiting to use the data set because of contention for the volume where the data set resides.

## Report options

```

RMF DSNV Report Options                               Line 1 of 2
Command ==>                                         Scroll ==> HALF

Change or verify parameters for the DSNV report. To exit press END.

VOLSER ==> MVSLIB      DASD Volume to be reported

Available DASD Volumes
MVSDOC  MVSJOB  MVSLIB  MVSSMP  MVSTGT  RMFUSR  RMFUS02
SYSCAT  SYSPAG  510948

```

Figure 49. DSNV Report Options Panel

The Report Options panel displays a list of all DASD volumes which have been found active during the report interval.

The current selection is displayed on top of the volume list.

### VOLSER

The volume serial number of the device for which data set level reporting is being requested.

One of the volumes listed under Available DASD Volumes, or, the volume serial number of a device that will be in the system at a later time.

The selection is saved across sessions in the current option set.

### Available DASD Volumes

The list of the online DASD volumes in the system. The volumes are sorted in alphabetical order.

### Note:

1. The *RESET* command is not supported.
2. If the volume name is blanked out, it is possible to leave the panel, but the fields in the report will be empty.

---

## ENCLAVE - Enclave Report

The ENCLAVE report provides detailed information about the activities of enclaves.

An enclave is a transaction that can span multiple dispatchable units (SRBs and tasks) in one or more address spaces and is reported on and managed as a unit. It is managed separately from the address space it runs in. CPU and I/O resources associated with processing the transaction are managed by the transaction's performance goal and reported to the transaction.

New types of applications (for example, DDF or ICSS Webserver) create enclave transactions executing in several address spaces, but they need to be managed as own single business units of work. Therefore, a report showing resource consumption and delays by enclave will improve significantly performance management for these new applications.

## How to request this report

To request the ENCLAVE report, select 1 on the Primary Menu, and then a 6 on the Overview Report Selection Menu (shown in Figure 6 on page 25), or enter the following command:

```
ENCLAVE [subsystem-type]
```

## Contents of the report

RMF V2R2 Enclave Report										Line 1 of 16	
Command ==>										Scroll ==> HALF	
Samples: 120	System: SYS5	Date: 09/28/16	Time: 13.13.30	Range: 120	Sec						
Current options:		Subsystem Type: ALL			-- CPU Util --						
		Enclave Owner:			App1%	EApp1%					
		Class/Group:			18.1	90.1					
Enclave	Attribute	CLS/GRP	P Goal	% D X	EApp1%	TCPU	USG	DLY	IDL		
*SUMMARY					70.04						
ENC00003	CTT	PG004	1	0	18.75	26.78	12	88	0.0		
	DDF										
	JOEGEE										
ENC00001	CTT	PG004	1		16.27	23.12	11	89	0.0		
	DDF										
	JOEGEE										
ENC00004	CTT	PG004	1	F	14.83	21.12	10	90	0.0		
	DDF										
	JOEGEE										
ENC00005	CTT	PG004	1	F	14.13	20.00	8.9	91	0.0		
	DDF										
	JOEGEE										

Figure 50. ENCLAVE Report

**Note:** There may be enclave activity in your system (for example, indicated by EApp1% > App1% in the SYSINFO report), but the ENCLAVE report issues the message Enclave data is not currently available. The reason is that only those enclaves are shown in the report that have been sampled at least twice and that are active or inactive at the end of the Monitor III MINTIME. Therefore, short-running enclaves will not appear in the report.

When the report interval spans more than one Monitor III MINTIME, the above criteria must match for the last MINTIME in the report interval.

The graphic version of this report provides information about CPU utilization of the enclaves.

## Field descriptions

Table 29. Fields in the ENCLAVE Report

Field Heading	Meaning
Subheader Section - You can define the setting of each field in the Enclave Report Options panel.	
Subsystem Type	Reporting only on enclaves that belong to this subsystem type, for example, DB2 or DDF.
Enclave Owner	Reporting only on enclaves that are owned by the address space with this jobname.
Class/Group	Reporting only on enclaves that run in this service class.

Table 29. Fields in the ENCLAVE Report (continued)

Field Heading	Meaning
Appl%	Percentage of the maximum general purpose processor capacity used by all address spaces during the report interval.  This value is divided by the number of logical processors or cores that have been active during this interval.
EAppl%	Percentage of the maximum general purpose processor capacity used by all address spaces and enclaves during the report interval.  This value is divided by the number of logical processors or cores that have been active during this interval.
Enclave Identification.	
Enclave	Generated name to allow association of an enclave with instances shown on other Monitor III reports.  *SUMMARY is shown in the summary line that totals up the CPU time for the reported enclaves.
Attribute	Dynamic list of attributes. The reporter lists the attributes (maximal eight characters) in the order at which they are specified in the Enclave Report Options panel.
CLS/GRP	Service class the enclave is associated with.
P	Service class period the enclave is currently running in.
Goal	Response time goal.
%	Response time percentile or velocity.
D	Dependent enclave indication. 'Y' if the enclave is an extension of an address space transaction, otherwise blank.
X	Multi-system Indicator  This column gives an indication about the origin of the enclave: <b>O</b> The enclave originated on this system. <b>F</b> The enclave originated on another system in the sysplex but is participating on this system. <b>blank</b> The enclave is a single-system enclave.
Enclave Performance.	
EAppl%	Percentage of the maximum general purpose processor capacity consumed by the individual enclave or by all reported enclaves (in the Monitor III range).
TCPU	Total CPU time (in seconds) consumed by the enclave (see 'Detailed Performance Statistics').
USG	Percentage of total USING samples (use samples for CPU and I/O), based on #STS (total number of state samples in the enclave).
DLY	Percentage of total DELAY samples (delay samples for CPU, I/O, capping, storage, queuing) based on #STS.
IDL	Percentage of idle samples based on #STS.

## Cursor-sensitive control on the Enclave Report

There are two ways of cursor-sensitive control in the *Enclave Report* that bring up the following pop-up windows:

- “Enclave Details”
- “Enclave Classification Attributes” on page 93

### Enclave Details

If you place the cursor on one of the values in the **EAPPL%**, **TCPU**, **USG**, **DLY**, or **IDL** columns, a pop-up window is shown, containing the enclave details for the corresponding enclave.

## Mon III - ENCLAVE

RMF Enclave Details														
Details for enclave ENC00003 with token 000000B0 0000008A. Press Enter to return to the Report panel.														
- CPU Time -				-zAAP Time -				-zIIP Time -						
Total	26.78				Total	6.33				Total	0.00			
Delta	22.50				Delta	1.01				Delta	0.00			
State	---	Using%	----	----	Execution Delays%	----	IDL	UNK						
Samples	CPU	AAP	IIP	I/O	CPU	AAP	IIP	I/O	STO	CAP	QUE	IDL	UNK	
592	11	1.0	0.0	0.0	88	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.3	

Figure 51. ENCLAVE Report - Enclave Details

Table 30. Fields in the ENCLAVE Report - Enclave Details

Field Heading	Meaning
CPU Time	<p><b>Total</b> Total CPU time (in seconds) consumed by the enclave on general purpose processors and special purpose processors.</p> <p><b>Delta</b> CPU time (in seconds) consumed by the enclave on general purpose processors and special purpose processors in the reported Monitor III range.</p>
zAAP Time	<p><b>Total</b> Total CPU time (in seconds) consumed by the enclave on zAAPs.</p> <p><b>Delta</b> CPU time (in seconds) consumed by the enclave on zAAPs in the reported Monitor III range.</p>
zIIP Time	<p><b>Total</b> Total CPU time (in seconds) consumed by the enclave on zIIPs.</p> <p><b>Delta</b> CPU time (in seconds) consumed by the enclave on zIIPs in the reported Monitor III range.</p>
State Samples	Total number of state samples in the enclave.
Using% and Execution Delays% - In contrast to other Monitor III fields, these states shown in the pop-up panel are multistate. This means, they reflect the real amount of work executing in the enclave.	
All percentages are based on the number of state samples.	
Using%	Percentage of: <p><b>CPU</b> CPU using samples</p> <p><b>AAP</b> zAAP using samples</p> <p><b>IIP</b> zIIP using samples</p> <p><b>I/O</b> I/O using samples</p>

Table 30. Fields in the ENCLAVE Report - Enclave Details (continued)

Field Heading	Meaning
Execution Delays%	Percentage of: <b>CPU</b> CPU delay samples <b>AAP</b> zAAP delay samples <b>IIP</b> zIIP delay samples <b>I/O</b> I/O delay samples <b>STO</b> Storage delay samples. This includes: <ul style="list-style-type: none"> <li>• Waiting for paging I/O from common</li> <li>• Waiting for cross memory page fault</li> <li>• Waiting for shared paging</li> <li>• Server private paging delay</li> <li>• Server VIO paging delay</li> <li>• Server hiperspace paging delay</li> <li>• Server MPL delay</li> <li>• Server swap-in delay</li> </ul> <b>CAP</b> CPU capping samples <b>QUE</b> Queue delay samples
IDL	Percentage of idle samples.
UNK	Percentage of unknown samples.

### Enclave Classification Attributes

If you place the cursor on a selected enclave name in the **Enclave** column and press Enter, a pop-up window appears showing all available classification attributes for the selected enclave. If not all attributes can be displayed on one screen, you can see an indication: More: +. In this case, press PF8 to see further attributes.

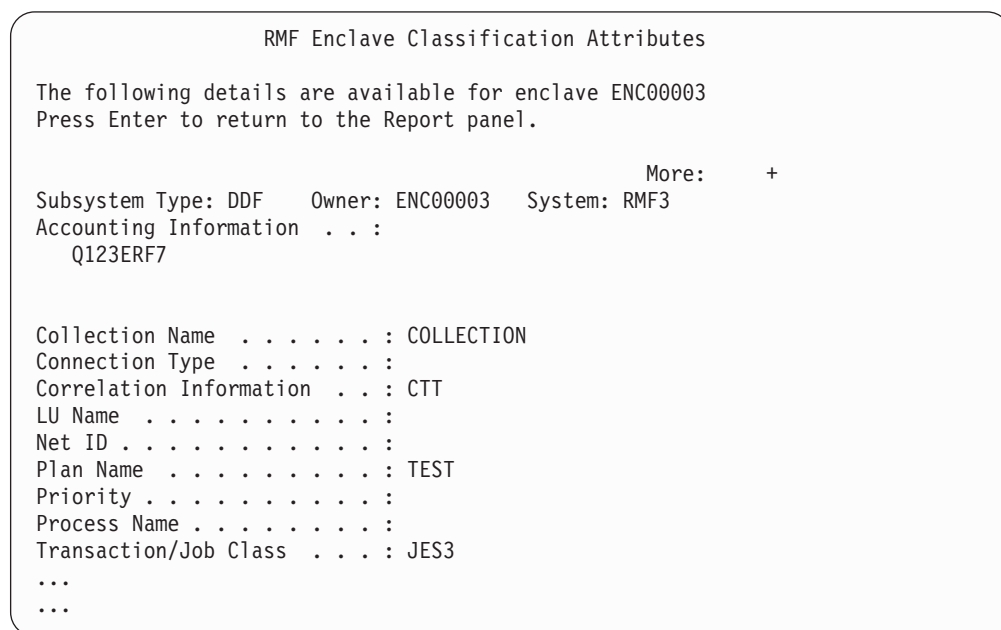


Figure 52. ENCLAVE Report - Enclave Classification Attributes (1)

### Report options

On the *ENCLAVE Report Options* menu, you can select:

- An enclave filter by one of the following criteria:

## Mon III - ENCLAVE

- Subsystem type, for example DDF, IWEB, or SOM
- Enclave owner job name, for example DB2MSTR
- Service class
- Performance group
- A list of classification attributes.

By default, the report is generated for every type of subsystem showing no attribute. Going through the options allows you to restrict the report to one subsystem only and to select only the attributes meaningful or of interest for that subsystem type. You can find details about supported attributes by subsystem type in *z/OS MVS Planning: Workload Management*.

```
RMF Enclave Report Options
Command ==>                               Scroll ==> CSR

Select one of the following options:

  1. Subsystem Type   ==> ALL           Specify a subsystem type or ALL
  2. Enclave Owner   ==>                Jobname of the enclave owner
  3. Service Class   ==>
  4. Performance Group ==>

Select (S) one or more classification attributes:

- Accounting Information           - Scheduling Environment
- Collection Name                 - Subsystem Collection Name
- Connection Type                 - Subsystem Instance
S Correlation Information         - Subsystem Parameter
- LU Name                         S Subsystem Type
- Net ID                          - Package Name
- Plan Name                       - Procedure Name
- Priority                         - Client IP Address
- Process Name                    - Client User ID
- Transaction/Job Class           - Client Transaction Name
- Transaction/Job Name           - Client Workstation/Host Name
S User ID                         - Client Accounting Information
```

Figure 53. ENCLAVE Report Options

### Subsystem Type

Report on enclaves that belong to this subsystem type, for example, DB2 or DDF.

ALL selects all active subsystems.

### Enclave Owner

Report on enclaves that are owned by the address space with this jobname.

### Service Class / Performance Group

Report on enclaves that run in this service class or performance group.

### Classification Attributes

You can select one or more classification attributes to be displayed in the *Attribute* column of the *Enclave Report* in a length of maximum eight characters.



## ENQ - Enqueue Delays Report

The Enqueue Delays report (ENQ) contains jobs waiting for a resource, the resources associated with each waiting job, and the jobs currently holding each resource. RMF lists the jobs by descending delay percentages.

### How to request this report

To request the Enqueue Delays report, select **3** from the Primary Menu, and then select **4** on the Resource Report Selection Menu (shown in Figure 8 on page 27), or enter the following command:

```
ENQ [job_class,service_class]
```

### Contents of the report

```

RMF V2R2  ENQ Delays                               Line 1 of 14
Command ==>                                         Scroll ==> HALF

Samples: 100      System: MVS1  Date: 09/28/16  Time: 10.03.20  Range: 100  Sec

Jobname  DLY  ----- Resource Waiting -----  ---- Holding ----
   %      %  STAT Major/Minor Names (Scope)  % Name/SYS  STAT
SPEWAK2 100  100  SW  SYSDSN          (SYS)          100 AMOLLOY  SO
                                     SYS1.NUCLEUS          100 SCHMATE  SO
                                     100 DRAGON   SO
                                     100 SCHUMAC  SO
SWARRENA 100  100  EW  SYSDSN          (SYS)          100 D71SJH1  SO
                                     SYS1.NUCLEUS          100 AMOLLOY  SO
                                     100 SCHMATE  SO
                                     100 DRAGON   SO
                                     100 SCHUMAC  SO
LUCKYSM   1    1  EW  SYSZVVD        (SYS)           1 STEVEB   EO
+CATALOG          SYS1.MVS3.MCAT
*MASTER*  1    1  EW  SYSIKJBC       (SYSS)          1 *MASTER* EO
                                     PETEG
    
```

Figure 54. ENQ Report

The graphic form of this report shows the percentage of each user's time spent waiting for a resource.

### Field descriptions

Table 31. Fields in the ENQ Report

Field Heading	Meaning
Jobname	Name of a job that is waiting for a resource. The ENQ delay report does not summarize data by job groups; all jobs within a job group are reported individually. If the catalog system address space is processing a catalog request on behalf of the job that is enqueued on a resource, the jobname of the catalog address space, (usually CATALOG) will appear below the jobname preceded by a +.

## Mon III - ENQ

Table 31. Fields in the ENQ Report (continued)

Field Heading	Meaning
DLY %	<p>Delay the waiting job is experiencing because of contention for any enqueued resource during the report interval. This value is calculated as follows:</p> $\text{DLY \%} = \frac{\text{Delay samples}}{\text{\# Samples}} * 100$ <p><b>Delay samples</b> The number of samples when the job was delayed for one or more enqueued resources.</p> <p><b>Note:</b> This DLY% value is also found in the ENQ field on the job delay report.</p>
Resource Waiting %	<p>Indicates how much of the overall delay of the job for enqueued resources is caused by a specific resource. This value is calculated as follows:</p> $\text{Waiting \%} = \frac{\text{Delay samples}}{\text{\# Samples}} * 100$ <p><b>Delay samples</b> The number of samples when the job was delayed for the resource.</p> <p><b>Note:</b> If there is no overlap in delay states, the WAITING% value(s) for a job add up to the DLY % value of the job.</p>
Resource Waiting STAT	<p>The status indicates whether the waiting job wants exclusive (EW) or shared (SW) use of the resource.</p>

Table 31. Fields in the ENQ Report (continued)

Field Heading	Meaning
Major/Minor Names	<p>The Major name and Minor name of the resource delaying the job The major name is listed above the minor name. The major name is up to eight characters long, and the minor name is up to 36 characters long. If the minor name contains unprintable characters it will be up to 18 characters long (represented by 36 hexadecimal digits). If the minor name is longer than 36 characters, RMF only displays the first 36 characters. If there are two resources with the same major name and their minor name differs only after the first 36 characters, then RMF considers them as the same resource. Shown on the same line as the major name, Scope shows whether the scope of the resource is system (SYS) or systems (SYSS). It is possible that two resources with the same major and minor name, but different scopes, might exist in the system.</p> <p>The following are the most common enqueue major names and their associated resources:</p> <p><b>Major Name</b>  <b>Resources</b>  <b>MSFDC</b> Service processor TP port  <b>SPFDSN</b>              Data set name  <b>SPFEDIT</b>              Data set name  <b>SYSDSN</b>              System data sets  <b>SYSIAT</b> JES3 CHKPNT data set  <b>SYSIEA01</b>              Dump data set  <b>SYSIEFSD</b>              Serializes device allocations  <b>SYSIEWLP</b>              SYSLMOD data set (Minor name is data set name)  <b>SYSIGGV1</b>              Master catalog  <b>SYSIGGV2</b>              Catalogs (Minor name is catalog name)  <b>SYSIKJBC</b>              TSO broadcast data set (Minor name is relative block address)  <b>SYSIKJUA</b>              User attribute data set  <b>SYSSEMF01</b>              SMF SYS1.MANx data set  <b>SYSVSAM</b>              VSAM data sets  <b>SYSVTOC</b>              VTOC (Minor name is volser)  <b>SYSAVM</b>              AVM queue or data areas  <b>SYSZBDT</b>              z/OS bulk data transfer (Minor name is node name)  <b>SYSZCAXW</b>              Catalog auxiliary work area  <b>SYSZCMD5</b>              Master trace command or Message loss detection  <b>SYSZCOMM</b>              Global Resource Serialization ring processing table  <b>SYSZCSD</b>              CSD control block field  <b>SYSZEC16</b>              Purge data set  <b>SYSZIGGI</b>              TSB (Minor name is ASID)  <b>SYSZIST0C</b>              Configuration restart data set (Minor name is ddname)  <b>SYSZJES2</b>              JES2 buffer or data set  <b>SYSZJWTP</b>              Job step messages</p>

Table 31. Fields in the ENQ Report (continued)

Field Heading	Meaning
Major/Minor Names continued	<p><b>Major Name</b></p> <p><b>Resources</b></p> <p><b>SYSZOPEN</b> System data sets</p> <p><b>SYSZPCCB</b> Private catalog control block</p> <p><b>SYSZPGAD</b> PAGEADD command</p> <p><b>SYSZPSWD</b> Password data set</p> <p><b>SYSZRPLW</b> Catalog name (Minor name is catalog name)</p> <p><b>SYSZSIPS</b> SYSEVENT</p> <p><b>SYSZSMF1</b> SMF buffer</p> <p><b>SYSZTIOT</b> Device allocation</p> <p><b>SYSZTRC</b> System trace</p> <p><b>SYSZUSRL</b> User label tracks</p> <p><b>SYSZVARY</b> Reconfiguration commands</p> <p><b>SYSZVMV</b> Volume mount and verify</p> <p><b>SYSZVOLS</b> System volumes (Minor name is volser)</p> <p><b>SYSZWTOR</b> WTOR reply (Minor name is REPLYxx, where xx is the message ID)</p>
Holding %	<p>Indicates how much a specific job is contributing to the holding of a resource. The value is expressed as a percentage. For example, a Holding % of 100 indicates that the specified job was enqueued on the resource and delaying the waiting job for the entire report interval. This value is calculated as follows:</p> $\text{Holding \%} = \frac{\text{Holding samples}}{\text{\# Samples}} * 100$ <p><b>Holding samples</b> The number of samples when the holding job held the resource and the delayed job was waiting for it. For primary source fields used in this calculation see the DELAY % field in this report description.</p>
Holding Name/SYS	<p>The name of the job that is holding the resource that the delayed job is waiting for. If the holding job is from another system, RMF also provides the system name (global resource serialization system identifier) which will appear below the holding jobname preceded by a /. If the catalog system address space is processing a catalog request on behalf of the job that is holding the resource that the delayed job is waiting for, the jobname of the catalog address space (usually CATALOG) will appear below the jobname preceded by a +.</p>
Holding STAT	<p>The status indicates whether the holding job has exclusive (EO) or shared (SO) use of the resource.</p>

### Report options

The ENQ Report Options panel is similar to the Device Report Options panel. See Figure 39 on page 72 for an example. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed. See Figure 37 on page 70 for an example.

## ENQR - Enqueue Resource Delays Report

The Enqueue Resource Delays report (ENQR) is similar to the Enqueue Delays report, but the information about a specific resource is kept together. RMF reports the resources according to the number of waiting jobs in descending order, the jobs waiting for each resource in descending delay percentage order, and the jobs holding the resource in descending holding percentages.

### How to request this report

To request the ENQR report, select **3**, and then select **5** on the Resource Report Selection Menu (shown in Figure 8 on page 27), or enter the following command:

```
ENQR [resourcename]
```

### Contents of the report

```

RMF V2R2  ENQ Resource Delays                               Line 1 of 14
Command ==>                                               Scroll ==> HALF

Samples: 100      System: MVS1  Date: 09/28/16  Time: 10.03.20  Range: 100  Sec

----- Resource Name -----      ---- Delayed ----      ---- Holding ----
Major/Minor      (Scope)              % Name      STAT      % Name/SYS  STAT
SYSDSN          (SYS)              100 SPEWAK2  SW      100 AMOLLOY  SO
SYS1.NUCLEUS                                100 SWARRENA  EW      100 SCHMATE  SO
                                                    100 DRAGON   SO
                                                    100 SCHUMAC  SO
                                                    100 D71SJH1  SO
                                                    100 AMOLLOY  SO
                                                    100 SCHMATE  SO
                                                    100 DRAGON   SO
                                                    100 SCHUMAC  SO
                                                    100 D71SJH1  SO
SYSIKJBC        (SYSS)              1 *MASTER*  EW      1 *MASTER*  EO
PETEG
SYSZVDS         (SYS)              1 LUCKYSM   EW      1 STEVEB    EO
SYS1.MVS3.MCAT      +CATALOG          +CATALOG
    
```

Figure 55. ENQR Report

The graphic form of this report shows the average number of active users for waiting for each resource.

### Field descriptions

Table 32. Fields in the ENQR Report

Field Heading	Meaning
Resource Name	The Major name and Minor name of the resource delaying the job. The major name is listed above the minor name. The major name is up to eight characters long and the minor name is up to 36 characters long. If the minor name contains unprintable characters, it will be up to 18 characters long (represented by 36 hexadecimal digits). If the minor name is longer than 36 characters, RMF only displays the first 36 characters. If there are two resource with the same major name and their minor name differs only after the first 36 characters, then RMF considers them as the same resource. Shown on the same line as the major name, SCOPE shows whether the scope of the resource is system (SYS) or systems (SYSS). It is possible that two resources with the same major and minor name, but different scopes, might exist in the system.

Table 32. Fields in the ENQR Report (continued)

Field Heading	Meaning
Delayed %	<p>The delay percentage of the job for a specific enqueued resource. This value is calculated as follows:</p> $\text{Delayed \%} = \frac{\text{\# Delay Samples}}{\text{\# Samples}} * 100$ <p><b>Delay samples</b> The number of samples when the job was delayed for a specific enqueued resource. RMF calculates the number of samples delayed by incrementing a counter once for each sample when the job is delayed for that resource.</p>
Delayed Name	Name of the job delayed for the resource. RMF lists all jobs delayed for the resource. If the catalog system address space is processing a catalog request on behalf of the delayed job, the jobname of the catalog address space (usually CATALOG) will appear below the jobname preceded by a +.
Delayed STAT	The status indicates whether the waiting job wants exclusive (EW) or shared (SW) use of the resource.
Holding %	<p>The percent of the range that a specific job was holding the resource while the named job was delayed. For example, a Holding % of 100 indicates that the specified job held the resource for the entire range period. This value is calculated as follows:</p> $\text{Holding \%} = \frac{\text{\# Holding Samples}}{\text{\# Samples}} * 100$ <p><b>Holding samples</b> The number of samples when the holding job was holding the resource while the named job was delayed.</p>
Holding Name	The name of the job that is holding the resource that the delayed job is waiting for. If the holding job is from another system, RMF also provides the system name (global resource serialization system identifier) which will appear below the holding jobname preceded by a /. If the catalog system address space is processing a catalog request on behalf of the job that is holding the resource that the delayed job is waiting for, the jobname of the catalog address space (usually CATALOG) will appear below the jobname preceded by a +.
Holding STAT	The status indicates whether the holding job has exclusive (EO) or shared (SO) use of the resource.

### Report options

```

RMF ENQR Report Options                               Line 1 of 1
Command ==>>                                         Scroll ==>> HALF

Change or verify parameters for the ENQR report. To exit press END.

Major ==>> SYSIEFSD   ENQ major name for report or ALL

                               Available ENQ Major Names
CLRLOG00  SYSIEFSD  SYSZVVDS
    
```

Figure 56. ENQR Report Options Panel

The Report Options panel allows you to select from a list of available major names, resources to be included in the report.

For MAJOR, specify the major name of the serially reusable resource for which you want information, or ALL for information about all serially reusable resources in the system. The major name you specify is saved across sessions in the current option set.

A list of all serially reusable resources that had any enqueue contention during the current report interval appears under Available ENQ Major Names.

**Major** The major name of the serially reusable resource that you want information about.

Enter all, for information about all serially reusable resources with enqueue contention during the report interval, one of the names listed under AVAILABLE ENQ MAJOR NAMES, or the major name of a serially reusable resource that might experience contention at a later time.

A resource only appears on the Enqueue Resource Report when it experiences enqueue contention during the report interval.

Your selection is saved across sessions in the current option set.

#### **Available ENQ Major Names**

The list of the serially reusable resources that had enqueue contention during the current report interval.

## **GROUP - Group Response Time Report**

The Group Response Time (GROUP) report presents information about using and delay values for a specific service or report class. The using and delay values are average values for all transactions processed during the report interval. The report presents the total using and delay value and a breakdown of this total value into each defined resource.

### **How to request this report**

To request the Group Response Time report, select **1** on the Primary Menu, and then select **5** on the Overview Report menu (shown in Figure 6 on page 25), or enter one of the following commands:

```
GROUP service_class,period
```

```
GROUP report_class,period
```

For example, to get a Group Response Time report for first period of service class HOTBATCH, enter:

```
GROUP HOTBATCH,1
```

## Contents of the report

```

RMF V2R2 Group Response Time
Command ==>                               Scroll ==> HALF
Samples: 100   System: MVS1 Date: 09/28/16 Time: 10.03.20 Range: 100 Sec
Class: HOTBATCH   Period: 1   Description: Very Important Batch
Primary Response Time Component: Using the processor

WFL   Users   Frames   Vector   EXCP   PGIN   TRANS   --- Response Time ---
%   TOT ACT   %ACT   UTIL   Rate   Rate   Rate   -- Ended TRANS--(Sec) -
100   1   0   2   0   2.7   0.1   0.117   0.302   1.447   1.749

                                -AVG USG-   -----Average Delay-----
                                Total  PROC DEV  PROC DEV STOR SUBS OPER ENQ OTHER
Average Users                   0.169  0.08 0.03  0.00 0.00 0.00 0.00 0.00 0.00 0.05
Response Time ACT 1.749  0.82 0.31  0.00 0.00 0.00 0.00 0.00 0.00 0.52

                                ---STOR Delay---   ---OUTR Swap Reason---   ---SUBS Delay---
                                Page Swap OTR   TI TO LW XS   JES HSM XCF
Average Users                   0.00 0.00 0.00  0.00 0.00 0.00 0.00 0.00 0.00 0.00
Response Time ACT 0.00 0.00 0.00  0.00 0.00 0.00 0.00 0.00 0.00 0.00
    
```

Figure 57. GROUP Report

There is no graphic version of this report available.

If you place the cursor on any of the fields named **WAIT**, **EXECUT**, or **ACTUAL**, the pop-up panel appears showing a detailed breakdown of the different wait reasons and their average duration.



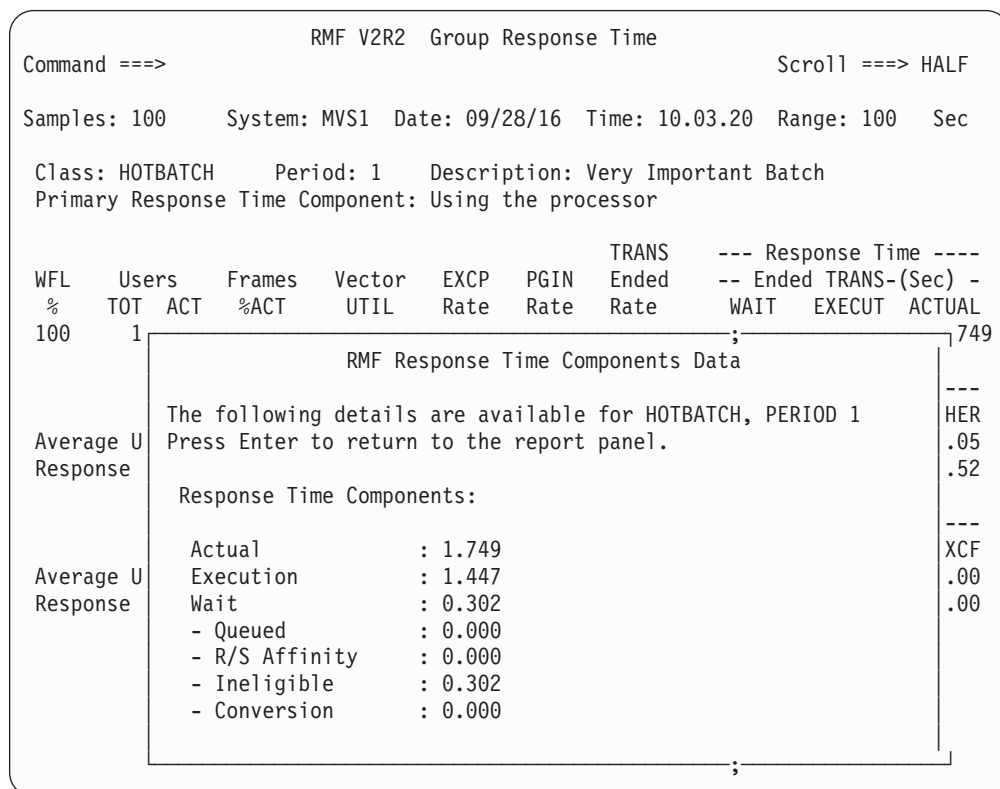


Figure 58. GROUP Report - Response Time Components

### Field descriptions

Table 33. Fields in the GROUP Report

Field Heading	Meaning
Class	The name of the service or report class
Period	The period number
Description	The description of the specified class, it is derived from the service policy.
Primary Response Time Component	<p>A description of the component contributing to the largest percentage of total response time. The description can be:</p> <ul style="list-style-type: none"> <li>• Using the processor</li> <li>• Using I/O devices</li> <li>• Processor delay</li> <li>• Device delay</li> <li>• Storage delay for               <ul style="list-style-type: none"> <li>- common paging</li> <li>- local paging</li> <li>- virtual I/O</li> <li>- XMEM</li> <li>- HIPR</li> <li>- swap in</li> <li>- out and ready</li> </ul> </li> <li>• Waiting for               <ul style="list-style-type: none"> <li>- JES</li> <li>- HSM</li> <li>- XCF</li> </ul> </li> <li>• Waiting for volume mount</li> <li>• Waiting for operator reply</li> <li>• ENQ serialization delay</li> <li>• Delayed for unmonitored reasons</li> </ul>

## Mon III - GROUP

Table 33. Fields in the GROUP Report (continued)

Field Heading	Meaning
WFL %	The workflow percentage of the specified class. A value of 100% indicates no workload contention, while a value of 0% indicates that all requests for system resources are delayed.
Users	<p>The number of users in the specified class. This category includes the following headings:</p> <p><b>TOT</b> Total number of users equals the number of different users found in all address spaces for the specified class during the report interval.</p> <p><b>ACT</b> Average number of active users found in all address spaces for the specified class during the report interval.</p> <p>An active user is either using a monitored resource, delayed for a monitored resource, or performing an activity that RMF does not measure.</p> <p>See the definition of Users/Active under "WFEX - Workflow/Exceptions Report" on page 210 for more details.</p>
Frames %ACT	<p>The percentage of central storage frames used by active users in the specified class during the report interval. For a definition of active users, see Users/Active under "WFEX - Workflow/Exceptions Report" on page 210.</p> <p>RMF accumulates the number of central storage frames for all active users during the report interval, then calculates the percentage as follows:</p> $\text{Frames \%ACT} = \frac{\text{ACSF}}{\text{OCSF}} * 100$ <p><b>ACSF</b> Accumulated central storage frames</p> <p><b>OCSF</b> Online central storage frames</p>
Vector UTIL	The vector time for the specified class as a percentage of total system vector capability. The field contains data only when measured on a system with a vector processor online, otherwise the field contains zeros.
EXCP Rate	<p>The rate of EXCP requests per second for the specified class:</p> $\text{EXCP Rate} = \frac{\sum \text{all EXCP Requests}}{\text{Range Time}}$
PgIn Rate	<p>The rate at which pages are being swapped:</p> $\text{PgIn Rate} = \frac{\sum \text{Page-in Counts for Class}}{\text{Range Time}}$
TRANS Ended Rate	<p>The average number of ended transactions per second that occurred for the specified class during the report interval:</p> $\text{TRANS Ended Rate} = \frac{\# \text{ Ended Transactions}}{\text{Range Time}}$

Table 33. Fields in the GROUP Report (continued)

Field Heading	Meaning
Response Time	<p>The average response time for all transactions that ended during the report interval. The field is divided into WAIT, EXECUT, and ACTUAL response time. The time a job was delayed due to TYPRUN=HOLD or TYPRUN=JCLHOLD is NOT included in any of the transaction times.</p> <p><b>Note:</b> The response times reported are for ended transactions only. If there is a delay while the transaction is queued or running, the problem will not be reported until after the transaction has ended. The WFL % field and the Average Users line can be used to identify the bottleneck.</p> <p><b>WAIT</b> The average time (in seconds) that a transaction spent waiting because of one of these reasons:</p> <ul style="list-style-type: none"> <li>• <b>Queued:</b> Average time a job was delayed for reasons other than the ones mentioned below. This field therefore basically includes the time a job was delayed for initiation. For TSO users, this can be a portion of LOGON processing. For APPC, this is the time the transaction spent on an APPC queue.</li> <li>• <b>R/S Affinity - Resource affinity scheduling delay:</b> Average time the job was delayed due to resource or system affinity scheduling. This means that resource(s) required for the job to run were not available at some point while the job was queued to JES2.</li> <li>• <b>Ineligible - Operational or JES scheduling delay:</b> Average time a job was delayed due to operational delays or JES scheduling delays, examples are: <ul style="list-style-type: none"> <li>- Job held by operator</li> <li>- Job class or job queue held</li> <li>- Duplicate jobname serialization</li> <li>- Job class execution limits</li> </ul> </li> <li>• <b>Conversion - JCL conversion delay:</b> Average time a job was delayed for JCL conversion. Jobs held during conversion (due to affinity, HSM recall, or enqueue contention) contribute only to conversion time, not to ineligible or R/S affinity times. Conversion time is not part of the total response time.</li> </ul> <p><b>EXECUT</b> The average time that a transaction was active in the system.</p> <p><b>ACTUAL</b> The sum of the execution time and the wait time, but does not include conversion time.</p> <p>If you place the cursor on one of these fields and press Enter, a <b>Response Time Components Data</b> pop-up panel will show a detailed breakdown of the different wait reasons and their average duration.</p>

## Mon III - GROUP

Table 33. Fields in the GROUP Report (continued)

Field Heading	Meaning
Average Users	<p>The average number of active users in the class during the report interval. The Average Users line is displayed in dark blue to differentiate it from the Response Time ACT line below it. The line is divided into:</p> <p><b>Total</b> Average number of active users in the class. An active user is either using a resource or is delayed by a resource and includes unmonitored reasons reported in the 'OTHER' column.</p> <p><b>Total</b> can be less than the sum of the individual categories if a user was found using or delayed in more than one category.</p> <p><b>AVG USG</b> Average number of users is summarized for the specified class. RMF takes the sum of using samples for the address space(s) associated with the class and divides by the number of samples. The average number of users is reported for the following categories:  <b>PROC</b> The average number of users using the processor during the report interval.  <b>DEV</b> The average number of users using devices during the report interval.</p> <p><b>Average Delay</b> Average number of delayed users is summarized for the specified class. RMF takes the sum of delay samples for the address space(s) associated with the class and divides by the number of samples in the Range.</p> <p>The average number delayed for is reported for the following categories:  <b>PROC</b> Waiting for a processor  <b>DEV</b> Waiting for a DASD or tape  <b>STOR</b> Waiting for a COMM, LOCL, SWAP, XMEM, HIPR, or VIO page, or on the out/ready queue  <b>SUBS</b> Waiting for services from JES, HSM, or XCF  <b>OPER</b> Waiting for the operator to reply to a message or mount a tape  <b>QUIESCE</b>  The operator has quiesced the address space. A quiesced address space can show unexpected data: <ul style="list-style-type: none"> <li>• A swappable address space will be swapped out, thus it can be OUTR and show storage delays.</li> <li>• A non-swappable address space will get lowest priority, thus it can show CPU delay, paging delay, or other delays, and even some USG % from time to time depending on the load on the system.</li> </ul> Cursor-sensitive control on this field gives you the Quiesce delay variation of the Job Delay report.  <b>ENQ</b> Waiting to use serially reusable resources that other jobs were using  <b>OTHER</b> Unknown time</p> <p><b>STOR Delay</b> A breakdown of the <b>Average Delay - STOR</b> field into the average number of users delayed for paging (Page), swapping (Swap), and swapped out and ready (OUTR). See the % <b>Delayed for</b> field description in Table 64 on page 155 for more information about the storage delays.</p>

Table 33. Fields in the GROUP Report (continued)

Field Heading	Meaning
Average Users (continued)	<p><b>OUTR Swap Reason</b>            A breakdown of the <b>STOR Delay OUTR</b> field into the average number of users delayed for specific swap reasons. The swap reasons are sorted by descending swap count; that is, the swap reason having the largest swap count is reported first. The report always displays four swap reason headings.</p> <p>The swap reasons can be:</p> <ul style="list-style-type: none"> <li><b>TI</b> Terminal input wait</li> <li><b>TO</b> Terminal output wait</li> <li><b>LW</b> Long wait</li> <li><b>XS</b> Auxiliary storage shortage</li> <li><b>RS</b> Real storage shortage</li> <li><b>DW</b> Detected long wait</li> <li><b>RQ</b> Requested swap</li> <li><b>NQ</b> Enqueue exchange swap</li> <li><b>EX</b> Exchange swap</li> <li><b>US</b> Unilateral swap</li> <li><b>TS</b> Transition swap</li> <li><b>IC</b> Improve central storage usage</li> <li><b>IP</b> Improve system paging rate</li> <li><b>MR</b> Make room for an out-too-long user</li> <li><b>AW</b> APPC wait</li> <li><b>IW</b> OMVS input wait</li> <li><b>OW</b> OMVS output wait</li> <li><b>SR</b> In-real swap</li> </ul> <p><b>SUBS Delay</b>            A breakdown of the <b>Average Delay - SUBS</b> field into the average number of users delayed by each subsystem (JES, HSM, and XCF). The subsystems are sorted by descending delay count; that is, the subsystem causing the largest delay is reported first.</p>

Mon III - GROUP

Table 33. Fields in the GROUP Report (continued)

Field Heading	Meaning
Response Time ACT	<p>The average response time (in seconds) spent in each delay or using category for a transaction that was active during the report interval. The <b>Response Time ACT</b> line is divided into:</p> <p><b>Total</b> Average time (in seconds) that an ended transaction was active in the system. The value is the same as the <b>Response Time, Ended TRANS (Sec), Active</b> value.  <b>Note:</b> The value for Total and the sum of the individual using and delay values can be different. <b>Total</b> represents only ended transactions, while the breakdown of using and delay values represents all active transactions during the report interval.</p> <p>In addition, <b>Total</b> can be less than the sum of the individual categories if some of the users are delayed and using at the same time.</p> <p><b>AVG USG</b>  Average time (in seconds) that a transaction was using a processor (PROC) or device (DEV) during the report interval.</p> <p><b>Average Delay</b>  Average time (in seconds) that a transaction was delayed for the following reasons:  <b>PROC</b> Waiting for a processor  <b>DEV</b> Waiting for a DASD or tape  <b>STOR</b> Waiting for a COMM, LOCL, SWAP, XMEM, HIPER, or VIO page, or on the out/ready queue  <b>SUBS</b> Waiting for services from job-entry subsystem (JES), Hierarchical Storage Manager (HSM), or Cross-System Coupling Facility (XCF)  <b>OPER</b> Waiting for the operator to reply to a message or mount a tape  <b>QUIESCE</b>  The operator has quiesced the address space. Cursor-sensitive control on this field gives you the Quiesce delay variation of the Job Delay report.  <b>ENQ</b> Waiting to use serially reusable resources that other jobs were using  <b>OTHER</b> Unknown time</p> <p><b>STOR Delay</b>  Breakdown of the <b>Average Delay - STOR</b> field into the average time (in seconds) that a transaction was delayed for paging (Page), swapping (Swap), and swapped out and ready (OUTR). See the % Delayed for field description in Table 64 on page 155 for more information about the storage delays.</p> <p><b>OUTR Swap Reason</b>  Breakdown of the <b>STOR Delay OUTR</b> field into the average time (in seconds) that a transaction was delayed for specific swap reasons. The report always displays four swap reason headings. The four swap reasons are determined by the <b>Average Users OUTR Swap Reason</b> field. See the field description for <b>Average Users OUTR Swap Reason</b> for a list of possible swap reasons.</p> <p>See the field descriptions for the Monitor I "PAGING - Paging Activity report" on page 418 for detailed information about the swap reasons.</p> <p><b>SUBS Delay</b>  Breakdown of the <b>Average Delay - SUBS</b> field into the average time (in seconds) that a transaction was delayed by each subsystem (JES, HSM, and XCF). The subsystems are sorted by descending delay count; that is, the subsystem causing the largest delay is reported first.</p>
Statistical error	<p>The measure of the statistical validity of the reported data. The statistical error heading and value only appear if the value is 25% or higher. The field is displayed as a warning to you that data reported is not statistically significant.</p> <p>Because you want to ensure that the data reported is an accurate representation of what happened during the report interval, the data collected should include response times that are short compared to the report interval. and a large enough number of ended transactions to provide a statistically significant sample.</p> <p>To decrease statistical error, try increasing the report interval.</p>

## Report options

```

RMF Service/Report Class Options: GROUP      Line 1 of 14
Command ==>                                Scroll ==> HALF

Change or verify parameters. To exit press END.
Changes will apply to SYSRTD, SYSWKM and GROUP reports.

Type      ==> S          Service or Report class (S R)
Class     ==> HOTBATCH  Class name
Period    ==> 3         Period number

                Available Service and Report Classes
BATCHHI  S  BATCHLOW S  BATCHMED S  BATCHRSP S  BATCH1  S  BATCH2  S
BTCHDEF  S  DISCRET  S  ENC_SC   S  ENCLAVE  S  ENCRMF_L S  ENCRMF_U S
IRLM     S  ZOSUN    S  OEICTWLM S  OMVS     S  OMVSKERN S  OTHDONRS S
OTHTRAN  S  PRDBAT   S  PRDBATHI S  PRDTSO  S  RESP20  S  RESP60  S
RMF      S  RMFGAT   S  SERVERS  S  STCCMD  S  STCDEF  S  STCLO   S
STORPROC S  SYSOTHER S  SYSSTC   S  SYSTEM  S  TSODEF  S  TSOEVEN S
TSOHI    S  TSOMED   S  TSOODD  S  TSOREG  S  TSOSLOW S  TSTBATHI S
TSTBATLO S  TSTBATMD S  APPC     R  ASCH    R  BCP     R  BERD    R
MASTER  R  MVSNFS   R  ZOSUN   R  OMVS    R  RDXXIRLM R  REPORT  R
RMF      R  RMFGAT   R  RPTDUMP R  RPTSTC  R  RRS     R  RSM     R
THRASHER R  TPNS     R  WLM     R  XCFAS   R

```

Figure 59. GROUP Report Options Panel

**Type** Here you specify whether you want to select a service or report class.

**Class** The service or report class for which you want data reported. You can specify any of the classes listed under Available Service and Report Classes.

If the class you want is not listed, it was not active during the current report interval. If you specify the class, it will appear on the report when it is available.

Your selection applies to all delay and common storage reports and is saved across sessions in the current option set.

**Period** Enter the number (between 1 and 8) of the period you want reported.

### Available Service and Report Classes

This list includes all service and report classes that had any activity during the current report interval.

---

## HSM - Hierarchical Storage Manager Delays Report

The Hierarchical Storage Manager (HSM) Delays report allows you to investigate situations where jobs are delayed when requesting service from HSM.

RMF lists all jobs delayed during the refresh period in order by descending delay percentage.

### How to request this report

To request the HSM report, select **4** on the Primary menu, and then select **1** on the Subsystem Report menu (shown in Figure 9 on page 27), or enter the following command:

```
HSM [job_class,service_class]
```

## Contents of the report

```

RMF V2R2  HSM Delays                               Line 1 of 3
Command ==>>>                                     Scroll ==>> HALF

Samples: 100      System: MVS1  Date: 09/28/16  Time: 10.03.20  Range: 100  Sec

Jobname      DLY  ----- Main Delay Reason -----
              %   % F-Code Explanation
AUDTRPTZ    94  94   3  Dataset recall from auxiliary storage.
APETER      82  82   3  Dataset recall from auxiliary storage.
TJSMITH     77  77   3  Dataset recall from auxiliary storage.
    
```

Figure 60. HSM Report

The graphic form of this report shows the percentage of each user's time spent waiting for HSM services.

### Field descriptions

Table 34. Fields in the HSM Report

Field Heading	Meaning
Jobname	Name of the job delayed when requesting service from HSM. The HSM Delays report does not summarize data by job groups; all jobs within a job group are reported individually.
DLY %	<p>Delay the waiting job is experiencing because of contention for HSM during the report interval. This value is calculated as follows:</p> $\text{DLY \%} = \frac{\text{\# Delay Samples}}{\text{\# Samples}} * 100$ <p><b>Delay samples</b>                      The number of samples when the job was delayed by HSM. RMF calculates this value by incrementing its counter once for each sample when one or more units of work in the address space had HSM delay. RMF considers the user delayed if all of the following conditions are met:</p> <ul style="list-style-type: none"> <li>• The MWE has a request from the user</li> <li>• The request is a "waited-on" request</li> <li>• The request has not completed processing</li> <li>• You receive one of the function codes listed under Main Delay Reason(s).</li> </ul> <p><b>Note:</b> This DLY % value is also found in the HSM field on the job delay report.</p>



Table 34. Fields in the HSM Report (continued)

Field Heading	Meaning																
Main Delay Reason(s)	<p>The subsystem function code that indicates the main reason for the delay. RMF reports the one or two function codes with the highest counts as the main delay reasons.</p> <p>% indicates how much of the HSM delay of the job is caused by the reported subsystem function. This value is calculated as follows:</p> $\% = \frac{\text{\# Delay Samples}}{\text{\# Samples}} * 100$ <p><b>Delay samples</b> The number of samples when the job was delayed for HSM for a specific subsystem function.</p> <p><b>Note:</b> The Main Delay Reason % values add up to the DLY % value of the job if there is no overlap in delay states and there are no more than two function codes responsible for the delay.</p> <p>The HSM F-codes (in decimal) and their explanations are as follows:</p> <table border="0"> <thead> <tr> <th>F-Code</th> <th>Explanation</th> </tr> </thead> <tbody> <tr> <td>03</td> <td>A data set is being recalled from auxiliary storage.</td> </tr> <tr> <td>05</td> <td>A data set is being recovered.</td> </tr> <tr> <td>06</td> <td>A data set is being migrated.</td> </tr> <tr> <td>07</td> <td>A data set is being backed up.</td> </tr> <tr> <td>08</td> <td>A control data set record is being read.</td> </tr> <tr> <td>08</td> <td>A JES3 C/I locate is being done.</td> </tr> <tr> <td>12</td> <td>A data set is being deleted.</td> </tr> </tbody> </table>	F-Code	Explanation	03	A data set is being recalled from auxiliary storage.	05	A data set is being recovered.	06	A data set is being migrated.	07	A data set is being backed up.	08	A control data set record is being read.	08	A JES3 C/I locate is being done.	12	A data set is being deleted.
F-Code	Explanation																
03	A data set is being recalled from auxiliary storage.																
05	A data set is being recovered.																
06	A data set is being migrated.																
07	A data set is being backed up.																
08	A control data set record is being read.																
08	A JES3 C/I locate is being done.																
12	A data set is being deleted.																

### Report options

The HSM Report Options panel is similar to the Device Report Options panel. See Figure 39 on page 72 for an example. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed. See Figure 37 on page 70 for an example.

---

## IOQUEUE - I/O Queuing Activity Report

The I/O Queuing Activity report (IOQUEUE) provides information, grouped by LCU (logical control unit), on the I/O configuration. The information includes contention rate, queue lengths, and percentages of time when one or more I/O components were busy. Information about the LCU is useful because the LCU is the focus of I/O configuration and path management measurements for a related group of I/O devices.

For all channels that are managed by **Dynamic Channel Path Management (DCM)**, additional information is available. DCM allows an installation to identify channels which they wish to be managed dynamically. These channels are not assigned permanently to a specific control unit, but belong to a pool of channels. Based on workload requirements in the system, these channels will be assigned dynamically by DCM. For each LCU with DCM managed channels, a summary line displays the minimum and maximum number of connected DCM managed channels, the number of defined DCM managed channels and accumulated activity data.

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 (System Assist Processor (SAP)), which is part of the channel subsystem, manages and schedules I/O work requests to the various devices within the LCU.

This report can tell you about the cause of performance problems associated with channel paths and devices. You could, for example, find the reason for an unusually long pending time reported on the device report. Check the relationship between the percentage of requests deferred for device busy and control unit busy for the LCU on the I/O Queuing Activity report.

### How to request this report

To request the I/O Queuing Activity report, select 3 from the Primary Menu, and then select 13 on the Resource Report Selection Menu (shown in Figure 8 on page 27), or enter the following command:

```
IOQUEUE
```

### Special considerations of report output

In a VM guest system environment, the report for a z/OS system that is authorized via the VM RMCHINFO directory option, shows static configuration data, only. Measurement data is not available.

Data items that are not valid are marked by dashes (---) in the output display.

**Note:** The report contains data for DASD control units only.

### Contents of the report

RMF V2R2 I/O Queuing Activity										Line 1 of 54			
Command ==>										Scroll ==> HALF			
Samples: 30		System: S5C		Date: 09/28/16		Time: 03.23.30		Range: 30 Sec					
Path	DCM	CTL	Units	DCM Group MN MX DEF	LCU	Cont Rate	De1 Lngth	Q CSS	AVG CHPID Taken	%DP Busy	%CU Busy	AVG CUB	AVG CMR
D7		5F00			0048				1.13	0.0	0.0	0.0	0.2
D6		5F00			0048				0.97	0.0	0.0	0.0	0.2
					0048	0.0	0.00	0.3	2.10	0.0	0.0	0.0	0.2
B0	PF	8000			0069				82.17	0.0	0.0	0.0	0.2
B1	PF	8000			0069				83.83	0.0	0.0	0.0	0.2
B2	NP	8000			0069				0.00	0.0	0.0	---	---
B3	NP	8000			0069				0.00	0.0	0.0	---	---
95	PF	8000			0069				83.17	0.0	0.0	0.0	0.2
					0069	0.0	0.00	0.4	249.17	0.0	0.0	0.0	0.2
B0	NP	8100			006A				0.00	0.0	0.0	---	---
B1	NP	8100			006A				0.00	0.0	0.0	---	---
B2	PF	8100			006A				124.53	0.0	0.0	0.0	0.2
B3	PF	8100			006A				124.87	0.0	0.0	0.0	0.2

Figure 61. IOQUEUE Report

The graphic form of this report shows the contention rate of each LCU.

### Field descriptions

Table 35. Fields in the IOQUEUE Report

Field Heading	Meaning
Path	<p>The two-digit hexadecimal channel path identifier (CHPID) of the online channel path attached to the physical control units in the LCU. There can be up to eight channel paths in an LCU.</p> <p>If applicable, the path attribute is indicated with the CHPID:  <b>PF</b> preferred path  <b>NP</b> non-preferred path  <b>NS</b> path attribute not specified</p> <p>For devices residing in control units that do not support path attributes, only the CHPID is displayed.</p>
DCM	<p>If the channel path is under control of Dynamic Channel Path Management (DCM) , this is indicated by a Y in this column. The activities of all DCM channels belonging to the same LCU will be summarized in a separate line.</p>
CTL Units	<p>The hexadecimal identifier of each physical control unit associated with an online channel path in the LCU group.</p>
DCM Group	<p>The values in columns MN MX DEF report the minimum and maximum number of DCM managed channels for one LCU (in this interval) as well as the installation-specified definition for this LCU.</p> <p>The line with these values is available only for LCUs with DCM managed channels. It contains in addition the accumulated values of the I/O activity rate, the director port contention, and the control unit contention of all DCM managed channels. These values may include also measurements of managed channels which were partially online.</p>
LCU	<p>The hexadecimal number that identifies the logical control unit (LCU).</p> <p>An LCU is the set of devices attached to the same physical control unit or a group of physical control units with one or more devices in common. Each physical control unit and each device can belong to only one LCU. They cannot be shared between LCUs.</p> <p>For each LCU, a summary line is reported in addition.</p>
Cont Rate	<p>The rate per second at which the SAP places delayed I/O requests on the CU-HDR for this LCU. This is done when all paths to the subchannel are busy and at least one path to the control unit is busy. For devices with only one path, or for devices where multiple paths exist and the busy condition is immediately resolved, the IOP does not count the condition.</p> $\text{Cont Rate} = \frac{\# \text{ Enqueued Requests}}{\text{Range Time}}$
Del Q Lngth	<p>The average number of delayed requests on the control unit header (CU-HDR). Each time a request is enqueued from the CU-HDR, RMF counts the number of requests remaining on the queue and adds that number to the accumulator. At the end of the interval, RMF divides the total number of accumulated queued requests by the number of times a request was enqueued.</p> $\text{Del Q Lngth} = \frac{\text{Accumulated Queue Length} - \# \text{ Enqueued Requests}}{\# \text{ Enqueued Requests}}$
AVG CSS	<p>The average number of milliseconds of delay that an I/O request encountered after the acceptance of the start or resume function at the subchannel for the LCU, until the channel subsystem's first attempt to initiate the operation.</p> $\text{AVG CSS} = \frac{\text{Channel Subsystem Time}}{\# \text{ I/O Operations Accepted}}$
CHPID Taken	<p>The rate at which I/O requests to devices of this LCU are satisfied by each CHPID during the interval. By reviewing the rate at which each channel path of the LCU satisfies I/O requests, you can see how evenly the work requests are distributed among the available paths and how effectively those paths are arranged for the LCU.</p> $\text{CHPID Taken} = \frac{\# \text{ I/O Operations Accepted on that Path}}{\text{Range Time}}$

## Mon III - IOQUEUE

Table 35. Fields in the IOQUEUE Report (continued)

Field Heading	Meaning
% DP Busy	<p>This field indicates director port contention. It is the number of times an I/O request was deferred because the director port was busy during the measurement interval.</p> $\% \text{ DP Busy} = \frac{\text{DPB}}{\text{DPB} + \text{CUB} + \text{SUC}} * 100$ <p><b>DPB</b> Number of deferred I/O requests due to director port busy  <b>CUB</b> Number of deferred I/O requests due to control unit busy  <b>SUC</b> Number of successful I/O requests on that path</p>
% CU Busy	<p>This field shows the relationship for each channel path of the LCU, between requests deferred due to control unit busy and total successful requests serviced by that path. Each CHPID of the LCU measures the distribution of control unit contention.</p> $\% \text{ CU Busy} = \frac{\text{CUB}}{\text{DPB} + \text{CUB} + \text{SUC}} * 100$ <p><b>DPB</b> Number of deferred I/O requests due to director port busy  <b>CUB</b> Number of deferred I/O requests due to control unit busy  <b>SUC</b> Number of successful I/O requests on that path</p>
AVG CUB	<p>The average number of milliseconds of delay that an I/O request encountered for the channel path because the control unit was busy.</p> $\text{AVG CUB} = \frac{\text{Control Unit Busy Time}}{\# \text{ I/O Operations Accepted on that Path}}$
AVG CMR	<p>The average number of milliseconds of delay that a successfully initiated start or resume function needs until the first command is indicated as accepted by the device. It allows to distinguish between real H/W errors versus workload spikes (contention in the fabric and at the destination port).</p> $\text{AVG CMR} = \frac{\text{Initial Command Response Time}}{\# \text{ I/O Operations Accepted on that Path}}$

## JES - Job Entry Subsystem Delays Report

The JES Delays report allows you to investigate situations where executing jobs are delayed when requesting service from JES. RMF lists all jobs delayed during the report interval in descending delay percentages.

### How to request this report

To request the JES report, select **4** on the Primary menu, and then select **2** on the Subsystem Report menu (shown in Figure 9 on page 27), or enter the following command:

```
JES [job_class,service_class]
```

## Contents of the report

```

RMF V2R2  JES Delays                               Line 1 of 3
Command ==>                                       Scroll ==> HALF
Samples: 120   System: SYSF  Date: 09/28/16  Time: 15.54.00  Range: 120  Sec
Jobname      DLY ----- Main Delay Reason -----
             %   % F-Code Explanation
AOR1         2   2  12  Waiting for job termination.
BXAB4        75  75 255  Waiting for SPOOL space.
    
```

Figure 62. JES Delays report

The graphic form of this report shows the percentage of each user's time spent waiting for JES services.

### Field descriptions

Table 36. Fields in the JES Report

Field Heading	Meaning
Jobname	Name of the job delayed when requesting service from JES. The JES Delays report does not summarize data by job groups; all jobs within a job group are reported individually.
DLY %	<p>Delay the waiting job is experiencing because of JES during the report interval. This value is calculated as follows:</p> $\text{DLY \%} = \frac{\text{\# Delay Samples}}{\text{\# Samples}} * 100$ <p><b>Delay samples</b> The number of samples when the job was delayed for JES. RMF calculates this number by incrementing its counter once for each sample when one or more units of work in the address space had JES delay.</p> <p>For JES2 delay, RMF scans all subsystem job blocks (SJBs).</p> <p>See the description of the <i>Main Delay Reason</i> field for valid JES function codes. <b>Note:</b> This DLY% value is also found in the JES field on the Job Delay report.</p>

Table 36. Fields in the JES Report (continued)

Field Heading	Meaning																																																																		
Main Delay Reason	<p>The subsystem function code that indicates the main reason for the delay. RMF reports the function code with the highest count as the main delay reason.</p> <p>The % column indicates how much of the JES delay of the job is caused by the reported subsystem function. This value is calculated as follows:</p> $\% = \frac{\text{\# Delay Samples}}{\text{\# Samples}} * 100$ <p><b>Delay samples</b> The number of samples when the job was delayed for JES for a specific subsystem function. For primary source fields used in this calculation see the DLY % field in this report description.</p> <p><b>Note:</b> The Main Delay Reason % values add up to the DLY % value of the job if there is no overlap in the delay states and there are no more than two function codes responsible for the delay.</p> <p>The JES F-codes (function codes in decimal) and their explanations are:</p> <table border="0"> <thead> <tr> <th data-bbox="492 726 573 747">F-Code</th> <th data-bbox="586 726 703 747">Explanation</th> </tr> </thead> <tbody> <tr><td>1</td><td>Processing TSO OUTPUT command request</td></tr> <tr><td>2</td><td>Waiting for JES to cancel a job</td></tr> <tr><td>3</td><td>Waiting for job status information</td></tr> <tr><td>9</td><td>Waiting for WTO/WTOR request</td></tr> <tr><td>12</td><td>Waiting for job termination</td></tr> <tr><td>13</td><td>Waiting for JES to restart a job</td></tr> <tr><td>23</td><td>Waiting for dynamic alloc via SSOBDYCD</td></tr> <tr><td>23</td><td>Waiting for SETUP request</td></tr> <tr><td>26</td><td>Waiting for change DD name</td></tr> <tr><td>27</td><td>Waiting for change ENQ use attribute</td></tr> <tr><td>75</td><td>Processing notify user</td></tr> <tr><td>132</td><td>Waiting for JDS access</td></tr> <tr><td>138</td><td>Validating SYSOUT destination</td></tr> <tr><td>138</td><td>Waiting for JES to cancel a job</td></tr> <tr><td>138</td><td>Waiting for job status information</td></tr> <tr><td>141</td><td>Waiting for ENDREQ</td></tr> <tr><td>144</td><td>Processing TSO OUTPUT command request</td></tr> <tr><td>149</td><td>Dynamically allocating data set to JES3</td></tr> <tr><td>151</td><td>Changing JES3 DD via dynamic allocation</td></tr> <tr><td>152</td><td>Waiting for FSS request</td></tr> <tr><td>153</td><td>Waiting for CI driver</td></tr> <tr><td>157</td><td>Waiting for SYSOUT API to process request</td></tr> <tr><td>158</td><td>Processing SSI Extended Status</td></tr> <tr><td>162</td><td>Waiting for TCPIP NJE global services</td></tr> <tr><td>163</td><td>Waiting for job class information</td></tr> <tr><td>164</td><td>Waiting for initiator information</td></tr> <tr><td>165</td><td>Waiting for NJE node information</td></tr> <tr><td>166</td><td>Waiting for spool partition information</td></tr> <tr><td>167</td><td>Waiting for JESPLEX information</td></tr> <tr><td>168</td><td>Waiting for SJF services</td></tr> <tr><td>169</td><td>Waiting for JES device information</td></tr> <tr><td>255</td><td>Waiting for SPOOL space</td></tr> </tbody> </table>	F-Code	Explanation	1	Processing TSO OUTPUT command request	2	Waiting for JES to cancel a job	3	Waiting for job status information	9	Waiting for WTO/WTOR request	12	Waiting for job termination	13	Waiting for JES to restart a job	23	Waiting for dynamic alloc via SSOBDYCD	23	Waiting for SETUP request	26	Waiting for change DD name	27	Waiting for change ENQ use attribute	75	Processing notify user	132	Waiting for JDS access	138	Validating SYSOUT destination	138	Waiting for JES to cancel a job	138	Waiting for job status information	141	Waiting for ENDREQ	144	Processing TSO OUTPUT command request	149	Dynamically allocating data set to JES3	151	Changing JES3 DD via dynamic allocation	152	Waiting for FSS request	153	Waiting for CI driver	157	Waiting for SYSOUT API to process request	158	Processing SSI Extended Status	162	Waiting for TCPIP NJE global services	163	Waiting for job class information	164	Waiting for initiator information	165	Waiting for NJE node information	166	Waiting for spool partition information	167	Waiting for JESPLEX information	168	Waiting for SJF services	169	Waiting for JES device information	255	Waiting for SPOOL space
F-Code	Explanation																																																																		
1	Processing TSO OUTPUT command request																																																																		
2	Waiting for JES to cancel a job																																																																		
3	Waiting for job status information																																																																		
9	Waiting for WTO/WTOR request																																																																		
12	Waiting for job termination																																																																		
13	Waiting for JES to restart a job																																																																		
23	Waiting for dynamic alloc via SSOBDYCD																																																																		
23	Waiting for SETUP request																																																																		
26	Waiting for change DD name																																																																		
27	Waiting for change ENQ use attribute																																																																		
75	Processing notify user																																																																		
132	Waiting for JDS access																																																																		
138	Validating SYSOUT destination																																																																		
138	Waiting for JES to cancel a job																																																																		
138	Waiting for job status information																																																																		
141	Waiting for ENDREQ																																																																		
144	Processing TSO OUTPUT command request																																																																		
149	Dynamically allocating data set to JES3																																																																		
151	Changing JES3 DD via dynamic allocation																																																																		
152	Waiting for FSS request																																																																		
153	Waiting for CI driver																																																																		
157	Waiting for SYSOUT API to process request																																																																		
158	Processing SSI Extended Status																																																																		
162	Waiting for TCPIP NJE global services																																																																		
163	Waiting for job class information																																																																		
164	Waiting for initiator information																																																																		
165	Waiting for NJE node information																																																																		
166	Waiting for spool partition information																																																																		
167	Waiting for JESPLEX information																																																																		
168	Waiting for SJF services																																																																		
169	Waiting for JES device information																																																																		
255	Waiting for SPOOL space																																																																		

### Report options

The JES Report Options panel is similar to the Device Report Options panel. See Figure 39 on page 72 for an example. Select YES for Jobs on the Report Options panel to display the Job Selection/Exclusion panel. See Figure 37 on page 70 for an example.

---

## JOB - Job Delay Report

The Job Delay report describes the reason why a specific job is delayed and provides possible causes leading to the delay.

### How to request this report

To request the Job Delay report, select **2** from the Primary Menu and then select **5** from the Job Report Selection Menu (shown in Figure 7 on page 26), or enter the following command:

```
JOB jobname
```

### Contents of the report

The Job Delay report is available in the following delay variations:

- Device delay
- Enqueue delay
- HSM address space delay
- JES address space delay
- Operator message delay
- Operator tape mount delay
- Processor delay
- Storage delay
- Quiesce delay
- XCF address space delay

The following sections explain each Job Delay report variation. The Job Delay report is divided into three sections. Each variation has the same fields in the top and bottom sections. The middle section varies depending on the type of delay being reported. Only the middle section will be described for each variation.

Delay can be either primary delay or requested delay. Primary delay shows information about the type of delay contributing most to the overall delay of a job. Requested delay shows information about a type of delay that you select. Use the Job Report Selection Menu, or cursor-sensitive control to select the type of delay to be reported. See “The Job Report Selection Menu” on page 25 for more information about requesting delay types.

### Displaying the Job Delays report

You can display the Job Delays report in several ways:

- Place the cursor on a jobname in any RMF report and press ENTER.
- Place the cursor on a % Delayed for field in a delay report or a Job Delay report and press ENTER.
- Issue the JOB <jobname> command on any command line.
- Use the Job Report Selection Menu to specify a particular job and the type of delay to be reported.

The highlighted line in the bottom section of the Job Delays report identifies the job which is analyzed in the top and middle sections.

Since more than one job with the same jobname may be in the system during the report interval, the address space identifier (ASID) is displayed instead of the jobname. To display a different job with the same name, place the cursor on the appropriate ASID field in the bottom section and press ENTER.

## The top and bottom parts of the Job Delays report

All variations of a Job Delays report have the same top part:

```

RMF V2R2 Job Delays
Command ==>>>
Line 1 of 1
Scroll ==>> HALF

Samples: 100 System: MVS1 Date: 09/28/16 Time: 10.03.20 Range: 100 Sec

Job: CHARLESR Primary delay: Excessive disconnect time on volume BPXLK1.

Probable causes: 1) Sequential access of data with short blocksize.
                  2) Active datasets spaced widely across volume.
                  3) Overloaded channel paths causing reconnect delays.
    
```

Figure 63. Top Part of Job Delays report

### Field descriptions

Table 37. Fields in the Job Delays report

Field Heading	Meaning
Job:	Name of delayed job.
Primary Delay: or Requested Delay:	A more detailed description of the reason stated in the Primary Reason field of the highlighted line in the bottom section of this report. Requested Delay appears instead of Primary Delay if you select device delays as the type of delay to be reported.
Probable causes:	Probable causes of the delay. These causes are selected according to the primary delay reason and by analysis of the data in the middle section of the report. Use the HELP key (PF1) for additional information about these causes and how they are determined.

```

----- Job Performance Summary -----
Service      WFL -Using%- DLY IDL UKN ----- % Delayed for ---- Primary
CX ASID Class P Cr % PRC DEV % % % PRC DEV STR SUB OPR ENQ Reason
B 0649 NRPRIME 1 S 29 7 13 39 0 16 15 33 0 0 0 0 BPXLK1
    
```

Figure 64. Bottom Part of Job Delay report

The bottom part of the Job Delay report has similar fields as the Delay report. Similar fields also appear on the Delay report. Table 17 on page 64 describes these fields. In the Job Delay report, the address space identifier (ASID) is displayed instead of the jobname since more than one job with the same jobname may have been in the system during the report interval. The Using% field shows the percentage of time the job was using a processor (PRC) and the percentage of time the job was using a device (DEV).

### Device Delay variation

This variation of the Job Delay report shows a job delayed by a device.

#### How to request this variation

To request the Device Delay variation, select 2 from the Primary Menu, and then select 1 on the Job Report Selection Menu (shown in Figure 7 on page 26), or enter the following command:

```
DEVJ jobname
```



----- Volume BPXLK1 Device Data -----						
Number:	0409	Active:	62%	Pending:	3%	Average Users
Device:	3390A	Connect:	17%	Delay DB:	1%	Delayed
Shared:	Yes	Disconnect:	42%	Delay CM:	0%	0.7
PAV:	1.6H					

Figure 65. Device Delay variation of the Job Delay report

The fields in the middle section of this report provide information about the device delaying the job.

### Field descriptions

Table 38. Fields in the Device Delay variation of the Job Delay report

Field Heading	Meaning
Number:	The device number where the volume is mounted.
Device:	The device type.
Shared:	Indicates if the device is shared.
PAV	<p>The number of parallel access volumes (base and alias) which were available at the end of the report interval. If the number has changed during the report interval, it is followed by an '*'. If the device is a HyperPAV base device, the number is followed by an 'H'. The value is the average number of HyperPAV volumes (base and alias) for that range.</p> $\text{Average \# of HPAV devices} = \frac{\text{Accumulated \# of HPAV devices}}{\text{Number of Samples}}$ <p>This field appears only for parallel access volumes.</p>
Active:	The percentage of time the device was active during the report interval. Dashes in this field indicate hardware data is not available. See Table 22 on page 76 for the calculation of this value.
Connect:	<p>The percentage of time the device was connected to a channel path. Dashes in this field indicate hardware data is not available. See Table 22 on page 76 for the calculation of this value. If the following condition exists on the DEVR report, this field will be highlighted to warn you about an excessive condition:</p> <p>CON % &gt; 40</p>
Disconnect:	<p>The percentage of time the device has an active channel program and is disconnected (not transferring data). Dashes in this field indicate hardware data is not available. See Table 22 on page 76 for the calculation of this value. If the following condition exists on the DEVR report, this field will be highlighted to warn you about an excessive condition:</p> <p>DSC % &gt; 40 and CON % &lt; 30</p>
Pending:	<p>The percentage of time all I/O requests wait before a path is available. Dashes in this field indicate hardware data is not available. See Table 22 on page 76 for the calculation of this value. If one of the following conditions exist on the DEVR report, this field will be highlighted to warn you about an excessive condition:</p> <p>PND % &gt; 40 and CON % = 0 PND % &gt; 40 and device is shared PND % &gt; 30 and device is not shared.</p>
Delay DB%:	<p>Device busy delay, which is the percentage of time during the report interval when the channel subsystem measured an I/O request delay because the device was busy. Device busy might mean that the volume is in use by another system, the device is reserved by another system, a head of string busy condition caused the contention, or some combination of these conditions has occurred.</p> $\text{Delay DB\%} = \frac{\text{Accumulated DB Delay Time}}{\text{Range Time}} * 100$

Table 38. Fields in the Device Delay variation of the Job Delay report (continued)

Field Heading	Meaning
Delay CM%:	Command response time delay, which is the percentage of time during the report interval, when the first command of an I/O instruction of the channel program is sent to the device, until the device indicates it has accepted the command.  $\text{Delay CM\%} = \frac{\text{Accumulated Command Response Delay Time}}{\text{Range Time}} * 100$
Average Users Delayed	The average number of users delayed by this device.  $\text{Average Users Delayed} = \frac{\sum \text{User Delay Counts}}{\# \text{ Valid Samples}} * 100$

## Enqueue Delay variation

This Job Delay report variation shows a job delayed by a serially reusable resource.

### How to request this variation

To request the Enqueue Delay variation, select **2** from the Primary Menu, and then select **2** from the Job Report Selection Menu (shown in Figure 7 on page 26), or enter the following command:

ENQJ jobname

----- Jobs Holding Resource -----					
Job:	AMOLLOY	Job:	SCHMATE	Job:	DRAGON
Holding:	100%	Holding:	100%	Holding:	100%
Status:	Shared	Status:	Shared	Status:	Shared

Figure 66. Enqueue Delay variation of the Job Delay report

The fields in the middle section of this report contain information about the three main holders of the resource for which this job is delayed.

### Field descriptions

Table 39. Fields in the Enqueue Delay variation of the Job Delay report

Field Heading	Meaning
Job:	Name of a job holding the resource. Up to three jobs can be displayed in this section.
Holding:	Indicates how much a specific job is contributing to the holding of the resource. See Table 32 on page 99 for the calculation of this value.
Status:	Indicates whether the job has exclusive or shared use of the resource.
System:	The name of the system the holding job is running on. This field appears only if the holding job is running on a different system.
Server:	The name of the catalog space which does the enqueue. This field appears only when the enqueue was done by a catalog space and is running on the same system.

## HSM and JES variations

The HSM and JES variations of the Job Delay report have the same format. They show a job delayed by either the HSM or JES address space.

## How to request these variations

Select 2 from the Primary Menu, and then select 3 for HSM or 4 for JES from the Job Report Selection Menu (shown in Figure 7 on page 26) or enter one of the following commands:

HSMJ jobname  
JESJ jobname

----- HSM Performance Summary -----			
Job: DFHSM	Workflow: 50%	Primary delay category:	OPER
	Using: 89%	Primary delay reason:	Mount
	Delay: 72%	Primary delay percent:	66%

Figure 67. HSM Delay variation of Job Delay report

## Field descriptions

The fields in the middle section of this report contain information about the HSM or JES address space.

Table 40. Fields in the HSM/JES delay variation of the Job Delay report

Field Heading	Meaning
Job:	Name of the JES/HSM address space.
Workflow:	The workflow percentage of the JES/HSM address space. See "Address space workflow (%)" on page 12 for the calculation.
Using:	The using percentage of the JES/HSM address space. See "Address space using (%)" on page 14 for the calculation.
Delay:	The delay percentage of the JES/HSM address space. See Table 17 on page 64 for the calculation.
Primary delay category:	The category of delay contributing most to the overall delay. See the % Delayed for field in Table 17 on page 64 for details.
Primary delay reason:	The contents of this field depend on the primary delay category. See the Primary Reason field in Table 17 on page 64 for an explanation.
Primary delay percent:	The percentage of delay for the primary delay category.

## Operator Message and Mount Delay variations

The message and mount delay variations of the Job Delay report have the same format. They show a job delayed by either an operator message request or mount request.

## How to request these variations

Select 2 from the Primary Menu, and then select 6 for delay caused by volume mount or select 7 for delay caused by operator reply from the Job Report Selection Menu (shown in Figure 7 on page 26) or enter one of the following commands:

MSGJ jobname  
MNTJ jobname

```

Job: MISTYDFS      Primary delay: Awaiting reply to operator request 53.

----- Job Performance Summary -----
      Service      WFL -Using%- DLY IDL UKN ---- % Delayed for ---- Primary
CX ASID Class    P Cr %  PRC DEV %  %  %  PRC DEV STR SUB OPR ENQ Reason
B 0167 NRPRIME  1   0  0  0 100  0  0  0  0  0  0  0 100  0 Message
    
```

Figure 68. Operator Message Delay variation of the Job Delay report

### Field descriptions

Table 41. Fields in the Operator Message and Mount Delay variations of the Job Delay report

Field Heading	Meaning
Job:	Name of delayed job.
Primary Delay: or Requested Delay:	A more detailed description of the reason stated in the Primary Reason field of the highlighted line in the bottom section of this report. Requested Delay appears instead of Primary Delay if you select operator delays as the type of delay to be reported.

### Processor Delay variation

This Job Delay report variation shows a job delayed by a processor.

#### How to request this variation

To request the Processor Delay variation, select **2** from the Primary Menu, and then select **8** from the Job Report Selection Menu (shown in Figure 7 on page 26) or enter the following command using the format:

PRO CJ jobname

```

----- Jobs Holding the Processor -----
Job:      D10PAV1      Job:      SAYLEKR      Job:      DAVEP
Holding:   18%        Holding:   15%        Holding:   12%
PROC Using: 25%        PROC Using: 21%        PROC Using: 25%
DEV Using:  0%         DEV Using:  5%         DEV Using:  9%
    
```

Figure 69. Processor Delay variation of the Job Delay report

The middle fields of this report provide information about the three main jobs holding the processor and causing delay.

### Field descriptions

Table 42. Fields in the Processor Delay variation of the Job Delay report

Field Heading	Meaning
Job:	Name of job holding the processor. Up to three jobs can be displayed in this section.
Holding:	The percentage of time the holding job used the processor while the delayed job was waiting for the processor.
PROC Using:	The percentage of time the holding job spent using the processor. See the USG % field in Table 54 on page 137 for the calculation of this value.
DEV Using:	The percentage of time the holding job spent using a DASD, tape, or MSC volume. See the USG % field in Table 19 on page 71 for the calculation of this value.

If the sum of the PROC DLY% (not shown on this report) and PROC Using fields of the holding job is 100, the PROC and DEV Using fields in the middle section will be highlighted. This indicates the main job holding the processor may be looping. In this case, "Job may be looping" is listed as a probable cause.

## Quiesce variation

### How to request this variation

To request the Quiesce Delay variation, select **2** from the Primary Menu, and then select **9** from the Job Report Selection Menu (shown in Figure 7 on page 26) or enter the following command using the format:

QSCJ jobname

This report shows a job delayed because it was quiesced by the operator.

----- Job Performance Summary -----																			
CX	ASID	Class	P	Cr	WFL %	-Using% PRC	DEV	DLY %	IDL %	UKN %	% Delayed for				----- Primary Reason				
B	0167	NRPRIME	1		0	0	0	100	0	0	0	PRC	DEV	STR	SUB	OPR	ENQ	Reason	
																		0	Quiesce

Figure 70. Quiesce Delay variation of the Job Delay report

### Field descriptions

Table 43. Fields in the Quiesce Delay variation of the Job Delay report

Field Heading	Meaning
Primary Reason QUIESCE	<p>The address space has been quiesced by the operator using the RESET command. A quiesced address space can show unexpected data:</p> <ul style="list-style-type: none"> <li>• A swappable address space will be swapped out, thus it can be OUTR and show storage delays.</li> <li>• A non-swappable address space will get lowest priority, thus it can show CPU delay, paging delay, or other delays, and even some USG % from time to time depending on the load on the system.</li> </ul>

## Storage Delay variation

This variation of the Job Delay report shows a job delayed by contention of storage.

### How to request this variation

To request the Storage Delay variation, select **2** from the Primary Menu, and then select **10** from the Job Report Selection Menu (shown in Figure 7 on page 26) or enter the following command using the format:

STORJ jobname

----- Job Storage Usage Data -----					
Average Frames:	294	Working set:	341	Fixed Frames:	38
Active Frames:	249	Aux Slots:	2928	DIV Frames:	0
Idle Frames:	45	Page In Rate:	3.3	ES Move Rate:	12.4

Figure 71. Storage Delay variation of Job Delay report

The middle fields of this report provide information about the storage usage of the delayed job.

### Field descriptions

Table 44. Fields in the Storage Delays variation of the Job Delay report

Field Heading	Meaning
Average Frames	The sum of active and idle frames.
Active Frames	The average number of frames held by the job while it was active. See Table 64 on page 155 for the calculation of this value.
Idle Frames	The average number of frames held by the job when it was idle. See Table 64 on page 155 for the calculation of this value.
Working Set	The average amount of storage a user occupied while in storage. See Table 64 on page 155 for the calculation of this value.
Aux Slots	The average number of auxiliary slots for each address space.
Page In Rate	<p>The rate at which pages are being read into central storage.</p> $\text{Page In Rate} = \frac{\sum \text{all Page-in Counts for Group}}{\text{Resident Time}}$ <p>The resident time is the total time the address space was swapped in.</p> <p>The page-in rate includes the shared storage page-ins.</p>
Fixed Frames	The average number of fixed frames the job was using during the range period including frames both above and below the 16 megabyte line. See Table 70 on page 162 for the calculation of this value.
DIV Frames	The number of central storage frames used by DIV. See Table 70 on page 162 for the calculation of this value.
ES Move Rate	The rate of pages moved from expanded storage. This includes both single and blocked pages; but does not include hiperspace or VIO pages.

### XCF variation

The cross-system coupling facility (XCF) variation of the Job Delay report shows a job delayed by XCF.

#### How to request this variation

To request the XCF variation, select **2** from the Primary Menu, and then select **11** from the Job Report Selection Menu (shown in Figure 7 on page 26) or enter the following command:

XCFJ jobname

----- XCF Path Summary -----				
	1st Path	2nd Path	3rd Path	4th Path
Delay:	4%	3%		
Dev Number:	0E80	0CA0		

Figure 72. XCF Delay variation of Job Delay report

### Field descriptions

Table 45. Fields in the XCF Delay variation of the Job Delay report

Field Heading	Meaning
Delay:	The delay percentage of the XCF address space. See Table 17 on page 64 for the calculation.
Dev Number:	The path number corresponding to the delay percentage.

### Monitor III Utility fields

You can use the Monitor III Utility to customize the Job Delay report. In addition to the delays previously described, you can use the Utility to have the following delay percentages shown.

Table 46. Additional Fields in the Job Delay report

Field Heading	Meaning
Using percentage	The percentage of time a job was using the processor or a device.
JES delay percentage	The percentage of time a job was delayed when requesting service from JES.
HSM delay percentage	The percentage of time a job was delayed when requesting service from HSM.
XCF delay percentage	The percentage of time a job was delayed when requesting service from XCF.
Operator mount delay percentage	The percentage of time a job was delayed by an operator mount request.
Operator message delay percentage	The percentage of time a job was delayed by an operator message request.
Operator quiesce delay percentage	The percentage of time a job was delayed because the operator quiesced the address space.
WLM resource capping delay percentage	The percentage of time a job was delayed because <ul style="list-style-type: none"> <li>• it has used up its CPU service as specified in the WLM policy for the resource group to which the job belongs</li> <li>• or because the work for which the job is running is overachieving its goal. So this work may be capped in order to divert its resources to run discretionary work (see also section 'Using Discretionary Goals' in <i>z/OS MVS Planning: Workload Management</i>).</li> </ul>

### Report options

Each variation of the Job Delays report uses the same Report Options panel. Use this panel to view and select an available jobname.

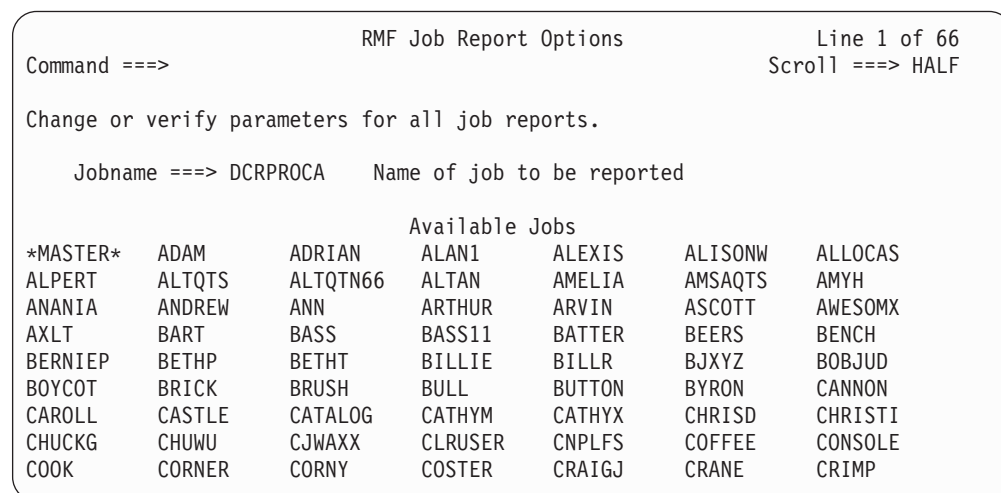


Figure 73. Job Report Options Panel

RMF saves the Jobname you enter across sessions.

#### Jobname

The name of the job for which you want data reported in your job delay report.

#### Available Jobs

The list of jobs that were active during the report interval.

If the job you want is not listed, it had no activity during the current report interval. If you specify the job, it will appear on your Job Delay reports when it is available.

---

## LOCKSP - Spin Lock Report

Through locking, the system serializes the use of system resources by authorized routines and, in a Parallel Sysplex, by processors. Lock holders can impede other work units that need the same lock and must wait until the lock holder releases the lock.

RMF reports about the various types of system resource locks in the **Spin Lock Report** described in this section and in the **Suspend Lock Report** described in "LOCKSU - Suspend Lock Report" on page 128.

If a spin lock is unavailable, the requesting processor continues testing the lock until the other processor releases it (spinning). As soon as the lock is released, the requesting processor can obtain the lock and thus can obtain control of the protected resource.

No symptoms for delays due to locks are visible except excessive spinning. Therefore, RMF periodically checks all types of system resource locks.

### How to request this report

To request the Spin Lock Report, select **1** from the Primary Menu and then select **12** on the Overview Report Selection Menu (shown in Figure 6 on page 25), or enter one of the following commands:

```
LOCKSP [HELD | SPIN | BOTH]
```

```
LSP [HELD | SPIN | BOTH]
```

### Contents of the report

The **Spin Lock Report** provides information about how often a spin lock is held and about jobs that are spinning because of a lock request. It consists of two sections:

- The upper part (Held section) displays information about spin locks which have been observed as held, either exclusively (EXCL) or shared (SHR).
- The lower part (Spin section) displays spin locks and address spaces which are spinning due to a request for this lock.



```

RMF V2R2 Spin Lock Report
Command ==>> Line 1 of 10
Scroll ==> PAGE

Samples: 120 System: CB88 Date: 09/28/16 Time: 09.21.00 Range: 120 Sec

Resource Type/ CPUID/ Held Spin
Jobname ASID Address % %
SRM EXCL 07 0.83
DISP EXCL 0E 0.41
SRM EXCL 09 0.41
DISP EXCL 11 0.41
DISP EXCL 07 0.41
DISP EXCL 04 0.41
VFIX EXCL 0C 0.41
VFIX EXCL 06 0.41
DISP EXCL 01 0.21
XCFQ SHR 06 0.21
IXLSHR SHR 01 0.21
IXLSHR SHR 0C 0.21
VFIX RMFGAT 0196 0147A128 0.05
IXSH S/E RMFGAT 0196 00FF14D4 0.05
    
```

Figure 74. Spin Lock Report

Table 47. Fields in the LSP Report

Field Heading	Meaning
Resource	The resource name or the address of the spin lock.
Type/Jobname	<b>Held</b> section: The type of the lock (exclusively or shared). <b>Spin</b> section: The jobname (address space), which is spinning due to the lock request.
CPUID/ASID	<b>Held</b> section: The ID of the logical CPU holding the lock. <b>Spin</b> section: The decimal address space identifier of the spinning job.
Address	The address of the instruction which obtained the lock.
Held %	The percentage of samples where the lock has been held.
Spin %	The percentage of samples where the requesting address space (ASID) has been found spinning due to the unavailable lock.

### Lock report options

On the *Lock Report Options* panel, you can specify options for the **Spin Lock Report** (LOCKSP) and the **Suspend Lock Report** (LOCKSU).

```

RMF Lock Report Options
Command ==>> Scroll ==> CSR

Change or verify parameters. To exit press END.

Changes will apply to the LOCKSP and the LOCKSU report.

Spin Lock ==> BOTH Information (HELD, SPIN or BOTH) in LOCKSP report
Lock Type ==> BOTH Lock type (GLOBAL, LOCAL or BOTH) in LOCKSU report
    
```

Figure 75. Lock Report Options

## LOCKSU - Suspend Lock Report

Through locking, the system serializes the use of system resources by authorized routines and, in a Parallel Sysplex, by processors. Lock holders can impede other work units that need the same lock and must wait until the lock holder releases the lock.

RMF reports about the various types of system resource locks in the **Suspend Lock Report** described in this section and in the **Spin Lock Report** described in "LOCKSP - Spin Lock Report" on page 126.

If a suspend lock is unavailable, the unit of work requesting the lock is delayed until the lock is available. Other work is dispatched on the requesting processor. All local locks are suspend locks.

No symptoms for delays due to suspend locks are visible. Therefore, this report provides information about the jobs that are holding a suspend lock, because the overall workflow can be impacted by contention situations for the same lock. Especially, if a work unit that is holding a lock is suspended for a longer period of time, other work units can be significantly delayed.

### How to request this report

To request the **Suspend Lock Report**, select **1** from the Primary Menu and then select **13** on the Overview Report Selection Menu (shown in Figure 6 on page 25), or enter one of the following commands:

LOCKSU [GLOBAL | LOCAL | BOTH]

LSU [GLOBAL | LOCAL | BOTH]

### Contents of the report

The report contains one segment for local suspend locks in the upper part and one for global suspend locks in the lower part of the report. Within the segments, the report lines are sorted by descending Held%. A separate work unit within the same address space can be identified by the value in column *Address*, which is the address of the instruction that obtained the lock.

RMF V2R2 Suspend Lock Report								Line 1 of 6
Command ==>								Scroll ==> PAGE
Samples: 120	System: CB88	Date: 09/28/16	Time: 09.21.00	Range: 120	Sec			
Resource	Type	Jobname	ASID	Address	Held %	Intr %	Disp %	Susp %
OS390R1	L	OS390R1	0045	07072C52	0.83	0.00	0.83	0.00
ZFS	L	ZFS	0234	012B2A66	0.83	0.00	0.83	0.00
GRS	L	GRS	0007	015EC438	0.83	0.00	0.83	0.00
ZFS	L	ZFS	0236	012B2A66	0.83	0.00	0.00	0.00
CMSSMFLK	G	ZFS	0062	00CE4652	0.83	0.00	0.00	0.00
CMSEDLK	G	OS390R1	0045	2A64085A	0.83	0.00	0.83	0.00
CMSLOCK	G	HZSPROC	0059	00D3EC70	0.83	0.00	0.00	0.00

Figure 76. Suspend Lock Report

Table 48. Fields in the LSU Report

Field Heading	Meaning
Resource	The resource name of the suspend lock. <ul style="list-style-type: none"> <li>• for local locks: the address space name where the local lock resides</li> <li>• for cross memory local (CML) locks: the primary address space name (which is different from the holder's job name)</li> <li>• for all types of cross memory services (CMS) locks: the lock word name.</li> </ul>
Type	The type of the suspend lock: <b>L</b> Local Suspend Lock <b>LX</b> Cross Memory Local (CML) Suspend Lock <b>G</b> Global CMS Suspend Lock.
Jobname	The name of the job/address space holding the lock.
ASID	The decimal address space identifier of the job holding the lock.
Address	The address of the instruction that obtained the lock. For local locks, the address of the instruction can be in the lock address space or the requestor address space.
Held %	The percentage of samples where the address space held the lock during the report interval.
Intr %	The percentage of samples where the address space was interrupted while holding the lock.
Disp %	The percentage of samples where the address space was dispatchable while holding the lock.
Susp %	The percentage of samples where the address space has been found suspended while another dispatchable unit was holding the lock.

**Note:** You can specify options for this report on the *Lock Report Options* panel described in “Lock report options” on page 127.

---

## OPD - OMVS Process Data Report

z/OS Unix address spaces can consist of several processes, which in turn might run one or more threads. Each process is typically associated with a UNIX command, consumes a certain amount of CPU, and also provides state information. UNIX System Services is the brand for UNIX on z/OS. In this context, it is referred to as open MVS or OMVS.

In addition to other reports that show OMVS address spaces with their jobname and using or delay information, the OPD report can be used for problem determination. It assists the performance analyst to find answers to the following questions:

- What are the delayed processes?
- What command is associated with them?
- What is the status of each of the processes?
- Which processes are high CPU consumers?

Address spaces under OMVS control are indicated by an additional letter **O** in the class column of the DELAY, the PROC, and the JOB report. The performance analyst can then use cursor sensitivity to navigate to this report or alternatively invoke it directly depending on the task he is trying to accomplish.

The report provides basic performance metrics on the first screen, while additional information specifically related to server processes can be shown by activating a pop-up panel.

## How to request this report

To request the OMVS Process Data report, select 1 from the Primary Menu and then select 7 on the Overview Report menu (shown in Figure 6 on page 25), or enter the following command:

OPD

## Contents of the report

RMF V2R2 OMVS Process Data										Line 1 of 24
Command ==>>					Scroll ==> HALF					
Samples: 18	System: SYS4	Date: 09/28/16	Time: 15.50.41	Range: 19	Sec					
Kernel Procedure: OMVS		Kernel ASID: 0014		Option: PID		ALL				
BPXPRM: OMVS=(71,04)										
-----										
Jobname	User	ASID	PID	PPID	LW	State	Appl%	Total	Server	
BPXOINIT	OMVSKERN	0030	1	0		MF	0.0	0.234	FILE	
INETD8	OMVSKERN	0047	5	1		1FI	0.0	0.052	N/A	
MVSNFSC	MVSNFS	5001	7	1		1A	0.0	0.229	N/A	
MVSNFSC	MVSNFS	5001	8	1		1A	0.0	0.229	N/A	
MVSNFSC	MVSNFS	5001	9	1		1A	0.0	0.229	N/A	
MVSNFSC	MVSNFS	5001	10	1		1A	0.0	0.229	N/A	
MVSNFSC	MVSNFS	5001	11	1		1A	0.0	0.229	N/A	
MVSNFSC	MVSNFS	5001	12	1		1A	0.0	0.229	N/A	
MVSNFSC	MVSNFS	5001	13	1		1A	0.0	0.229	N/A	
MVSNFSC	MVSNFS	5001	14	1		1A	0.0	0.229	N/A	
TCPIP	TCPIP	0044	15	1		MR	0.0	43.59	N/A	
TCPIP	TCPIP	0044	16	1		1R	0.0	43.59	N/A	
TCPIP	TCPIP	0044	17	1		1R	0.0	43.59	N/A	

Figure 77. OPD Report

There is no graphic version of this report available.

Table 49. Fields in the OPD Report

Field Heading	Meaning
Kernel Procedure	Name of the procedure used to start the OMVS kernel address space.
Kernel ASID	Decimal ID of the kernel address space.
BPXPRM	List of suffixes indicating the BPXPRM Parmlib member concatenation.
Option	Displays the current report option as specified on the Report Options panel.
Jobname	Jobname associated with the process.
User	User name associated with the process.
ASID	Decimal ID of the address space the process is associated with.
PID	Process ID.
PPID	Parent process ID.
LW	If the reported process is waiting for the process latch of another process, 'Y' is shown, otherwise blank.
State	Cumulated state information of the address space and process. You can place the cursor on any field (except Jobname and PPID) in a process line and press Enter - this will show you a pop-up panel with an explanation of the process state.
Appl%	Percentage of TCB and local/global SRB time consumed by the address space during the reporting range. <b>Note:</b> The calculated value is based on uniprocessor capacity and can exceed 100% on systems with more than one processor. To get the system utilization, this value has to be divided by the number of logical processors or cores.

Table 49. Fields in the OPD Report (continued)

Field Heading	Meaning
Total	Total computing time in seconds, consumed by the address space the process is running within. When only one process is running in the address space, this time represents the accumulated CPU time for that process. In case of multiple processes running in an address space, it is the sum of the CPU time used by all of the work running in that address space.
Server	If the process represents a server, one of the following is shown: <b>FILE</b> Network file server <b>LOCK</b> Network lock server <b>FEXP</b> Network file exporter <b>SFDS</b> Shared file server  For non-server processes, 'N/A' is shown.

The following pop-up panel shows an example of process details for a server process. For a non-server process, RMF displays 'N/A' in the fields below **Server Information**.

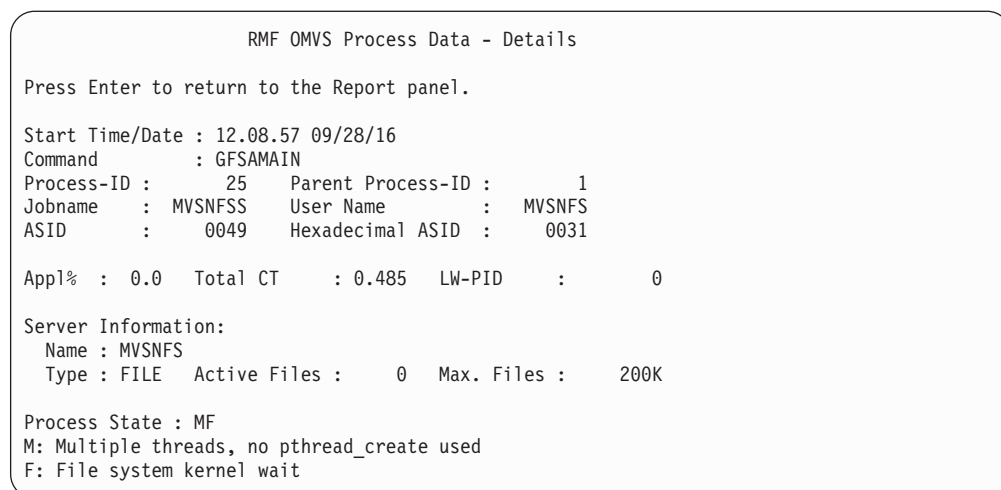


Figure 78. OPD Report - Details for Server Process

Table 50. Fields in the OPD Details Report

Field Heading	Meaning
Start Time/Date	Start time and date when the process has been started.
Command	The command that created the processes truncated to 40 characters.
Process-ID	Process ID.
Parent Process-ID	Parent process ID.
Jobname	Jobname associated with the process.
User Name	User name associated with the process.
ASID	Decimal ID of the address space the process is associated with.
Hexadecimal ASID	Hexadecimal identifier of the address space.
Appl%	Percentage of TCB and local/global SRB time consumed by the address space during the reporting range. <b>Note:</b> APPL% shows CPU utilization based on uniprocessor capacity. On systems with more than one processor this value has to be divided by the number of processors to get the system utilization.
Total CT	Total computing time in seconds, consumed by the address space the process is running in. When only one process is running in the address space, this time represents the accumulated CPU time for that process. In case of multiple processes running in an address space, it is the sum of the CPU time used by all of the work running in that address space.

Table 50. Fields in the OPD Details Report (continued)

Field Heading	Meaning																																												
LW-PID	Process ID of the process on whose latch the reported process is waiting for.																																												
Server Information. 'N/A' is shown next to each field if this is not a server process.																																													
Name	The name of the server process.																																												
Type	If the process represents a server, one of the following is shown: <b>FILE</b> Network file server <b>LOCK</b> Network lock server <b>FEXP</b> Network file exporter <b>SFDS</b> Shared file server																																												
Active Files	The number of active server file tokens.																																												
Max. Files	The maximum number of active server file tokens allowed.																																												
Process State	Cumulated state information of the address space and process. For each possible state a separate line is shown below the field. The following translation table is used:  <table border="0"> <thead> <tr> <th>State</th> <th>Meaning</th> </tr> </thead> <tbody> <tr><td>1</td><td>Single thread</td></tr> <tr><td>A</td><td>Message queue receive wait</td></tr> <tr><td>B</td><td>Message queue send wait</td></tr> <tr><td>C</td><td>Communication system kernel wait</td></tr> <tr><td>D</td><td>Semaphore operation wait</td></tr> <tr><td>E</td><td>Quiesce frozen</td></tr> <tr><td>F</td><td>File system kernel wait</td></tr> <tr><td>G</td><td>MVS pause wait</td></tr> <tr><td>H</td><td>Multiple threads, pthread_create used</td></tr> <tr><td>I</td><td>Swapped out</td></tr> <tr><td>K</td><td>Other kernel wait</td></tr> <tr><td>L</td><td>Cancelled, parent waits</td></tr> <tr><td>M</td><td>Multiple threads, no pthread_create used</td></tr> <tr><td>P</td><td>Ptrace kernel wait</td></tr> <tr><td>Q</td><td>Quiesce termination wait</td></tr> <tr><td>R</td><td>Running</td></tr> <tr><td>S</td><td>Sleeping</td></tr> <tr><td>T</td><td>Stopped</td></tr> <tr><td>W</td><td>Waiting for child</td></tr> <tr><td>X</td><td>Creating new process</td></tr> <tr><td>Z</td><td>Zombie. Cancelled, parent does not wait</td></tr> </tbody> </table>	State	Meaning	1	Single thread	A	Message queue receive wait	B	Message queue send wait	C	Communication system kernel wait	D	Semaphore operation wait	E	Quiesce frozen	F	File system kernel wait	G	MVS pause wait	H	Multiple threads, pthread_create used	I	Swapped out	K	Other kernel wait	L	Cancelled, parent waits	M	Multiple threads, no pthread_create used	P	Ptrace kernel wait	Q	Quiesce termination wait	R	Running	S	Sleeping	T	Stopped	W	Waiting for child	X	Creating new process	Z	Zombie. Cancelled, parent does not wait
State	Meaning																																												
1	Single thread																																												
A	Message queue receive wait																																												
B	Message queue send wait																																												
C	Communication system kernel wait																																												
D	Semaphore operation wait																																												
E	Quiesce frozen																																												
F	File system kernel wait																																												
G	MVS pause wait																																												
H	Multiple threads, pthread_create used																																												
I	Swapped out																																												
K	Other kernel wait																																												
L	Cancelled, parent waits																																												
M	Multiple threads, no pthread_create used																																												
P	Ptrace kernel wait																																												
Q	Quiesce termination wait																																												
R	Running																																												
S	Sleeping																																												
T	Stopped																																												
W	Waiting for child																																												
X	Creating new process																																												
Z	Zombie. Cancelled, parent does not wait																																												

### Report options

```

RMF OMVS Process Data Report Options
Command ==>                               Scroll ==> HALF

Change or verify parameters. To exit press END.
Select one of the following options:

1 1. Process ID ==> 1           ALL or a process ID
2. ASID      ==> 0100         ID of an address space in decimal or
                             hexadecimal (with preceding X) format
3. Jobname   ==>              Jobname associated with a process
4. User      ==>              User name associated with a process
    
```

Figure 79. OPD Report Options Panel

You can specify a process ID, an address space ID (in decimal or hexadecimal format), a jobname, or a user name to tailor the OPD report.

## PCIE - PCIE Activity Report

The PCIE Activity Report can be used to investigate performance problems that are related to PCI Express based functions (PCIE functions). This report provides these types of information:

- General PCIE activity metrics that are independent of the type of the exploited hardware feature and reflect the activity of the z/OS system on which RMF data collection took place. These metrics include data rates for the communication between z/OS programs and the PCIE functions (PCI LOAD, PCI STORE, PCI STORE BLOCK, REFRESH PCI TRANSLATIONS, and Read/Write Transfer data rates).
- Additional metrics that can be displayed on a pop-up panel for each Hardware Accelerator and zEnterprise Data Compression (zEDC) device. These metrics have single system scope and include device driver buffer statistics, common accelerator metrics (for example, total request execution time or the amount of transferred data), as well as compression specific metrics (for example, the amount of compressed data and the number and throughput of compression requests).

The PCIE Activity Report provides statistics and performance measurements on PCIE functions allocated by at least one z/OS address space. A PCIE function is captured by the report if one of the following feature activities has been detected:

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

### How to request this report

To request the PCIE Activity Report, select a **3** from the Primary Menu, then select **14** from the Resource Report Selection Menu (shown in Figure 8 on page 27) or enter the following command:

```
PCIE
```

### Contents of the report

RMF V2R2 PCIE Activity											Line 1 of 3	
Command ==>											Scroll ==> CSR	
Samples: 300		System: TA2		Date: 12/18/16		Time: 22.20.00		Range: 300		Sec		
ID	CHID	Type	Jobname	ASID	Status	Alloc Time%	PCI Load	Operations Store	Rate Block	-Xfer Refr	Rate Read	Rate Write
0022	037C	zEDC	FPGHWAM	0012	Alloc	100	0	0.008	0	0.533	0	0
0029	03BC	RoCE	TCPIP	0030	Alloc	100	0	0.108	0	1.234	0	0
0030	0400	ISM	TCPIP	0030	Alloc	100	0	0.123	0	2.345		2.4

Figure 80. PCIE Activity Report

Table 51. Fields in the PCIE Activity Report

Field Heading	Meaning
Function ID	Hexadecimal identifier of the PCIE Function for which performance data is reported.
Function CHID	Physical or virtual channel identifier for the PCIE function.

## Mon III - PCIE

Table 51. Fields in the PCIE Activity Report (continued)

Field Heading	Meaning
Function Type	Device type for the PCIE function which can be one of the following: <b>HWA</b> Hardware Accelerator <b>ISM</b> SMC-Direct over Internal Shared Memory <b>Oth</b> Unknown device type <b>RoCE</b> Remote Direct Memory Management <b>zEDC</b> zEnterprise Data Compression
Function Jobname	Name of the job who allocated the PCIE function.
Function ASID	Address space ID of the job who allocated the PCIE function.
Function Status	The PCIE function status at the end of this reporting interval: <b>Alloc</b> The function is allocated and in use <b>DPend</b> The function is in the process of de-allocation <b>Error</b> The function is in permanent error <b>DeAlloc</b> The function is de-allocated <b>Unknown</b> The function status is unknown
Alloc Time %	Percentage of the reporting interval for which the PCIE function was allocated or in the process of de-allocation.
PCI Operations Rate Load	Rate of PCI Load operations that were executed within this reporting interval.
PCI Operations Rate Store	Rate of PCI Store operations that were executed within this reporting interval.
PCI Operations Rate Block	Rate of PCI Store Block operations that were executed within this reporting interval.
PCI Operations Rate Refresh	Rate of Refresh PCI Translations operations that were executed within this reporting interval.
Xfer Rate Read	The number of megabytes per second that a RoCE device received on the external Ethernet interface. On zEC12 or zBC12, this field designates the number of megabytes per second that were transferred by DMA reads from all defined DMA address spaces to the PCIE function.
Xfer Rate Write	The number of megabytes per second transmitted on a RoCE or SMC-D device. On zEC12 or zBC12, this field designates the number of megabytes per second that were transferred by DMA writes from the PCIE function to all defined DMA address spaces.

### Monitor III Utility fields

You can use the Monitor III Utility to customize the PCIE Activity report to show the shown in Table 52.

Table 52. Monitor III Utility fields

Field Heading	Meaning
PCIEDWUP	The number of work units per second that were processed by this zEDC device. The field is valid only if the PCIE function is a zEDC device that is configured on post-zEC12 or post-zBC12 hardware.
PCIEDAUT	Utilization of this zEDC device. The field is valid only if the PCIE function is a zEDC device that is configured on post-zEC12 or post-zBC12 hardware.
PCIEDPKR	The number of packets received per second. The field is valid only if the PCIE function is a RoCE device that is configured on post-zEC12 or post-zBC12 hardware.
PCIEDPKT	Number of packets transmitted per second. The field is valid only if the PCIE function is a zEDC device that is configured on post-zEC12 or post-zBC12 hardware.
PCIENET1	Physical-network identifier (PNET ID) that identifies the first port of the RoCE device or virtual PCIE function
PCIENET2	Physical-network identifier (PNET ID) that identifies the second port of the RoCE device



## Cursor sensitive control

If the cursor is placed on one of the values in the Function ID, Function PCID, or Function Type columns, additional metrics can be displayed on a pop-up panel if the selected PCIE function is used as a Hardware Accelerator or zEDC Accelerator. Depending on the which performance metrics are available for the selected PCIE function, these metrics are displayed either on the RMF Hardware Accelerator And Compression Activity pop-up panel (see “Cursor sensitive control”) or RMF Hardware Accelerator Activity pop-up panel (see “Cursor sensitive control”).

RMF Hardware Accelerator And Compression Activity		
Press Enter to return to the Report panel.		
Function ID : 006C	Alloc Time % :	100
Allocated : 22.03.22	on 12/18/14	
Hardware Accelerator		
Time Busy % :	0.286	Transfer Rate : 4.87
Request		
Execution Time :	28.0	Std. Deviation: 8.07
Queue Time :	65.7	Std. Deviation: 140
Size :	47.6	
Buffer Pool		
Memory Size :	16	Utilization % : 0
	Compression	Decompression
Request Rate :	102	0.437
Throughput :	2.91	0.009
Ratio :	2.79	0.652

Figure 81. RMF Hardware Accelerator And Compression Activity panel

“Cursor sensitive control” provides the meanings of the fields that appear on the RMF Hardware Accelerator And Compression Activity and RMF Hardware Accelerator Activity pop-up panels.

Table 53. Fields on pop-up panels.

Field Heading	Meaning
Function ID	PCIE Function ID that represents the Hardware Accelerator.
Alloc Time %	Percentage of the reporting interval for which the PCIE function was allocated or de-allocate-pending.
Allocated	Date and time when the Hardware Accelerator was allocated.
Hardware Accelerator Time Busy %	The percentage of time that this partition kept the Hardware Accelerator busy.
Hardware Accelerator Transfer Rate	The number of megabytes per second that were transferred by DMA operations.
Request Execution Time	The average time in microseconds the Hardware Accelerator took to process a request.
Request Execution Time Std. Deviation	The standard deviation of the request execution time.

Table 53. Fields on pop-up panels. (continued)

Field Heading	Meaning
Request Queue Time	The average queue time in microseconds that was spent for a request.
Request Queue Time Std. Deviation	The standard deviation of the request queue time.
Request Size	The average number of kilobytes that were transferred per request.
Buffer Pool Memory Size	The total size of memory in megabytes that is allocated to the buffer pool.
Buffer Pool Utilization %	The average utilization percentage of the buffer pool that z/OS kept for in-use buffers.
Request Rate	The number of compression or decompression requests per second.
Throughput	The number of megabytes that were compressed or decompressed per second.
Ratio	The ratio between input and output bytes that were compressed or decompressed within this reporting interval.

---

## PROC - Processor Delays Report

The Processor Delays report (PROC) displays all jobs that were waiting for or using the processor during the report interval.

RMF reports the jobs by descending overall delay percentages. Because use of the processor by many jobs might contribute to the delay of another job, RMF reports up to three jobs in the Holding Job(s) field. The jobs in this field are those that were most often found using the processor while the job was delayed.

### How to request this report

To request the Processor Delay report, select 3 from the Primary Menu, then 1 from the Resource Report Selection Menu (shown in Figure 8 on page 27) or enter the following command:

```
PROC [job_class,service_class]
```

For example, to get a Processor Delays report for TSO service class TSOPRIME, enter:

```
PROC T, TSOPRIME
```

## Contents of the report

```

RMF V2R2 Processor Delays
Command ==>>>
Line 1 of 138
Scroll ==>> HALF

Samples: 60      System: MVS1  Date: 09/28/16  Time: 09.10.00  Range: 60 Sec

Jobname  Service  CPU  DLY  USG  EApp1  -----  Holding  Job(s)  -----
          CX  Class  Type  %   %   %      %  Name    %  Name    %  Name
WSWS7    0  OMVS   CP   11  46  59.4   9 *ENCLAVE  7 DBS3DIST  7 WSP1S2F
WSP1S2FS SO  WASCR  CP    4   4  42.5   2 DBS3DIST  2 WSWS7    2 VTAM44
          AAP    6   0  98.4   6 *ENCLAVE
WSP1S6FS SO  WASCR  CP    0   0   5.3    6 *ENCLAVE
          AAP    6   0   7.7    2 XCFAS    2 DBS3DIST  2 WSP1S2F
DBS3DBM1 S  DB2HIGH CP    2   6   0.8
WSP1S6F  SO  WASCR  CP    0   2   1.9
          AAP    2   2   0.7    2 *ENCLAVE
U078069  0  OMVS   CP    2   4   1.2    2 WSWS7    2 DBS3DIST  2 U078069
WSP1S4F  SO  WASCR  CP    0   0   0.1
          AAP    2   0   0.4    2 WSP1S6F
U078068  0  OMVS   CP    2   0   0.2    2 XCFAS    2 WSWS7    2 *ENCLAVE
DBS3DIST SO  DB2HIGH CP    0  78 111.0
          IIP    0   2  21.3
XCFAS    S  SYSTEM CP    0  28 24.1
TCP/IP   SO  SYSSTC CP    0  22 16.1
VTAM44   S  SYSSTC CP    0  19 14.5
WSP1S2F  SO  WASCR  CP    0  15 14.0
    
```

Figure 82. PROC - Processor Delay Report

The graphic form of this report shows each user's processor delay percentage and processor using percentage.

### Field descriptions

Table 54. Fields in the PROC Report

Field Heading	Meaning
Jobname	The name of a job. The processor delay report does not summarize data by job groups; all jobs within a job group are reported individually.
CX	Abbreviation for the job class as follows: <b>S</b> Started task <b>T</b> TSO <b>B</b> Batch <b>A</b> ASCH <b>O</b> OMVS  An <b>O</b> as second character indicates that the address space is using OMVS services.
Service Class	The name of the service class that a specified job has been running in.
CPU Type	The processor type: <b>CP</b> general purpose processor <b>AAP</b> Application Assist Processor (zAAP) <b>IIP</b> Integrated Information Processors (zIIP)

## Mon III - PROC

Table 54. Fields in the PROC Report (continued)

Field Heading	Meaning
DLY %	<p>Delay percentage that the waiting job (address space) is experiencing because of contention for the processor of the type indicated in column <i>CPU Type</i> during the report interval.</p> $\text{DLY \%} = \frac{\text{\# Delay Samples}}{\text{\# Samples}} * 100$ <p><b>Delay samples</b> The single state count of samples being delayed by the processor. RMF increments this count only once for each sample when one or more units of work (TCBs, SRBs, interrupted ready task or asynchronous exit) associated with the address space are delayed for the processor.</p> <p><b>Note:</b> This DLY% value is also found in the PROC field on the Job Delay report.</p>
USG%	<p>The percentage of time when the job is receiving service from the processor of the type indicated in column <i>CPU Type</i>.</p> $\text{USG \%} = \frac{\text{\# Using Samples}}{\text{\# Samples}} * 100$ <p><b>Using samples</b> The number of samples when the job was found using the processor. If the processor running Monitor III has other ready work to do (any ready SRB, interrupted ready task, asynchronous exit routine, or TCB is on the dispatching queue), then it looks for the first address space having a unit of work on the dispatching queue that is not already using another processor. Then the number of samples is incremented by one for the address space having the first dispatchable unit of work according to the dispatcher sequence search order. The processor running Monitor III is not counted as a processor in use if there is no other ready work to do.</p>
EApp1%	<p>Percentage of the processor time used by transactions that executed on the type of processor indicated in column <i>CPU Type</i>. This calculation is based on uniprocessor capacity which means that this value can exceed 100% in systems with more than one processor. To get the system utilization, this value has to be divided by the number of logical processors or cores.</p> <p><b>Note:</b> The processor times that are used to calculate this value is the sum of TCB time, global and local SRB time, preemptable or client SRB time, and enclave CPU time consumed within this address space.</p>
Holding Job(s)	<p>Up to three jobs that, by their use of the processor, contributed most to the delay of the job listed under Jobname.</p> <p><b>%</b> The percentage of delay caused by the named job to the job waiting to use the processor.</p> <p><b>Name</b> The name of a job contributing to the delay of the job waiting to use the processor.</p> <p>The name <b>*ENCLAVE</b> in this field means that one or more enclaves are active on the processor. The percentage shown for enclaves is the sum of all enclave using samples found while the reported job was delayed.</p> $\% = \frac{\text{\# Holding Samples}}{\text{\# Samples}} * 100$ <p><b>Holding samples</b> The number of samples when the job was using the processor and delaying the other job (indicated in the Jobname field).</p> <p><b>Note:</b> In a multiprocessor environment, there is a holding job for each processor. For example, in a two-processor environment, two jobs can each account for 100% of the delay of the job waiting for the processor.</p>

## Monitor III Utility fields

Table 55 shows the additional fields you can select for this report.

Table 55. Additional Fields in the PROC Report

Field Heading	Meaning
PRCPCAP	This column contains the actual delay caused by WLM <ul style="list-style-type: none"> <li>• due to a resource group maximum</li> <li>• or due to discretionary goal management. This means that the work in question may be overachieving its goal. It may be capped to divert its resources to run discretionary work (see also section 'Using Discretionary Goals' in <i>z/OS MVS Planning: Workload Management</i>).</li> </ul> It shows the Dispatchable Unit (TCB or SRB) capped delay, which should be distinguished from the address space capping state shown in the CAPP field on the Work Manager Delays report.
PRCPODEL	Overall delay percentage for this address space.
PRCPOUSE	Overall using percentage for this address space.
PRCPTST	Percentage of the processor time used by non-enclave work that executed on behalf of this address space.
PRCPETST	Percentage of the processor time used by enclave and non-enclave work that executed within this address space.
PRCPAPPL	Percentage of the processor time used by non-enclave work that executed on behalf of this address space and processor type.
PRCPTWFL	Overall workflow percentage of this address space and processor type.
PRCPTUSE	Overall using percentage for this address space and processor type.
PRCPUCP	Overall using percentage on general purpose processors for zAAPs and zIIPs.
PRCPASI	Address space ID of the job.

### Report options

The PROC Report Options panel is similar to the DEV Report Options panel. See Figure 39 on page 72 for an example. Selecting YES for Jobs on the Report Options panel displays the Job Selection/Exclusion panel (see Figure 37 on page 70).

---

## PROCU - Processor Usage Report

The Processor Usage report (PROCU) displays all jobs that were using a general purpose or special purpose processor during the report interval. RMF reports the jobs by descending overall CPU time. The report gives you information about the percentage of CPU time on general purpose processors consumed on behalf of the job. In addition, the percentage of CPU time used by work that is eligible for being offloaded to an Application Assist (zAAP) or Integrated Information (zIIP) processor is shown. You can use this information to understand the benefit of adding a zAAP or zIIP into the configuration.

The EAppl fields also display the percentage of task time, SRB and enclave CPU time consumed within the address space on general purpose processor or special purpose (zAAP and zIIP) processors.

### How to request this report

To request the Processor Usage report, select **3** from the Primary Menu, then **1A** from the Resource Report Selection Menu (shown in Figure 8 on page 27) or enter the following command:

```
PROCU [job_class,service_class]
```

For example, to get a Processor Usage report for TSO service class TSOPRIME, enter:

## Contents of the report

```

RMF V2R2 Processor Usage
Command ==>>>
Line 1 of 20
Scroll ==>> PAGE

Samples: 60 System: MVS1 Date: 09/28/16 Time: 09.10.00 Range: 60 Sec

Jobname  Service  --- Time on CP % ---  ----- EAppl % -----
         CX Class  Total  AAP  IIP  CP  AAP  IIP
DBS3DIST SO DB2HIGH  100.6  0.0  0.1  111.0  0.0  21.3
WSWS7   O  OMVS      59.4  0.0  0.0  59.4  0.0  0.0
XCFAS   S  SYSTEM    24.1  0.0  0.0  24.1  0.0  0.0
WSP1S2F SO WASCR    14.0  0.1  0.0  14.0  2.6  0.0
TCP/IP  SO SYSSTC    16.1  0.0  0.0  16.1  0.0  0.0
VTAM44  S  SYSSTC    14.5  0.0  0.0  14.5  0.0  0.0
WSP1S2FS SO WASCR    10.2  2.7  0.0  42.5  98.4  0.0
OMVS    S  SYSTEM     3.7  0.0  0.0  3.7  0.0  0.0
CICS2TFA SO CICS2RGN  3.4  0.0  0.0  3.4  0.0  0.0
WSP1S6F SO WASCR     1.9  0.1  0.0  1.9  0.7  0.0
*MASTER* S  SYSTEM     2.4  0.0  0.0  2.4  0.0  0.0
CICS2AFA SO CICS2RGN  2.0  0.0  0.0  2.0  0.0  0.0
RMFGAT  SO SYSSTC     1.7  0.0  0.0  1.7  0.0  0.0
WSP1S6FS SO WASCR     1.0  0.3  0.0  5.3  7.7  0.0
    
```

Figure 83. PROCU - Processor Usage Report

## Field descriptions

Table 56. Fields in the PROCU Report

Field Heading	Meaning
Jobname	The name of a job using processor time. The processor usage report does not summarize data by job groups; all jobs within a job group are reported individually.
CX	Abbreviation for the job class as follows: <b>S</b> Started task <b>T</b> TSO <b>B</b> Batch <b>A</b> ASCH <b>O</b> OMVS  An <b>O</b> as second character indicates that the address space is using OMVS services.
Service Class	The name of the service class that a specified job has been running in.
Time on CP %	<p><b>Total</b> Percentage of CPU time spent on general purpose processors as sum of TCB time, global and local SRB time, and preemptable or client SRB time consumed on behalf of this address space.</p> <p><b>AAP</b> Percentage of CPU time on general purpose processors by this address space which was used by zAAP eligible work. This is a subset of the <b>Total</b> percentage.</p> <p><b>IIP</b> Percentage of CPU time on general purpose processors by this address space which was used by zIIP eligible work. This is a subset of the <b>Total</b> percentage.</p> <p><b>Note:</b> The calculated values are based on uniprocessor capacity, which means that they can exceed 100% on systems with more than one processor. To get the system utilization, this value has to be divided by the number of logical processors or cores.</p>

Table 56. Fields in the PROCU Report (continued)

Field Heading	Meaning
EAppl %	<p><b>CP</b> Percentage of CPU time on general purpose processors (CPs) as sum of TCB time, global and local SRB time, preemptable or client SRB time, and enclave CPU time consumed within this address space.</p> <p><b>AAP</b> Percentage of CPU time consumed on zAAPs within this address space.</p> <p><b>IIP</b> Percentage of CPU time consumed on zIIPs within this address space.</p> <p><b>Note:</b> EAppl% shows CPU utilization based on uniprocessor capacity. This means that the value can exceed 100% in systems with more than one processor. To get the system utilization, this value has to be divided by the number of logical processors or cores.</p>

### Monitor III Utility fields

Table 57 shows the additional fields you can select for this report.

Table 57. Additional Fields in the PROCU Report

Field Heading	Meaning
PRUPCLA	Class (A, B, E, O, S, or T)
PRUPTOTC	Percentage of total accumulated CPU time as sum of TCB time, global and local SRB time and preemptable or client SRB time, consumed on behalf of this address space.
PRUPTOTE	Percentage of total accumulated CPU time as sum of TCB time, global and local SRB time, preemptable or client SRB time, and enclave time consumed within this address space.
PRUPTCB	Percentage of TCB time consumed in this address space.
PRUPSRB	Percentage of SRB time consumed in this address space by local or global SRBs.
PRUPPCS	Percentage of preemptable or client SRB time consumed on behalf of this address space.
PRUPEPS	Percentage of preemptable or client SRB and enclave CPU time consumed within this address space.
PRUPASI	Address space ID of the job.

### Report options

The PROCU Report Options panel is similar to the Device Report Options panel. See Figure 39 on page 72 for an example. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed (see Figure 37 on page 70).

---

## RG - Resource Group Data Report

This report provides additional information about resource groups. It can also be used with historical data to get information about service policy definitions that are no longer available.

### How to request this report

To request the Resource Group Data report, select **U** from the Primary Menu, and then a **3** from the User menu, or you can enter the following command:

```
RG
```

### Contents of the report

The Resource Group Data report (RG report) is a modification of the Sysplex Summary report. It replaces the Sysplex Summary report columns containing **actual values** and the **Performance Index** with columns **Duration**, **Resource Group** and **Min**, **Max** and **Act Capacity**. Figure 84 on page 142 shows the RG

report modification of the Sysplex Summary report.

```

RMF V2R2 Sysplex Summary - RMFPLEX1           Line 1 of 35
Command ==>                                     Scroll ==> HALF
WLM Samples: 100   Systems: 8 Date: 09/28/16 Time: 10:31:40 Range: 100   Sec
Service Definition: SLA_2002                     Installed at: 07/01/15, 06.02.00
Active Policy: WORKDAYS                          Activated at: 07/01/15, 23.17.54
Exec Vel Resp. Time                             Resource ----- Capacity -----
Name      T I  Goal    Goal    Duration  Group      Min    Max    Act
ATMLABOR S 1          0.80 90%
BATCH.1  S          1.00M AVG    1000    PRODUCT    2000    4500 1471.007
        1 3
        2 D
BATCH.2  S D          2000    4500 1439.990
CICSTOR S 1          90          CICS      N/A     2100 752.0183
CICSAOR S 1          90          CICS      N/A     2100 846.8753
CICSDOR S 1          90          CICS      N/A     2100 492.0953
DEVELOP S          1.00 AVG    500
        1 1          3.00 90%    2000
        2 3          40
        3 5
POSMULTI S 1          0.50 90%
SYSSTC  S          N/A     N/A
    
```

Figure 84. RG Report

### Field descriptions

Fields particular to the RG report are described here. For a detailed description of other fields, see Table 82 on page 191.

Table 58. Fields in the RG Report

Field Heading	Meaning
Duration	Service class period duration in unweighted CPU service units (that means, not multiplied with the service coefficients) per second.  A duration is required in all but the last service class period. For single periods or for the last period of multiple periods this value is always zero.
Resource Group	Name of the resource group associated with the work in this service class.
Capacity Min	When capacity is defined in terms of service units with a sysplex scope, MIN is the minimum capacity in unweighted CPU service units per second of task or SRB execution time that should always be available to this resource group.  When capacity is defined in terms of percentage of the LPAR share on a system scope, MIN is the percentage of the LPAR share of this system that should always be available to this resource group.  When capacity is defined as a number of general purpose processors (CPs), MIN is the number of CPs that should always be available to this resource group. A number of 100 represents the capacity of one CP.
Capacity Max	When capacity is defined in terms of service units with a sysplex scope, MAX is the maximum capacity in unweighted CPU service units per second of task or SRB execution time this resource group may use across the sysplex.  When capacity is defined in terms of percentage of the LPAR share on a system scope, MAX is the percentage of the LPAR share of this system that this resource group should never exceed.  When capacity is defined as a number of general purpose processors (CPs), MAX is the number of CPs that this resource group may use, whereby a number of 100 represents the capacity of one CP.



Table 58. Fields in the RG Report (continued)

Field Heading	Meaning
Capacity Actual	Actual capacity, in unweighted CPU service units per second, as consumed within that resource group. The data presented is related to the specified service class.

## RLSLRU - VSAM LRU Overview Report

This report provides Local Buffer Manager LRU statistics for each system. The data in this report can help you in adjusting the goal and the limit for the local cache size.

In Parmlib member IGDSMSxx, there is a goal response limit for the local cache size, it defaults to 100 MB. You can specify a limit up to 1.5 GB (if a bigger value will be given, the report will display MAX as buffer size goal). Each LRU cycle, it is determined whether the system is over the goal and the buffer aging algorithms are accelerated. If the system is 5 times over the goal or reaches the 1.5 GB limit, the system starts clearing the buffers. If systems appear where BMF is over the goal (status *Accelerated* or *Reclaimed*), you could adapt the goal in Parmlib member IGDSMSxx by changing the RLS\_MAX\_POOL\_SIZE value.

### How to request this report

To request the VSAM LRU Overview report, select **S** from the Primary Menu and then select **12** on the Sysplex Report menu (shown in Figure 5 on page 24), or enter the following command:

```
RLSLRU
```

In addition, you can navigate to this report with cursor-sensitive control from the VSAM RLS Activity report.

### Contents of the report

```

RMF V2R2 VSAM LRU Overview - SYSPLEX           Line 1 of 2
Command ==>>>                               Scroll == => HALF

Samples: 120   Systems: 2   Date: 09/28/16   Time: 13.25.00   Range: 120   Sec

MVS          Avg CPU - Buffer Size - Accl  Reclaim  ----- Read -----
System       Time    Goal High   %      %      BMF%  CF%  DASD%

SYS4
Below 2GB   0.023    MAX   1M    0.0   0.0   0.0  0.0  0.0
Above 2GB   3.543    MAX   1M    0.0   0.0  97.5  0.0  2.5
SYS5
Below 2GB   4.457    MAX   1M    0.0   0.0   0.0  0.0  0.0
    
```

Figure 85. RLSLRU Report

Cursor-sensitive control on a system line displays a pop-up panel with buffer counts by pool for the selected system. There are sixteen storage pools (2K, 4K, ... 32K) available.

RMF VSAM LRU Overview - Buffer Counts by Pool						
The following details are available for MVS System: SYSF						
Press Enter to return to the Report panel.						
Fixed Pages	Low	:	305	Fixed Storage	:	500
	High	:	305	Real Storage %	:	46
	Average	:	305			
	----- Below 2 GB -----			----- Above 2 GB -----		
Size	Low	High	Avg	Low	High	Avg
						More: +
2K	163	315	226	0	0	0
4K	713	1537	1299	0	0	0
6K	0	0	0	0	0	0
8K	0	0	0	0	0	0
10K	0	0	0	0	0	0
12K	0	0	0	347	458	412
14K	0	0	0	0	0	0
16K	460	678	656	0	0	0
18K	0	0	0	0	0	0
20K	0	0	0	0	0	0
22K	0	0	0	0	0	0
24K	0	0	0	0	0	0
26K	0	0	0	0	0	0
28K	0	0	0	0	0	0
30K	0	0	0	0	0	0
32K	0	0	0	0	0	0

Figure 86. VSAM LRU Overview - Buffer Counts by Pool

### Field descriptions

**Note:** If applicable, the measurements in the RLRLRU Report are presented for storage addresses below and above the 2GB bar. Two lines of data are then displayed for each system and a label *Above 2GB* or *Below 2GB* precedes the corresponding line as shown in Figure 85 on page 143. Also, all measurements in the Buffer Counts by Pool panel are presented for storage addresses below and above the 2GB bar. This is indicated by the corresponding column headings as shown in Figure 86.

Table 59. Fields in the RLRLRU Report

Field Heading	Meaning
MVS System	System name.
Avg CPU Time	Average CPU time spent by BMF LRU processing during each report interval (milliseconds).
Buffer Size: Goal	Buffer size goal (MB). If no valid goal has been defined, MAX will be shown.
Buffer Size: High	Buffer size actual high value (MB).
Accel%	Percentage of Buffer Manager LRU intervals when BMF was over the goal and buffer aging algorithms were accelerated.
Reclaim%	Percentage of Buffer Manager LRU intervals when BMF was over the goal and buffer aging algorithms were bypassed to reclaim buffers.
Read	<p><b>BMF%</b> Percentage of READ requests that could be satisfied from local buffers being managed by SMSVSAM.</p> <p><b>CF%</b> Percentage of CF cache structure READ requests.</p> <p><b>DASD Read %</b> Percentage of READ requests to DASD.</p>

**Fields in the Buffer Counts by Pool panel:** There are 16 buffer pools with different buffer sizes between 2K and 32K incremented by 2K. For each pool, this panel presents a line with the high, low and average numbers of BMF buffers during this interval. These values are provided for storage pools allocated to addresses below and above the 2GB-bar. If values above 2GB are not available, '-' is displayed.

Field Heading	Meaning
Fixed Storage	The amount of buffer pool storage that is specified to be fixed. If no value is available, a '-' is shown.
Real Storage %	The value specified by the RLSFIXEDPOOLSIZE parameter divided by the amount of real storage in the system (in percent). If no value is available, a - is shown.
Fixed Pages	Low, high and average actual number of fixed 4KB-pages. If no value is available, a '-' is shown.

## RLSSC/RLSDS - VSAM RLS Activity Report

With VSAM RLS, GETs and PUTs are executed by SMSVSAM on behalf of the application. When the application's data request can be satisfied from SMSVSAM's local buffers, no I/O is necessary. If the data in the buffers is invalid, SMSVSAM accesses the VSAM RLS cache structures in the coupling facility in order to fulfill the request. If data in the cache structures is invalid, a DASD I/O is performed.

The report is providing VSAM RLS activity data regarding READ and WRITE requests accessing the local buffers, the CF cache structures and DASD. This data might help you in answering important questions like

- Are there problems with LRU (Least Recently Used algorithms) or buffer pool sizes?
- Are the CF cache structures too small?

### How to request this report

The VSAM RLS Activity report has two different versions:

- VSAM RLS Activity by Storage Class
- VSAM RLS Activity by Data Set

Both versions of the report have a similar structure, but a different scope of data. From the initial Sysplex Total View for each version, you can navigate to a System/CF Structure View.

To request the VSAM RLS Activity report, select **S** from the Primary Menu and then select **10** (for storage class) or **11** (for data set) on the Sysplex Report menu (shown in Figure 5 on page 24), or enter one of the following commands:

```
RLSSC [storage_class]
```

```
RLSDS [data_set_name]
```

In addition, you can navigate between the two versions of the report with cursor-sensitive control.

### Contents of the report

#### VSAM RLS Activity by Storage Class

This report provides a VSAM RLS activity view by storage class. For each storage class, sysplex wide totals will be displayed for direct and sequential access.

```

RMF V2R2  VSAM RLS Activity - SYSPLEX          Line 1 of 12
Command ==>                               Scroll == => HALF

Samples: 59      Systems: 2      Date: 09/28/16  Time: 13.16.00  Range: 60  sec

          < 2GB / > 2GB
LRU Status : Good / Accl
Contention % : 0.0 / 0.0
False Cont % : 0.0 / 0.0

Stor Class  Access  Resp  ----- Read -----  ----- BMF -----  Write
            Access  Time   Rate  BMF%  CF%  DASD%  Valid%  False Inv%  Rate

RLS
Below 2GB  DIR    0.004  665.6  88.2  0.5  11.3  100    0.01    0.00
           SEQ    0.000   0.00  0.0  0.0  0.0  0.0    0.00    0.00
Above 2GB  DIR    0.004  665.6  88.2  0.5  11.3  100    0.01    0.00
           SEQ    0.000   0.00  0.0  0.0  0.0  0.0    0.00    0.00
RLS1
Below 2GB  DIR    0.005  200.0  90.5  0.0  9.5   100    0.00    0.00
           SEQ    0.000   0.00  0.0  0.0  0.0  0.0    0.00    0.00
RLS2
Below 2GB  DIR    0.003  213.3  90.5  0.0  9.5   100    0.00    0.00
           SEQ    0.000   0.00  0.0  0.0  0.0  0.0    0.00    0.00
RL3
Above 2GB  DIR    0.004  665.6  88.2  0.5  11.3  100    0.01    0.00
           SEQ    0.000   0.00  0.0  0.0  0.0  0.0    0.00    0.00
    
```

Figure 87. VSAM RLS Activity by Storage Class - Sysplex Total View

Cursor-sensitive control on the LRU STATUS field displays the VSAM LRU Overview report (see Figure 85 on page 143) with the Local Buffer Manager LRU statistics for each system.

Cursor-sensitive control on a storage class name redisplay the report with a system and CF cache structure breakdown for the selected storage class.

```

RMF V2R2 VSAM RLS Activity - SYSPLEX Line 1 of 23
Command ==> Scroll == => HALF
Samples: 120 Systems: 2 Date: 09/28/16 Time: 13.25.00 Range: 120 Sec

< 2GB / > 2GB
LRU Status : Good / Acce1 Stor Class : RLS
Contention % : 0.0 / 0.0 Cache Set : PUBLIC1
False Cont % : 0.0 / 0.0 Lock Set : RLSLOCKSET
Lock Struct: RLSLOCKSTR

```

System/CF	Access	Resp Time	Read Rate	Read BMF%	Read CF%	Read DASD%	BMF Valid%	Write False Inv%	Write Rate
*ALL									
Below 2GB	DIR	0.000	14.98	83.0	0.0	17.0	100	0.00	0.00
	SEQ	0.000	0.00	0.0	0.0	0.0	0.0	0.00	0.00
Above 2GB	DIR	0.000	14.98	83.0	0.0	17.0	100	0.00	0.00
	SEQ	0.000	0.00	0.0	0.0	0.0	0.0	0.00	0.00
SYS4									
CACHE01									
Below 2GB	DIR	0.000	7.49	83.0	0.0	17.0	100	0.00	0.00
	SEQ	0.000	0.00	0.0	0.0	0.0	0.0	0.00	0.00
Above 2GB	DIR	0.000	7.49	83.0	0.0	17.0	100	0.00	0.00
	SEQ	0.000	0.00	0.0	0.0	0.0	0.0	0.00	0.00
CACHE02									
Below 2GB	DIR	0.000	0.00	0.0	0.0	0.0	0.0	0.00	0.00
	SEQ	0.000	0.00	0.0	0.0	0.0	0.0	0.00	0.00
CACHE03									
Above 2GB	DIR	0.000	0.00	0.0	0.0	0.0	0.0	0.00	0.00
	SEQ	0.000	0.00	0.0	0.0	0.0	0.0	0.00	0.00

Figure 88. VSAM RLS Activity by Storage Class - System/CF Structure View

At the top of the report, the sysplex wide totals for the storage class will be displayed (indicated by \*ALL) followed by report lines per system and CF cache structure.

Cursor-sensitive control on a CF structure name displays the CF Activity report (see Figure 23 on page 43) for this structure.

### VSAM RLS Activity by Data Set

This report provides a VSAM RLS activity view by VSAM data sets. The information will be grouped by VSAM spheres. A sphere consists of components, i.e. data sets like BASE.DATA, BASE.INDEX, ALT.DATA and ALT.INDEX. For each data set, sysplex wide totals will be displayed for direct and sequential access. Only those VSAM spheres will be presented for which data collection has been requested.

```

RMF V2R2 VSAM RLS Activity - SYSPLEX          Line 1 of 20
Command ==>                               Scroll == => HALF

Samples: 120   Systems: 2   Date: 09/28/16   Time: 13.25.00   Range: 120   Sec

          < 2GB / > 2GB
LRU Status : Good / Acce1
Contention % : 0.0 / 0.0
False Cont % : 0.0 / 0.0

Sphere/DS  Access  Resp  ----- Read -----  ----- BMF -----  Write
           Access  Time   Rate  BMF%  CF%  DASD%  Valid%  False Inv%  Rate

BMAI.VSAMIN.MEGA
BMAI.VSAMIN.MEGA.AIX.DATA
  Below 2GB DIR    0.003  0.01  0.0  0.0   100    0.0    0.00    0.00
            SEQ    0.000  0.00  0.0  0.0    0.0    0.0    0.00    0.00
  Above 2GB DIR    0.003  0.01  0.0  0.0   100    0.0    0.00    0.00
            SEQ    0.000  0.00  0.0  0.0    0.0    0.0    0.00    0.00
BMAI.VSAMIN.MEGA.AIX.INDEX
  Below 2GB DIR    0.003  0.03  50.0  0.0   50.0   100    0.00    0.00
            SEQ    0.000  0.00  0.0  0.0    0.0    0.0    0.00    0.00
  Above 2GB DIR    0.003  0.03  50.0  0.0   50.0   100    0.00    0.00
            SEQ    0.000  0.00  0.0  0.0    0.0    0.0    0.00    0.00
BMAI.VSAMIN.MEGA.DATA
  Below 2GB DIR    0.000  7.45  83.2  0.0   16.8   100    0.00    0.00
            SEQ    0.000  0.00  0.0  0.0    0.0    0.0    0.00    0.00
MAI1.VSAMIN.MEGA
MAI1.VSAMIN.MEGA.AIX.DATA
  Above 2GB DIR    0.003  0.01  0.0  0.0   100    0.0    0.00    0.00
            SEQ    0.000  0.00  0.0  0.0    0.0    0.0    0.00    0.00
    
```

Figure 89. VSAM RLS Activity by Data Set - Sysplex Total View

Cursor-sensitive control on the LRU STATUS field displays the VSAM LRU Overview report (see Figure 85 on page 143) with the Local Buffer Manager LRU statistics for each system.

Cursor-sensitive control on a data set name redisplay the report with a system breakdown for the selected VSAM data set. At the top of the report, the sysplex wide totals for the data set will be displayed (indicated by \*ALL) followed by report lines per system.

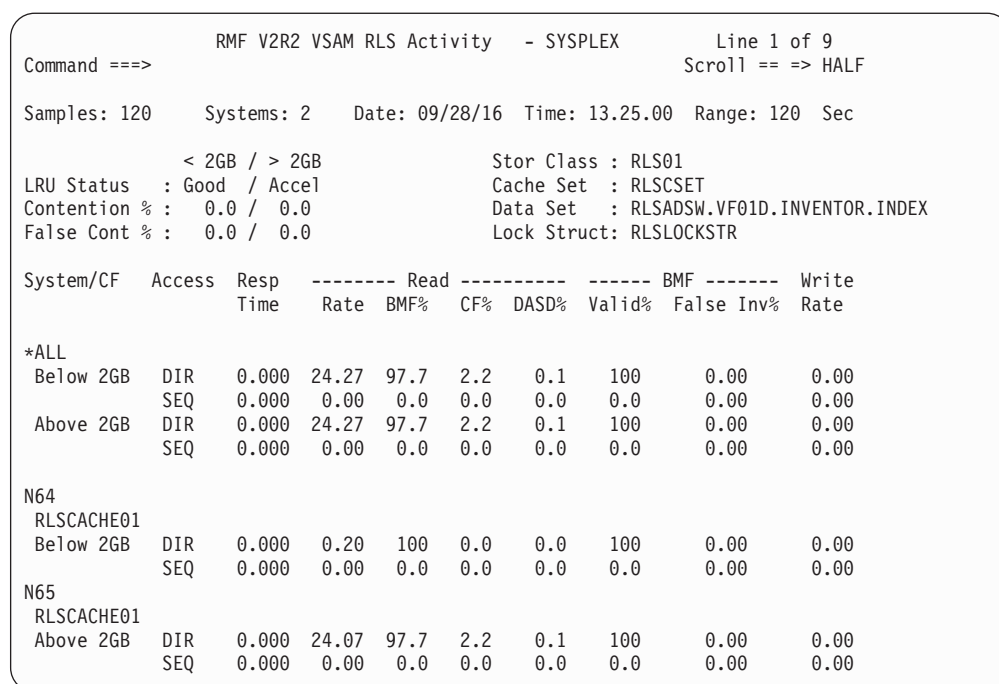


Figure 90. VSAM RLS Activity by Data Set - System/CF Structure View

Cursor-sensitive control on a CF structure name displays the CF Activity report (see Figure 23 on page 43) for this structure.

### Field descriptions

**Note:** If applicable, all measurements in the VSAM RLS Activity reports are presented for storage addresses below and above the 2GB bar. This is either indicated by the heading *< 2GB / > 2GB* or by the labels *Above 2GB* or *Below 2GB*.

Table 60. Fields in the VSAM RLS Activity report

Field Heading	Meaning
The following information is displayed in the report header depending on the report view:	
LRU Status	LRU status of local buffers under control of BMF (Buffer Management Facility). <b>Good</b> BMF is at or below its goal on all systems. <b>Accelerated</b> BMF is over the goal on at least one system, and the buffer aging algorithms were accelerated. <b>Reclaimed</b> BMF is over the goal on at least one system, and the buffer aging algorithms were bypassed to reclaim buffers.
Contention %	Percentage of true LOCK contentions: all external requests issued by connectors delayed due to contention on a lock. If the value above 2 GB is not available, '-' is displayed.
False Cont %	Percentage of false LOCK contentions: All external requests issued by connectors that experience "hash contention".  This occurs because a hashing algorithm is used to map a lock request to a lock table entry. When more than one lock request maps to a lock table entry, there is the potential for contention delay. You may need to increase the size of the lock table.  If the value above 2 GB is not available, '-' is displayed.
Stor Class	Storage class name.
Cache Set	DFSMS cache set name.
Data Set	VSAM data set name.

## Mon III - RLSSC/RLSDS

Table 60. Fields in the VSAM RLS Activity report (continued)

Field Heading	Meaning
Lock Set	Lock set name. Multiple Lock Structure support allows defining an additional Coupling Facility DFSMS lock structure to be associated with a single SMS storage class. Each lock set can contain a single lock structure name.  If multiple lock structures are not supported, the line remains blank.
Lock Struct	Name of the lock structure associated with the lock set. If no lock structure is associated, IGWLOCK00 is used and reported. If multiple lock structures are not supported, IGWLOCK00 is also used, however, the line remains blank.
The body of the report contains the following columns. The first column depends on the report flavour and view.	
Stor Class	Storage class name (first column in the Sysplex Total View in the Storage Class flavour).
System/CF	System name (*ALL indicates a sysplex-wide view) and the CF cache structure name (first column in the System/CF Structure View in both flavours).
Sphere/DS	VSAM sphere name and the VSAM data set name (first column in the Sysplex Total View in the Data Set flavour).
Access	Indicates whether the values in this row are shown for direct access (DIR) or sequential access (SEQ). There may be up to two sets of data depending on whether buffers above or below 2 GB were accessed.
Resp Time	Average response time of all requests (seconds).
Read Rate	Total number of BMF READ requests per second. BMF READ requests is sum of BMF valid READ hits, CF READ hits and DASD READs.  The value is reported as '<0.01' if the rate is greater than 0 but below 0.01.
BMF Read%	Percentage of BMF valid READ hits that is the percentage of READ requests that were satisfied from local buffers being managed by SMSVSAM.
CF Read%	Percentage of READ requests that were satisfied by the CF cache structure.
DASD Read%	Percentage of READ requests to DASD.
BMF Valid%	Percentage of BMF READ hits that were valid. If a buffer is found in the local cache and is determined to be valid according to the information in local control blocks, this counts as a BMF valid READ hit. <b>Note:</b> A BMF READ hit is determined to be valid based on the IXLVECTR local vector service TestLocalCache. If it is invalid based on IXLVECTR, this counts as a BMF invalid READ hit. BMF READ hits is the sum of valid and invalid READ hits. If IXLVECTR indicates the buffer to be valid, it can be used. If invalid, the buffer can not be used. There are two reasons for indicating a buffer to be invalid: <ul style="list-style-type: none"> <li>• Another system has altered the data which has been locally buffered. Thus, the copy in the BMF local cache became out-of-date (BMF true invalid READ hits).</li> <li>• The coupling facility has lost track of the integrity status of the buffer (BMF false invalid READ hits).</li> </ul> To make use of BMF Valid%, following formulas are helpful: <ul style="list-style-type: none"> <li>• <math>\text{BMF Invalid Read Hits} = \text{True} + \text{False Invalid Read Hits}</math></li> <li>• <math>\text{BMF Read Hit\%} = \text{BMF Read\%} / \text{BMF Valid\%} * 100</math></li> <li>• <math>\text{BMF Invalid Read Hit\%} = \text{BMF Read Hit\%} - \text{BMF Read\%}</math></li> </ul>
BMF False Inv%	Percentage of READ requests when the copy in the BMF local cache was invalid because the coupling facility has lost track of the integrity status of the buffer.
Write Rate	Total number of BMF WRITE requests per second.  The value is reported as '<0.01' if the rate is greater than 0 but below 0.01.

## SCM - Storage Class Memory (SCM) Activity Report

The storage class memory (SCM) activity report can be used to investigate performance problems that are related to storage class memory (SCM).



## How to request this report

To request the SCM Activity Report, select a 3 from the Primary Menu, then select 15 from the Resource Report Selection Menu (shown in Figure 8 on page 27) or enter the following command:

SCM

## Contents of the report

The storage class memory (SCM) activity report provides these types of information:

### EADM device/subchannel level information

The EADM (extended asynchronous data mover) device summary section at the top of the report provides the rate of start subchannel (SSCH) instructions for all EADM devices together with response time statistics consisting of pending, IOP queue and initial command response time.

### Flash Express card level information

For each Flash Express card, the report provides measurements at both the LPAR and CPC level. The rate at which internal requests are processed by the adapter card, the rate at which data units where read and written, the average response and IOP queue time is displayed.

Command ==>>		RMF V2R2 SCM Activity		Line 1 of 2	
				Scroll ==> CSR	
Samples: 120	System: CB8C	Date: 04/24/15	Time: 04.07.00	Range: 120	Sec
----- EADM Device/Subchannel Summary -----					
SSCH Total	SSCH Rate	PEND Time	IOPQ Time	ICMR Time	
1	0.01	0.128	0.000	0.000	
Card ID	Util(%)	Read(B/s)	Write(B/s)	Req Rate	Resp Time
	Part Total	Part Total	Part Total	Part Total	Part Total
0001-0580	0.00 0.01	0.00 36932	0.00 152K	0.00 46.08	0.000 0.379
0001-05AC	0.00 0.01	34.13 37410	0.00 189K	0.01 55.27	0.256 0.327

Figure 91. SCM Activity Report

Table 61. Fields in the SCM Activity Report

Field Heading	Meaning
EADM device/subchannel summary.	
This section provides summary information about the extended asynchronous data mover (EADM) devices or subchannels. EADM subchannels are similar to I/O subchannels in a way that I/O instructions can be issued. However, they do not have channel paths or device numbers assigned, and they are not defined in the I/O configuration. They are created automatically during IPL.	
SSCH Total	The total number of SSCH instructions to all EADM devices in the report interval.
SSCH Rate	The number of SSCH instructions to all EADM devices per second.
PEND Time	The average function pending time across all EADM devices in milliseconds. This is similar to function pending time for traditional I/O devices, which is the amount of time between when the SSCH is issued and the first command in the channel program is accepted.  $P\text{END} = \frac{\text{Sum( Function Pending Time )}}{\text{Measurement Event Count}}$

## Mon III - SCM

Table 61. Fields in the SCM Activity Report (continued)

Field Heading	Meaning
IOPQ Time	The average IOP queue time across all EADM devices in milliseconds. This is unique to EADM devices. It represents the amount of time the request is not accepted by the adapter because it would exceed its maximum capacity. For a particular I/O request, this may occur multiple times.  $\text{IOPQ} = \frac{\text{Sum( IOP Queue Time )}}{\text{Measurement Event Count}}$
ICMR Time	The average initial command response time across all EADM devices in milliseconds. This is the time from when the first command does not immediately proceed to execute until the successful start of execution at the SCM resource part.  $\text{ICMR} = \frac{\text{Sum( Initial Command Response Time )}}{\text{Measurement Event Count}}$
Flash Adapter measurements	
Card ID	The identifier of the flash adapter card.
Following fields are displayed at a system-wide level (Total) and for the current LPAR (Part) whereby IOPQ Time is only available at the total level.	
Util(%)	The average utilization of the flash card during the interval as reported by the SCM measurement facility.
Read(B/s)	Bytes read per second.
Write(B/s)	Bytes written per second.
Req Rate	The requests processed per second.
Resp Time	The average response time per request in milliseconds.
IOPQ Time	The average IOP queue time per request in milliseconds.

## SPACED - Disk Space Report

The Disk Space Report displays capacity and disk space information for volumes. This report displays only those volumes that belong to storage groups specified with the Monitor III SGSPACE gatherer option. You can use this information to decide whether a certain volume provides sufficient free disk space for new allocation requests.

Together with the Storage Space Report, this report can help to make decisions for long-term disk space capacity planning.

Although the Disk Space Report is a single system report, the report combines the data collected from all systems within the sysplex. This allows you to gather the data for an SMS Storage Group only on one system of the sysplex so that the collection of redundant data can be avoided.

### How to request this report

To request the Disk Space Report, select **1** from the Primary Menu, then select **11** from the Overview Report Selection Menu (shown in Figure 6 on page 25) or enter one of the following commands:

```
SPACED
SPD
```

In addition, you can invoke the Disk Space Report for all volumes of a selected storage group by using cursor-sensitive control from the Storage Space Report (SPACEG).

## Contents of the report

RMF V2R2 Disk Space Report						Line 1 of 90
Samples: 60	System: TRX2	Date: 09/28/16	Time: 08.42.00	Range: 60	Sec	
Volume	Total (MB)	Free (MB)	Free (%)	Largest Ext (MB)	Storage Group	
SYSSD1	8120	2922	36.0	2922	DB2	
SYSSD3	8120	2291	28.2	2291	DB2	
SYSSD2	8120	2074	25.5	2074	DB2	
SYSOPE	8120	6326	77.9	6326	OMVSSYS	
SYSSM5	8120	1164	14.3	40	SMS	
SYSSM3	8120	1034	12.7	233	SMS	
SYSSM6	8120	1017	12.5	294	SMS	
SYSSM2	8120	1004	12.4	198	SMS	
SYSSMS	8120	982	12.1	62	SMS	
SYSSM4	8120	947	11.7	34	SMS	
SYSSM7	8120	728	9.0	139	SMS	

Figure 92. Disk Space Report

### Field descriptions

Table 62. Fields in the Disk Space Report

Field Heading	Meaning
Volume	Name of the volume belonging to a monitored storage group.
Total (MB)	Total amount of disk space (in megabytes) on the volume.
Free (MB)	Total amount of free disk space (in megabytes) on the volume.
Free (%)	Percentage of free disk space on the volume.
Largest Ext (MB)	Largest block (extent) in megabytes of unallocated disk space available on the volume.
Storage Group	Name of the storage group to which the volume belongs.

## SPACEG - Storage Space Report

A storage group is a collection of storage volumes and attributes, defined by the storage administrator and treated as a single object storage hierarchy. The Storage Space Report allows you to keep track of disk space consumption on a storage group level. This report displays only those volumes that belong to storage groups specified with the Monitor III SGSPACE gatherer option.

From this report, you can see whether the system can provide sufficient disk space for new allocation requests. This report may also be useful for making decisions in long-term disk space capacity planning.

Although the Storage Space Report is a single system report, the report combines the data collected from all systems within the sysplex. This allows you to gather the data for an SMS Storage Group only on one system of the sysplex so that the collection of redundant data can be avoided.

### How to request this report

To request the Storage Space report, select **1** from the Primary Menu, then select **10** from the Overview Report Selection Menu (shown in Figure 6 on page 25) or enter one of the following commands:

SPACEG  
SPG

## Contents of the report

RMF V2R2 Storage Space Report				
Samples: 30	System: TRX2	Date: 09/28/16	Time: 08.42.00	Range: 30 Sec
SGroup	Total (MB)	Free (MB)	Free (%)	Volumes
*ALL	322116	165322	51.3	100 *
SGSMB	70382	35610	50.6	28 *
SGZFS	78503	39376	50.2	32 *
SMSCAT1	2707	398	14.7	1
SMS3390B	16240	8683	53.5	2
S1P01	18947	18802	99.2	3
S1P03	24360	23946	98.3	3
TSODA1	21654	1473	6.8	4
TSODA2	8120	411	5.1	1
TSODA3	5414	139	2.6	2
TSODA4	5414	86	1.6	2
USSFS	70375	36398	51.7	12

Figure 93. Storage Space Report

Cursor-sensitive control on a storage group name displays the Disk Space Report with all volumes of this storage group.

### Field descriptions

Table 63. Fields in the Storage Space Report

Field Heading	Meaning
SGroup	Name of the storage group connected to the system. The line showing *ALL in this column presents the accumulated values or average percentage values for all storage groups.
Total (MB)	Total amount of disk space (in megabytes) on all online volumes in the storage group.
Free (MB)	Total amount of free disk space (in megabytes) on all online volumes in the storage group.
Free (%)	Percentage of free disk space in the storage group.
Volumes	Number of volumes in the storage group. If at least one volume did not return any space information, the number is followed by an *.

## STOR - Storage Delays Report

The Storage Delays report is job-oriented. It displays storage delay information for all jobs.

### How to request this report

To request the Storage Delays report, select **3**, from the Primary Menu, then select **6** from the Resource Report Selection Menu (shown in Figure 8 on page 27) or enter the following command:

```
STOR [job_class,service_class]
```

For example, to get a Storage Delays report for TSO service class TSOPRIME, enter:

```
STOR T, TSOPRIME
```

## Contents of the report

```

Command ==>>>
RMF V2R2 Storage Delays
Line 1 of 103
Scroll ==>> CSR

Samples: 119 System: SYSF Date: 09/28/16 Time: 10.52.00 Range: 120 Sec

Jobname C Service DLY ----- % Delayed for ----- -- Working Set --
          C Class %      COMM LOCL SWAP OUTHR OTHR Central Expanded
*MASTER* S SYSTEM 0 0 0 0 0 0 0 3514
PCAUTH S SYSSTC 0 0 0 0 0 0 0 119
RASP S SYSTEM 0 0 0 0 0 0 0 284
TRACE S SYSSTC 0 0 0 0 0 0 0 271
DUMPSRV S SYSTEM 0 0 0 0 0 0 0 423
XCFAS S SYSTEM 0 0 0 0 0 0 0 8817
GRS S SYSTEM 0 0 0 0 0 0 0 8048
SMSPDSE S SYSTEM 0 0 0 0 0 0 0 1972
SMSVSAM S SYSTEM 0 0 0 0 0 0 0 3350
CONSOLE S SYSTEM 0 0 0 0 0 0 0 3303
WLM S SYSTEM 0 0 0 0 0 0 0 2603
ANTMAIN S SYSTEM 0 0 0 0 0 0 0 1162
ANTAS000 S STCDEF 0 0 0 0 0 0 0 1194
DEVMAN S SYSTEM 0 0 0 0 0 0 0 178
    
```

Figure 94. STOR Report

The graphic form of this report shows the percentage of each user's time that COMM, LOCL, SWAP, OUTHR, and OTHR contributed to the delay of the job for storage.

### Field descriptions

Table 64. Fields in the STOR Report

Field Heading	Meaning
Jobname	Name of a job that is delayed for storage. The STOR report does not summarize data by job groups; all jobs within a job group are reported individually.
C	A one-character abbreviation for the job class as follows: <b>A</b> ASCH <b>B</b> Batch <b>O</b> OMVS <b>S</b> Started task <b>T</b> TSO <b>?</b> Data is missing or invalid.
Service Class	The name of the service class that a specified job has been running in.
DLY %	Delay the waiting job (address space) is experiencing because of contention for storage during the range period, expressed as a percentage. $DLY \% = \frac{\# \text{ Delay Samples}}{\# \text{ Samples}} * 100$ <b>Delay samples</b> The number of samples where the job was delayed for storage one or more times due to COMM, LOCL, OTHR, SWAP, and OUTHR <b>Note:</b> This value is also found in the STOR field on the Delay report.

## Mon III - STOR

Table 64. Fields in the STOR Report (continued)

Field Heading	Meaning
% Delayed for	<p>The percentage that COMM, LOCL, OTHR, SWAP, and OUTR contribute to the delay of the job for storage. If there is no overlap of the delay states, the percentages for all these resources add up to the DLY % value; if there is overlap, the percentages add up to more than the DLY % value.</p> <p>The meaning of each category follows:</p> <p><b>COMM</b> The percentage that common (CSA or LPA) storage paging contributes to the delay of the job from the time of the page fault until I/O is completed. The shared storage paging is also included in the COMM count.</p> <p><b>LOCL</b> The percentage that local (private) storage paging contributes to the delay of the job from the time of the page fault until I/O is completed. The shared storage paging is also included in the LOCL count.</p> <p><b>SWAP</b> The percentage that swapping contributes to the delay of the job from the time of swap initiation until the last swap page I/O is completed.</p> <p><b>OUTR</b> The percentage that being swapped out and ready contributes to the delay of the job.</p> <p><b>OTHR</b> The sum of:</p> <ul style="list-style-type: none"> <li>• VIO (virtual I/O)</li> <li>• XMEM (paging delays from cross memory address spaces)</li> <li>• HIPR (Paging delays from standard hiperspaces, except ESO hiperspaces) This is a single state sum, which means that whenever several of the delays are detected in the same cycle, they are counted as one occurrence of OTHR delay.</li> </ul> <p>Cursor-sensitive control on this field gives you the STORR report.</p> <p>For COMM, LOCL, OTHR, and SWAP, RMF scans all ASM AIA chains. If the address space has one or more incomplete page input requests, RMF updates the counter in the appropriate category once per sample.</p>
Working Set	<p>The working set value represents the average amount of storage (in frames) a user occupied while in central and expanded storage (not swapped), including dataspaces and hiperspaces. The second column is blank if the system is running in 64-bit mode.</p>

### Monitor III Utility fields

You can use the Monitor III Utility to customize the Storage Delays report. In addition to the delays previously described, you can use the Utility to have the delay percentages in Table 65 shown in the Storage Delays report.

Table 65. Additional Fields in the STOR Report

Field Heading	Meaning
% delayed for VIO	The percentage of time a job was delayed because of virtual I/O.
% delayed for XMEM	The paging delays from cross memory address spaces.
% delayed for HIPR	The paging delays from standard hiperspaces (including waits during scroll wait), but not ESO hiperspaces.
Average ACTV frames	The average number of central storage frames held by the job while it was active.
Average fixed frames total	The average number of fixed frames the job was using during the report interval including frames both above and below the 16 megabyte line.
Average IDLE frames	The average number of frames held by the job while it was idle.

### Report options

The STOR Report Options panel is similar to the Device Report Options panel. See Figure 39 on page 72 for an example. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed. See Figure 37 on page 70 for an example.

## STORC - Common Storage Report

This report provides information about the use of common storage (CSA, ECSA, SQA, and ESQA) within a system.

The top section of the report provides overall system information about the use of common storage. For more information about the fields in this section, see Table 66 on page 158.

The bottom section of the report provides job-related information about the use of common storage for jobs active during the specified report interval. The jobs are sorted by descending storage percentage; that is, for each job with the maximum of the four common storage percentages, the job with the highest maximum percentage is reported first. These fields are described in Table 67 on page 159.

**Note:**

1. The report can be incomplete for some jobs, this will be indicated by messages ERB617I, ERB618I, or ERB619I. They explain that CSA/ECSA or SQA/ESQA data needed for RMF reporting was not completely gathered. VSM common storage (CSA, ECSA, SQA, and ESQA) tracking was either not active or partially active since the job started.

The common storage data fields are reported in dark blue for those jobs that data was partially gathered.

If VSM common storage tracking was partially active, you can use the BREF command to select a range period when common storage data gathering was active.

If VSM common storage tracking was not active, contact your system programmer to activate VSM common storage tracking by issuing:

```
SET DIAG=01
```

The defaults in the SYS1.PARMLIB member DIAG01 are:

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

2. In the calculations used for this report, when CSA is converted to SQA but not allocated, the amount is still considered part of allocated CSA. Only when the converted CSA is allocated is it considered allocated SQA. Because CSA to SQA conversion can increase SQA to a value larger than defined at IPL, percent values of SQA can be greater than 100%. This also applies for ECSA to ESQA conversions.

## How to request this report

To request the Common Storage report, select **3**, from the Primary Menu, then select **10** from the Resource Report Selection Menu (shown in Figure 8 on page 27) or enter the following command:

```
STORC [job_class,service_class]
```

For example, to get a Common Storage report for TSO service class TSOPRIME, enter:

```
STORC T, TSOPRIME
```

## Contents of the report

```

RMF V2R2 Common Storage                               Line 1 of 680
Command ==>>>                                       Scroll ==>> HALF

Samples: 98      System: SYSF Date: 09/28/16 Time: 13.38.10 Range: 100 Sec

System Information
IPL Definitions          ----- Percent -----
                         CSA ECSA SQA ESQA      CSA ECSA SQA ESQA
Peak Allocation Values   17  12  55  18      578K  38M  914K  25M
Average CSA to SQA Conversion  0   0           0   0
Average Use Summary     17  12  55  18      578K  37M  914K  25M
Available at End of Range 83  88  45  82      2754K 263M  746K  117M

Unalloc Common Area: 3400K

Jobname  Service      ELAP  -- Percent Used -  ----- Amount Used -----
Act C Class  ASID  Time  CSA ECSA SQA ESQA  CSA ECSA SQA ESQA
*STCDEF                1   1   0   0   38285 2459K 5704 17126
%MVS                  3   5  25   8   99592  14M  418K  11M
%REMAIN              7   0  23   0   222K  308K  390K 20376
CSF           S STCDEF  0113  2.2D  1   0   0   0   32360 45832  3960  1976
RRS           S STCDEF  0093  2.2D  0   1   0   0     272 1544K    0  2552
    
```

Figure 95. STORC Report

There is no graphic version of this report available.

### Field descriptions

Table 66. Fields in the STORC Report - System Information Section

Field Heading	Meaning
IPL Definitions	The amount of common storage (CSA, ECSA, SQA, and ESQA) defined to the system at IPL.  For the definition of CSA/SQA amounts and the description of the IEASYSxx member, see <i>z/OS MVS Initialization and Tuning Reference</i> .
Peak Allocation Values	The peak common storage (CSA, ECSA, SQA, and ESQA) values since IPL.  The peak allocation values include common storage used by: <ul style="list-style-type: none"> <li>Active jobs</li> <li>System activity not related to a specific job</li> <li>Ended jobs that did not release all common storage</li> </ul> The percent values are calculated by dividing the peak allocation amount by the corresponding common storage IPL value.
Average CSA to SQA Conversion	The amount of CSA or ECSA converted to SQA or ESQA, respectively.  The percent value of CSA is calculated by dividing the amount of CSA converted to SQA by the amount of CSA defined at IPL.  The percent value of ECSA is calculated by dividing the amount of ECSA converted to ESQA by the amount of ECSA defined at IPL.  The percent and amount values for SQA and ESQA are blank.
Average Use Summary	The average common storage (CSA, ECSA, SQA, and ESQA) usage during the specified range.  The average use values include common storage used by: <ul style="list-style-type: none"> <li>Active jobs</li> <li>System activity not related to a specific job</li> <li>Ended job that did not release all common storage</li> </ul>
Available at End of Range	The amount and percentage of common storage (CSA, ECSA, SQA, and ESQA) available for allocation at the end of the specified range.



Table 66. Fields in the STORC Report - System Information Section (continued)

Field Heading	Meaning
Unalloc Common Area	The amount of unallocated common area below 16M (CSA + SQA) available at the end of the specified range.

Table 67. Fields in the STORC Report - Job Information Section

Field Heading	Meaning
Jobname	Name of a job or job group (*SYSTEM, *TSO, *BATCH, *STC, *ASCH, or *OMVS) using common storage. The name can also be: <b>%MVS</b> Summary information about common storage being requested with the GETMAIN parameter OWNER(SYSTEM). <b>%REMAIN</b> Summary information about common storage that was not released by ended jobs. <b>*srvcls</b> Summary information about this service class ('srvcls' is replaced by the name of the service class).  The %MVS and %REMAIN summary lines cannot be excluded from this report.  Cursor-sensitive control is only active on the '%REMAIN' Name field. If you select %REMAIN, the STORCR (Common Storage Remaining) report will be shown.
Act	Job status at end of the specified report interval. <b>N</b> A job ended during the report interval and released all common storage <b>H</b> A job ended during the report interval but is still holding some common storage  If the field is blank, the job was still active.  Cursor-sensitive control is only active on the 'H' ACT field. If you select <b>H</b> , the STORCR (Common Storage Remaining) report will be shown with the selected job reported first. <b>Note:</b> If VSM tracking is stopped during the report interval, the job status for all jobs, even those still holding common storage, will be N.
C	A one-character abbreviation for the job class as follows: <b>A</b> ASCH <b>B</b> Batch <b>O</b> OMVS <b>S</b> Started task <b>T</b> TSO <b>?</b> Data is missing or invalid.
Service Class	The name of the service class that a specified job has been running in.
ASID	The unique system-assigned identifier for the address space in which the job is running.
ELAP time	The time elapsed between a job's starting time and end of range time.  If a job's ELAP time is greater than 999.9 days, RMF displays asterisks (**.*D) in this field instead of the elapsed time.  The field is blank when the start time of a job is zero. For example, the field is blank for batch initiators. <b>Note:</b> The ELAP time value in the STORC report is not the same as the TET (transaction elapsed time) exception value in the WFEX report. The ELAP time is the time range from the start of the address space. TET is the transaction elapsed time for the last active transaction.
Percent Used	The average percentage of common storage (CSA, ECSA, SQA, and ESQA) used by a job during the specified report interval.  If data gathering was stopped between IPL and end of range, the field is reported in dark blue.
Amount Used	The average amount of common storage (CSA, ECSA, SQA, and ESQA) used by a job during the specified report interval (specified as bytes). <b>Note:</b> You might notice a difference between the Amount Used reported in a summary line and the sum of Amount Used values for all corresponding jobs. The Amount Used value for an individual job is rounded to the nearest whole number. However, the Amount Used value for a summary line is calculated by adding the exact value for all corresponding jobs and then rounding the value.  If data gathering was stopped between IPL and end of range, the field is reported in dark blue.

### Monitor III Utility fields

You can use the Monitor III Utility to customize the STORC report. In addition to the information previously described, you can use the Utility to have the following values shown.

Table 68. Additional Fields in the STORC Report

Field Heading	Meaning
Termination date	The date the job ended.
Termination time	The time the job ended.
ID	The ID for the ended job. The ID can be used to identify the job in the system logs.

### Report options

```

RMF STORC Report Options                               Line 1 of 1
Command ==>>>                                       Scroll ==>> HALF

Change or verify parameters. To exit press END.
Changes will apply to DELAY, DEV, ENQ, HSM, JES, PROC, STOR, STORC, STORF,
and XCF.
Class          ==>>> ALL          Classes: ALL TSO BATCH STC ASCH OMVS
Service class  ==>>> *ALL        *ALL or one of available service classes below
Summary        ==>>> NO          Class summary lines on STORC Report (YES NO)
Threshold      ==>>> 0           Minimum common storage use value (%) to
                                include job in STORC report
Jobs           ==>>> NO          View job selection/exclusion panel next (YES NO)

                                Available Service Classes
APPPRIME  NRPRIME  OMVS  TSOPRIME  SYSTEM  SYSSTC
    
```

Figure 96. STORC Report Options Panel

You can specify a threshold for common storage usage. If a selected job's use of CSA, ECSA, SQA, or ESQA is greater than or equal to the threshold value, that job will be displayed in the report.

If you select YES for Jobs on the Report Options panel, a Job Selection/Exclusion panel is displayed. See Figure 37 on page 70 for an example.

## STORCR - Common Storage Remaining Report

The Common Storage Remaining report identifies jobs that have ended but have not released all of their allocated common storage (CSA, ECSA, SQA, and ESQA) since IPL.

The report is a snapshot of the system at the end of the specified report interval.

The jobs are sorted by descending storage percentage; that is, for each job with the maximum of the four common storage percentages, the job with the highest maximum percentage is reported first. Cursor-sensitive control is not active on the report.

### How to request this report

To request the Common Storage Remaining report, select **3** from the Primary Menu, then select **11** from the Resource Report Selection Menu (shown in Figure 8 on page 27) or enter following command:

```
STORCR
```

## Contents of the report

```

RMF V2R2 Common Storage Remaining Line 1 of 30
Command ==> Scroll ==> CSR
Samples: 119 System: SYSF Date: 09/28/16 Time: 17.52.00 Range: 120 Sec

Jobname ID Job Ended Date Time Amount of Common Storage
Not Released at End of Job
CSA ECSA SQA ESQA
%REMAIN 0 1109K 1504 23952
CANFDSST STC04642 09/28/16 15.59.39 0 513K 96 128
CANFDSST STC04661 09/28/16 04.40.09 0 506K 0 128
CATALOG 09/28/16 15.01.02 0 0 768 8552
CATALOG 09/28/16 10.01.03 0 0 512 6624
IRRDP TAB STC04335 09/28/16 13.36.06 0 86632 0 0
CATALOG 09/28/16 13.35.32 0 0 128 1360
STARTMVS 09/28/16 13.47.17 0 0 0 5336
SMFDUMP STC04334 09/28/16 13.36.36 0 2416 0 0
BENK STC04637 09/28/16 16.15.44 0 2304 0 0
TAGE TSU04619 09/28/16 05.04.15 0 1024 0 0
BPXAS STC04881 09/28/16 09.33.13 0 0 0 96
BPXAS STC04865 09/28/16 07.23.36 0 0 0 96
    
```

Figure 97. STORCR Report

There is no graphic version of this report available.

**Note:** If you request the STORCR report and VSM tracking was stopped between IPL and end of range, an empty STORCR report can be displayed with the message 'No ended jobs found'.

### Field descriptions

Table 69. Fields in the STORCR Report

Field Heading	Meaning
Jobname	Name of the job that ended but did not release all of its common storage (CSA, ECSA, SQA, or ESQA).  The %REMAIN summary line is always the first reported line and cannot be excluded from the report.
ID	The ID for the ended job.  The ID can be used to identify the job in the system logs. The ID field is blank for some system-related address spaces.
Job Ended Date	The date the job ended.  You can use the language options panel to customize the date format.
Job Ended Time	The time the job ended.  You can use the language options panel to customize the time format.
Amount of Common Storage Not Released at End of Job	The amount of allocated common storage (CSA, ECSA, SQA, and ESQA) that was not released when the job ended.  If data gathering was stopped between IPL and end of range, the field can be reported in dark blue.

There are no report options to specify for the STORCR report.

## STORF - Storage Frames Report

The Storage Frames report contains detailed frame counts, auxiliary slot count, and page-in rate for each address space.

### How to request this report

To request the Storage Frames report, select **3** from the Primary Menu, then select **7** from the Resource Report Selection Menu (shown in Figure 8 on page 27) or enter the following command:

```
STORF [job_class,service_class]
```

### Contents of the report

```

RMF V2R2 Storage Frames                               Line 1 of 103
Command ==>>>                                       Scroll ==>> CSR

Samples: 9      System: SYSF  Date: 09/28/16  Time: 08.44.00  Range: 30  Sec

Jobname  Service  -- Frame Occup.-- - Active Frames - AUX  PGIN
C Class  Cr  TOTAL  ACTV  IDLE  WSET  FIXED  DIV  SLOTS  RATE

INIT     S  SYSSTC   197K   0  197K   0  197K   0   53   0
IOSAS    S  SYSTEM   133K  133K   0  133K  755   0   49   0
THR64GMD B  JESLOW5   115K  115K   0  115K  115K  0   43   0
STGTHR32 B  JESLOW5  77390 77390   0  77390  367   0  62388  0
STGTHR33 B  JESLOW5  77390 77390   0  77390  367   0  59546  0
STGTHR31 B  JESLOW5  77390 77390   0  77390  367   0  60125  0
STGTHR34 B  JESLOW5  77389 77389   0  77389  367   0  64878  0
INIT     S  SYSSTC   66778   0  66778   0  66627  0   24   0
THR64FMC B  JESLOW5  65927 65927   0  65927 65720  0   60   0
THR64AMB B  JESLOW5  65925 65925   0  65925 65720  0   62   0
THR64EMC B  JESLOW5  65925 65925   0  65925 65720  0   62   0
THR64DMC B  JESLOW5  65925 65925   0  65925 65720  0   62   0
THR64BMC B  JESLOW5  65847 65847   0  65847 65612  0   18   0
THR64CMC B  JESLOW5  65846 65846   0  65846 65612  0   18   0
    
```

Figure 98. STORF Report

There is no graphic version of this report available.

### Field descriptions

Table 70. Fields in the STORF Report

Field Heading	Meaning
Jobname	Name of a job that is delayed for storage. The STORF delay report does not summarize data by job groups; all jobs within a job group are reported individually.
C	A one-character abbreviation for the job class as follows: <b>A</b> ASCH <b>B</b> Batch <b>O</b> OMVS <b>S</b> Started task <b>T</b> TSO <b>?</b> Data is missing or invalid.
Service Class	The name of the service class that a specified job has been running in.
Cr	An <b>S</b> in this column indicates that WLM managed the address space as <i>storage critical</i> during the reporting interval.

Table 70. Fields in the STORF Report (continued)

Field Heading	Meaning
Frame Occup.	<p>This field shows the frame occupancy divided into three categories.</p> <p><b>TOTAL</b> The sum of the ACTV and IDLE frames. The shared page counts are not included in TOTAL.</p> <p><b>ACTV</b> The average number of frames held by the job while it was active. This value represents the average number of active central storage frames the job used during the report interval.</p> $\text{ACTV} = \frac{\sum \text{Central Storage Frames}}{\# \text{ Samples}}$ <p><b>IDLE</b> The average number of frames held by the job while it was idle. This value represents the average number of central storage frames the jobs used when it was idle during the report interval.</p> $\text{IDLE} = \frac{\sum \text{Central Storage Idle Frames}}{\# \text{ Samples}}$
Active Frames	<p>This field is broken into three categories.</p> <p><b>WSET</b> See Table 64 on page 155 for a description of this value. The shared page counts are not included in WSET.</p> <p><b>FIXED</b> The average number of fixed frames the job was using during the report interval including frames both above and below the 16 megabyte line. While a user is swapped in, it is the number of fixed frames being used. While a user is swapped out, it is the number of fixed frames that will be used when the user is swapped back in.</p> $\text{FIXED} = \frac{\sum \text{Fixed Frames}}{\# \text{ Samples}}$ <p><b>DIV</b> The number of central storage frames used by DIV. This count is accumulated only for jobs, not for service classes. A service class displayed on the report may not be the same as where the job was running when the DIV sample was taken.</p> $\text{DIV} = \frac{\sum \text{DIV Frames}}{\# \text{ DIV Samples}}$
AUX SLOTS	Number of auxiliary slots for each address space.
PGIN RATE	<p>The average number of page-ins per second for an address space.</p> $\text{PGIN RATE} = \frac{\sum \text{Page-in Counts for Group}}{\text{Resident Time}}$ <p>The calculation is the total number of non-swap page-ins (including VIO page-ins, hiperspace page-ins, shared page group page-ins, and page-ins caused by page faults) during the range period divided by the total time an address space was swapped-in (resident time) in seconds.</p>

### Report options

The STORF Report Options panel is similar to the DEV Report Options panel. See Figure 39 on page 72 for an example. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed. See Figure 37 on page 70 for an example.

## STORM - Storage Memory Objects Report

This report provides information about the use of memory objects within the system. A memory object is a contiguous range of virtual addresses that is allocated by jobs in units of megabytes on a megabyte boundary.

The *System Summary* section of the report provides overall system information about memory objects. The bottom section provides job-related information about the use of memory objects for jobs active during the specified report interval.

### How to request this report

To request the Storage Memory Objects report, select a **3** from the Primary Menu, then select **7A** from the Resource Report Selection Menu (shown in Figure 8 on page 27) or enter the following command:

```
STORM [job_class,service_class]
```

### Contents of the report

```

RMF V2R2 Storage Memory Objects Line 1 of 12
Command ==> Scroll ==> CSR
Samples: 60 System: TRX1 Date: 09/28/16 Time: 10.38.00 Range: 60 Sec

----- System Summary -----
--MemObj-- --Frames-- -1MB MemObj- --1MB Fixed-- -1MB Pageable-
Shared 4 Shared 177K Fixed 3 Total 200 Initial 560
Common 97 Common 37173 Shared 2 Common 30 Dynamic 129
                %Used 14.2 Common 1 %Used 100 %Used 100
-----
Service      ---- Memory Objects --- -1MB Frames- ----- Bytes -----
Jobname C Class ASID Total Comm Shr 1 MB Fixed Pgable Total Comm Shr
LOAD1MP B DISCRETN 0069 71 0 0 0 0 0 0 2538M 0 0
LOAD2GNP B DISCRETN 0070 71 0 0 0 0 0 0 4073M 0 0
LOAS1MP B DISCRETN 0072 70 0 1 0 0 4 2537M 0 1024K
JL1LE1MS B DISCRETN 0055 10 0 2 0 0 0 4109M 0 4096M
PGBLE1MP B DISCRETN 0056 10 0 0 0 0 0 21.0M 0 0
STOR641M B DISCRETN 0058 10 0 0 2 40 0 53.0M 0 0
STOR644K B DISCRETN 0059 10 0 0 0 0 0 113M 0 0
    
```

Figure 99. STORM Report

There is no graphic version of this report available.

### Field descriptions

Table 71. Fields in the STORM Report - System Summary Section

Field Heading	Meaning
MemObj	<p><b>Shared</b> Average number of memory objects allocated in the high virtual shared storage of the system.</p> <p><b>Common</b> Average total number of memory objects allocated in the high virtual common storage of the system. This value includes the memory objects that cannot be attributed to an address space.</p>
Frames	<p><b>Shared</b> Average number of high virtual shared storage pages backed in central storage (in units of 4 KB).</p> <p><b>Common</b> Average number of high virtual common storage pages backed in central storage (in units of 4 KB).</p> <p><b>%Used</b> Percentage of high virtual common storage used by the system.</p>

Table 71. Fields in the STORM Report - System Summary Section (continued)

Field Heading	Meaning
1MB MemObj	<p><b>Fixed</b> Average number of fixed memory objects that are allocated in the system and can be backed in 1 MB frames.</p> <p><b>Shared</b> Average number of memory objects that are allocated in high virtual shared storage and can be backed in 1 MB frames.</p> <p><b>Common</b> Average number of fixed memory objects that are allocated in high virtual common storage and can be backed in 1 MB frames. This value includes the memory objects that cannot be attributed to an address space.</p> <p>These fields are only available with Enhanced DAT Architecture.</p>
1MB Fixed	<p><b>Total</b> Total number of 1 MB frames that can be used by fixed memory objects. This value is equal to the size of the Large Frame Area in megabytes.</p> <p><b>Common</b> Average number of 1 MB high virtual common memory pages fixed in central storage. This value includes the pages that cannot be attributed to an address space or have not been freed during address space termination.</p> <p><b>%Used</b> Percentage of 1 MB frames that are used by fixed memory objects in the Large Frame Area, regardless of whether the frames are actually used for 1 MB pages or used to satisfy 4 KB space requests on a constrained system.</p> <p>These fields are only available with Enhanced DAT Architecture.</p>
1MB Pgable	<p><b>Initial</b> Number of 1 MB frames that can be used by pageable and DREF memory objects. This value is calculated by the system at IPL time.</p> <p><b>Dynamic</b> Average number of 1 MB frames in the LFAREA that were used to satisfy 1 MB pageable page requests.</p> <p><b>%Used</b> Percentage of 1 MB frames that are used by pageable and DREF memory objects, regardless of whether the frames are actually used for 1 MB pages or used to satisfy 4K space requests on a constrained system.</p> <p>These fields are only available with Enhanced DAT Architecture.</p>

Table 72. Fields in the STORM Report - Address Space Section

Field Heading	Meaning
Jobname	The name of a job using memory objects
C	<p>A one-character abbreviation for the job class as follows:</p> <p><b>A</b> ASCH</p> <p><b>B</b> Batch</p> <p><b>O</b> OMVS</p> <p><b>S</b> Started task</p> <p><b>T</b> TSO</p> <p><b>?</b> Data is missing or invalid.</p>
Service Class	The name of the service class that a specified job has been running in.
ASID	The decimal identifier of the address space in which the job is running.
Memory Objects	<p><b>Total</b> Average number of memory objects allocated by this address space.</p> <p><b>Comm</b> Average number of high virtual common memory objects allocated by this address space.</p> <p><b>Shr</b> Average number of high virtual shared memory objects allocated by this address space.</p> <p><b>1 MB</b> Average number of fixed memory objects allocated with this address space as the owner that can be backed in 1 MB frames. This field is only available with Enhanced DAT Architecture.</p>

## Mon III - STORM

Table 72. Fields in the STORM Report - Address Space Section (continued)

Field Heading	Meaning
1 MB Frames	<p><b>Fixed</b> Average number of 1 MB frames in the Large Frame Area owned by this address space. Frames that are used to satisfy 4 KB space requests on a constrained system are not included.</p> <p><b>Pgable</b> Average number of 1 MB frames that are used by pageable and DREF memory objects owned by this address space. Pageable memory objects that have been fixed after allocation, are also included. Frames that are either used by common 1 MB pages or to satisfy 4 KB space requests on a constrained system are not included.</p> <p>These fields are only available with Enhanced DAT Architecture.</p>
Bytes	<p><b>Total</b> Average amount of storage allocated from high virtual memory in memory objects with this address space as the owner.</p> <p><b>Comm</b> Average amount of high virtual common storage allocated with this address space as the owner.</p> <p><b>Shr</b> Average amount of high virtual shared storage allocated with this address space as the owner.</p>

### Report options

The STORM Report Options panel is similar to the Device Report Options panel. See Figure 39 on page 72 for an example. If you select YES for Jobs on the STORM Report Options panel, the Job Selection/Exclusion panel is displayed. See Figure 37 on page 70 for an example.

---

## STORR - Storage Resource Delays Report

The Storage Resource Delays report (STORR) provides information about storage problems and paging space delay by volume serial.

### How to request this report

To request the Storage Resource Delays report, select **3** from the Primary Menu, then select **8** from the Resource Report Selection Menu (shown in Figure 8 on page 27) or enter the following command:

```
STORR
```



## Contents of the report

```

RMF V2R2 Storage Resource Delays Line 1 of 7
Command ==> Scroll ==> PAGE

Samples: 9 System: SYSF Date: 09/28/16 Time: 08.44.00 Range: 30 Sec

----- Central Storage Summary -----
----- % Frames ----- Frames System
NUC SQA CSA LPA ACTV IDLE AVAIL SHR Online UIC
  0  3  0  0  38  14  44  1  6291706  65534

Total SQA + ESQA Overflow: 28K
----- Page/Swap Activity -----
Volume DEV CU ACT CON DSC PND Pend SPACE - AVG Active Users-
Serial Type Type PAV % % % % Reasons TYPE TOTL LOCL SWAP COMM

S53PG7 33903 2105 4 0 0 0 0 None LOCL 0.0 0.0 0.0 0.0
S53PG8 33903 2105 4 0 0 0 0 None LOCL 0.0 0.0 0.0 0.0
S53PG9 33903 2105 4 0 0 0 0 None LOCL 0.0 0.0 0.0 0.0
S53PG3 33903 2105 4 0 0 0 0 None LOCL 0.0 0.0 0.0 0.0
S53PG2 33903 2105 6 0 0 0 0 None LOCL 0.0 0.0 0.0 0.0
S53PG1 33903 2105 9 0 0 0 0 None COMM 0.0 0.0 0.0 0.0

```

Figure 100. STORR Report

The report has two sections.

### Central Storage Summary

This section includes general information about the use of central storage.

### Page/Swap Activity

This section includes information about page/swap activity and paging delays.

If RMF cannot provide data in the Page/Swap Activity section for ACT %, CON %, DSC %, and PND %, dashes appear in these fields.

If RMF is unable to obtain valid hardware data for a sub-channel, it prints dashes (---) instead of DLY DB% and DLY CU%.

The graphic form of this report shows the average number of active users connected (CON), disconnected (DSC), pending (PND), and delayed for LOCL, SWAP, and COMM.

### Field descriptions

Table 73. Fields in the STORR Report - Central Storage Summary Section

Field Heading	Meaning
% Frames	<p>The percentage of storage being used for NUC, SQA, CSA, LPA, ACTV, IDLE, and available. All percentages are based on the total number of online central storage frames during the report interval.</p> <p>The categories are as follows:</p> <p><b>NUC</b> Percentage of central storage frames allocated to the nucleus (NUC).  <b>SQA</b> Percentage of central storage frames allocated to the system queue area (SQA).  <b>CSA</b> Percentage of central storage frames allocated to the common storage area (CSA).  <b>LPA</b> Percentage of central storage frames allocated to the link pack area (LPA).  <b>ACTV</b> Percentage of private frames allocated to jobs that are active. This value represents the number of central storage frames allocated to all active address spaces. It includes idle, using, and unknown time.  <b>IDLE</b> Percentage of private frames allocated to jobs that are idle. This value represents the number of central storage frames allocated to all idle address spaces.  <b>AVAIL</b> Percentage of available central frames.  <b>SHR</b> Percentage of shared frames in central storage.</p>
Frames Online	The number of central storage frames, excluding read-only frames. Nucleus frames are included in this value.
System UIC	The system's unreferenced interval count indicates storage contention.
Total SQA + ESQA Overflow	<p>The amount of CSA and ECSA storage used to hold SQA and ESQA data when SQA and ESQA are full. If there is no overflow, the field heading for Total SQA + ESQA overflow will not appear.</p> $\text{Total SQA + ESQA Overflow} = \frac{\text{Total Overflow}}{\# \text{ Samples}}$ <p><b>Total overflow</b>                      Total overflow above and below 16M</p>

Table 74. Fields in the STORR Report - Page/Swap Activity Section

Field Heading	Meaning
Volume Serial	Name the volume that contains a page data set.
DEV Type	The device type.
CU Type	The control unit type.
PAV	<p>The number of parallel access volumes (base and alias) which were available at the end of the report interval. If the number has been changed during the report interval, it is followed by an '*'.                      If the device is a HyperPAV base device, the number is followed by an 'H'. The value is the average number of HyperPAV volumes (base and alias) for that interval.</p> $\text{Average \# of HPAV devices} = \frac{\text{Accumulated \# of HPAV devices}}{\text{Number of Samples}}$ <p>This field appears only for parallel access volumes.</p>

Table 74. Fields in the STORR Report - Page/Swap Activity Section (continued)

Field Heading	Meaning
ACT %	<p>The percentage of time during the report interval when the device was active. To derive this value, RMF computes the accumulated active time as follows:            Active Time = PND Time + CON Time + DSC Time</p> <p><b>PND Time</b>            The time all I/O requests wait in the logical control unit (CU-HDR) queue before there is an available path. Pending time includes the time spent waiting for a channel, control unit, or head of string, or for the actual device (if it is a shared device that is reserved by another processor).</p> <p><b>CON Time</b>            The time the device was connected to a channel path to actually transfer data between the device and storage.</p> <p><b>DSC Time</b>            The time the device has an active channel program and is disconnected (not transferring data). Disconnect time includes seek time, normal rotation delay time, and extra rotation delay because the channel was busy when the device needed to reconnect.</p>
CON %	<p>The percent connect time:</p> $\text{CON \%} = \frac{\text{Connect Time}}{\text{Range Time}} * 100$
DSC %	<p>The percent disconnect time:</p> $\text{DSC \%} = \frac{\text{Disconnect Time}}{\text{Range Time}} * 100$
PND %	<p>The percent pending time:</p> $\text{PND \%} = \frac{\text{Pending Time}}{\text{Range Time}} * 100$ <p><b>Note:</b></p> <ol style="list-style-type: none"> <li>The CON %, DSC %, and PND % values sum to the ACT % value. DB and CU delay are a subset of pending time and sum to PND % or less.</li> <li>IOS updates the data fields used to calculate CON %, DSC %, and PND % when the I/O operation completes. Therefore, some of the time from the previous report interval might be included in these values, while some of the time in the current range period might be absent from these values. This discrepancy is noticeable on paging devices because they have very long channel programs.</li> </ol>
Pend Reasons	<p>The reason for the delay and the percentage of delay.</p> <p><b>DB</b> Device busy delay, which is the percentage of time during the report interval when the channel subsystem measured I/O request delay because the device was busy. Device busy might mean that the volume is in use by another system, the device is reserved by another system, a head of string busy condition caused the contention, or some combination of these conditions has occurred.</p> $\text{DLY DB\%} = \frac{\text{Accumulated DB Delay Time}}{\text{Range Time}} * 100$ <p><b>CMR</b> Command response time delay, which is the percentage of time during the report interval, when the first command of an I/O instruction of the channel program is sent to the device, until the device indicates it has accepted the command.</p> $\text{DLY CMR\%} = \frac{\text{Accumulated Command Response Delay Time}}{\text{Range Time}} * 100$ <p><b>Note:</b> If either hardware data or volume related percentages are not available, this field is blank.</p>

## Mon III - STORR

Table 74. Fields in the STORR Report - Page/Swap Activity Section (continued)

Field Heading	Meaning
SPACE TYPE	The space types for which the percentage of the volume's delay is reported. The types appear in the following order: <b>LOCL</b> User private area <b>COMM</b> Common area <b>PLPA</b> Pageable link pack area
AVG Active Users	The average number of jobs waiting for the data set. $\text{AVG Active Users} = \frac{\sum \text{Waiting Jobs}}{\# \text{ Samples}}$ <p><b>Sum of all waiting jobs</b> Sum of all delay samples for all jobs waiting for the data set.</p> <p>This category is divided into:</p> <p><b>TOTL</b> The percentage COMM, LOCL, and SWAP contribute to the overall delay according to the SPACE TYPE specified. The percentages for all these resources add up to DLY % if there is no overlap of the delay states; if there is overlap, the percentages add up to more than DLY %.</p> <p><b>LOCL</b> The percentage that local (private) storage paging contributes to the delay from the time of the page fault until I/O is completed.</p> <p><b>SWAP</b> The percentage that swapping contributes to the delay from the time of swap initiation until the last swap page I/O is completed.</p> <p><b>COMM</b> The percentage that common (CSA or LPA) storage paging contributes to the delay from the time of the page fault until I/O is completed.</p> <p>For LOCL, SWAP, and COMM, RMF scans all ASM AIA chains. If the address space has one or more incomplete page input requests, RMF updates the counter in the appropriate category (LOCL, SWAP, or COMM) once per sample.</p>

### Monitor III Utility fields

You can use the Monitor III Utility to customize the STORR report. In addition to the delays previously described, you can use the Utility to have the following delay percentages shown.

Table 75. Additional Fields in the STORR Report

Field Heading	Meaning
Percentage of using	The percentage of time the device was found being used by an address space.
Percentage of DLY-DB	The percentage of time during the report interval when the channel subsystem measured I/O request delay because the device was busy.
Percentage of DLY-CUB	The percentage of time during the report interval when there is I/O request delay because the control unit was busy.
Percentage of DLY-DPB	The percentage of time during the report interval when there is I/O request delay because the ES/Connection Director port was busy.
Delay reason percentage	The percentage of time the device was delayed.

## STORS - Storage Delay Summary Report

This Storage Delay Summary (STORS) report provides you with an overview of storage usage by service classes, report classes, and workload groups.

### How to request this report

To request the Storage Delay Summary report, select **3** from the Primary Menu, then select **9** from the Resource Report Selection Menu (shown in Figure 8 on page 27) or enter the following command:

STORS workload\_group | service\_class | report\_class

## Contents of the report

```

RMF V2R2 Storage Delay Summary                               Line 1 of 19
Command ==>                                                Scroll ==> CSR

Samples: 119      System: SYSF Date: 09/28/16 Time: 10.52.00 Range: 120 Sec

----- Central Storage Summary -----
----- % Frames ----- Frames System
NUC  SQA  CSA  LPA  ACTV  IDLE  AVAIL SHR   Online  UIC
   0   3   0   0   38   14   44   1  6291706  65534

Group  T  -- Users --  - Average Number Delayed For-  - Average Frames-  PGIN
          TOTL ACTV  ANY COMM LOCL SWAP OTR  OTHR  ACTV  IDLE FIXED  RATE

BATCH  W    4    0    0    0    0    0    0    0    11325    0    352    0.0
BTCHDEF S    3    0    0    0    0    0    0    0    10998    0    298    0.0
OMVSKERN S    1    0    0    0    0    0    0    0    327    0    54    0.0
OMVS   W    2    0    0    0    0    0    0    13746    0    304    0.0
OE     S    2    0    0    0    0    0    0    13746    0    304    0.0
STC   W   16    0    0    0    0    0    0    19207   373  1153    0.0
GPMSEVE S    1    0    0    0    0    0    0    8666    0    110    0.0
STCDEF S   15    0    0    0    0    0    0   10541   373  1043    0.0
    
```

Figure 101. STORS Report

The top section on the report provides overall system information and is the same as the Central Storage Summary section of the STORR report. The bottom section of the report provides summary lines for service classes, report classes, and workload groups.

A graphic report shows the average number of users delayed for COMM, LOCL, SWAP, OTR, and OTHR.

### Field descriptions

Table 76. Fields in the STORS Report

Field Heading	Meaning
Central Storage Summary	Fields in this section are described in Table 73 on page 168.
Group	The name of the group, including: <ul style="list-style-type: none"> <li>• Workload group names</li> <li>• Service class names</li> <li>• Report class names</li> </ul>
T	A one-character abbreviation for the type of workload manager group as follows: <ul style="list-style-type: none"> <li><b>W</b> Workload group name</li> <li><b>S</b> Service class name</li> <li><b>R</b> Report class name</li> <li><b>n</b> Service/report class period</li> </ul>
Users	The number of users within the group. This category includes the following headings: <ul style="list-style-type: none"> <li><b>TOTL</b> The total number of users equals the number of different users found in all address spaces for the group listed during the report interval.</li> <li><b>ACTV</b> The average number of active users is a measure of system workload.</li> </ul> See "WFEX - Workflow/Exceptions Report" on page 210 for the definition of User/Active.

## Mon III - STORS

Table 76. Fields in the STORS Report (continued)

Field Heading	Meaning
Average Number Delayed For	<p>The average number of delayed users is summarized for the following categories:</p> <p><b>ANY</b> Delay the group experienced because of contention for any of the following measured storage reasons during the report interval.</p> <p><b>COMM — LOCL — SWAP — OTR — OTHR</b> For descriptions of these delays, see the corresponding field in the STOR report (Table 64 on page 155).</p>
Average Frames	<p>The average number of storage frames the group held during the report interval. This field reports on the following frame categories:</p> <p><b>ACTV — IDLE</b> See Table 70 on page 162 for a description of these counts.</p> <p><b>FIXED</b> The average number of fixed frames the job was using during the report interval including frames both above and below the 16 megabyte line.</p> $\text{Avg Fixed Frames} = \frac{\sum \text{Fixed Frames}}{\# \text{ Samples}}$
PGIN RATE	<p>The rate at which pages are being read into central storage.</p> $\text{PGIN RATE} = \frac{\sum \text{All Page-in Counts for Group}}{\text{Resident Time}}$ <p>The address-space related shared storage page-ins are included in the PGIN RATE.</p>

### Monitor III Utility fields

You can use the Monitor III Utility to customize the Storage Delay Summary report. In addition to the delays previously described, you can use the Utility to have the delays in Table 77 shown in the Storage Delay Summary report.

Table 77. Additional Fields in the STORS Report

Field Heading	Meaning
Average number delayed for VIO	The average number of delayed users due to virtual I/O.
Delayed for XMEM	The average number of users delayed due to cross memory address space services.
Delayed for HIPR	The average number of users delayed due to standard hiperspace services (including waits during scroll wait, but not ESO hiperspaces).

## Report options

```

RMF STORS Report Options                               Line 1 of 4
Command ==>                                           Scroll ==> HALF

Select (S), exclude (X) or fill-in groups for the STORS report. Press END.
Selections made here also affect the System Information (SYSINFO) and
the Sysplex Summary (SYSSUM) report.

Service class ==> YES  Service class and period lines on the report (YES NO)
Report class  ==> YES  Report class and period lines on the report (YES NO)
Period        ==> YES  Active periods for listed service classes (YES NO)

Sel Group  T      Sel Group  T      Sel Group  T      Sel Group  T
-  - - - - -      -  - - - - -      -  - - - - -      -  - - - - -
S  *ALL           PRIMEAPP W      PRIMEBAT W      PRIMETSO W
   PRIMOMVS W    SYSTEM W      APPPRIME S      HOTPRIME S
   NRPRIME S     OMVS S        OMVSKERN S      TSOPRIME S
   WLMPRIME S    SYSTEM S      SYSSTC S        SYSOTHER S
    
```

Figure 102. STORS Report Options Panel

The STORS report, the SYSINFO report, and the SYSSUM report use similar Report Options panels. Selections made for service classes, report classes, or workload groups on either options panel affect all reports.

### Service class

If you enter YES for Service Class, all service classes and service class periods (if you also specified YES for Period) are displayed below each workload group. Otherwise, no service classes are shown.

You can also specify any of the available service classes listed in the scrollable section at the bottom of this panel.

If the service class you want is not listed, it was not active during the current report interval. If you specify the service class, it will appear on the report when it is available.

### Report class

If you enter YES for Report Class, all report classes and report class periods (if you also specified YES for Period) are displayed. Otherwise, no report classes are shown.

**Period** Enter YES for Period to have all periods displayed below each class entry on the report.

Enter NO to have only the service or report class entries displayed on the report.

**Sel** Allows you to select or exclude specific classes on your STORS report.

**Group** The columns headed by Group include all the service class names, workload group names, and report class names currently in the system and any names that you have previously selected, whether or not they are currently in the system.

To request a report for several groups with similar names, use an asterisk (\*) as a "wild card" character. For example, to request a report for all groups starting with A, specify 's' under Sel, 'a\*' under Group and ensure that there is an 'x' beside \*ALL.

You can also specify multiple wild card entries, for example, to list all service classes starting with CICS® and all service classes starting with IMS™, specify

Se1	Group	T	Se1	Group	T
S	CICS*__	S	S	IMS*__	S
<hr/>					
X	*ALL	-			-

You can use the wild card to select by type, for example, to list service classes only, specify:

Se1	Group	T	Se1	Group	T
S	*_____	S		_____	-
<hr/>					
X	*ALL	-			-

**T - type**

Type can be:

- W** Workload group name
- S** Service class name
- R** Report class name

---

## SYSENQ - Sysplex Enqueue Delays Report

The SYSENQ report is similar to the ENQR report (see “ENQR - Enqueue Resource Delays Report” on page 99), but the information presents contentions for serially reusable resources in the sysplex. This can help in understanding bottlenecks in the sysplex not being caused by the current system.

**Note:** The report shows sysplex-wide enqueue delays only, you find all other enqueue delays in the ENQR report.

### How to request this report

To request the SYSENQ report, select **S** on the Primary Menu, and then select **4** on the Sysplex Report menu (shown in Figure 5 on page 24), or enter the following command:

SYSENQ



## Contents of the report

```

RMF V2R2 Sysplex ENQ Delays - RMFPLEX1 Line 1 of 10
Command ==> Scroll ==> HALF

Samples: 114 Systems: 3 Date: 09/28/16 Time: 12.58.30 Range: 100 Sec

----- Resource Name -----
Major/Minor % Jobname Sys-Name ST % Jobname Sys-Name ST
IGDCDSXS 99 SMS SYS1 EW 99 SMS SYS4 EO
SYS1.SMS.COMMDS
SYSVSAM 99 SMF SYS1 EW 99 FPB SYS1 EO
SYS1.SYS1.MAN3.DATASYS1.CA
SYSZMCS 99 CONSOLE SYS1 EW 99 ALC SYS1 EO
SYSZMCS#CL2
SYSZMCS 20 CONSOLE SYS1 EW 18 CATALOG SYS1 EO
SYSZMCS#CL1 2 CATALOG SYS1 EO
DSNJBSDS 99 S412MSTR SYS1 EW 99 S411MSTR SYS1 EO
MODIFY
    
```

Figure 103. SYSENG Report

The graphic form of this report shows the average number of active users waiting for each resource.

### Field descriptions

Table 78. Fields in the SYSENG Report

Field Heading	Meaning
Resource Name	The Major name and Minor name of the resource delaying the job. The major name is listed above the minor name. The major name is up to eight characters long and the minor name is up to 36 characters long. If the minor name contains unprintable characters, it will be up to 18 characters long (represented by 36 hexadecimal digits). If the minor name is longer than 26 characters, RMF only displays the first 26 characters. If there are two resources with the same major name and their minor names differ only after the first 36 characters, then RMF considers them as the same resource.
-- Delayed -- %	The delay percentage of the job for a specific enqueued resource.  $\text{Delayed \%} = \frac{\text{\# Delay Samples}}{\text{\# Samples}} * 100$ <p><b>Delay samples</b> The number of samples when the job was delayed for a specific enqueued resource.</p>
-- Delayed -- Jobname	Name of the job delayed for the resource. RMF lists all jobs delayed for the resource.  If the catalog system address space is processing a catalog request on behalf of the delayed job, the jobname of the catalog address space (usually CATALOG) will appear below the jobname preceded by a +.
-- Delayed -- Sys-Name	The z/OS system name where the job is running on.
-- Delayed -- ST	The status indicates whether the waiting job wants exclusive (EW) or shared (SW) use of the resource.

Table 78. Fields in the SYSEQ Report (continued)

Field Heading	Meaning
-- Holding -- %	<p>The percent of the range that a specific job was holding the resource while the named job was delayed.</p> $\text{Holding \%} = \frac{\text{\# Holding Samples}}{\text{\# Samples}} * 100$ <p><b>Holding samples</b> The number of samples when the holding job was holding the resource while the named job was delayed.</p> <p>Because more than one job can hold the resource at a time, these values can add up to more than 100%.</p>
-- Holding -- Jobname	<p>The name of the job that is holding the resource that the delayed job is waiting for.</p> <p>If the catalog system address space is processing a catalog request on behalf of the delayed job, the jobname of the catalog address space (usually CATALOG) will appear below the jobname preceded by a +.</p>
-- Holding -- Sys-Name	<p>The z/OS system name where the job is running on.</p>
-- Holding -- ST	<p>The status indicates whether the holding job has exclusive (EO) or shared (SO) use of the resource.</p>

## SYSINFO - System Information Report

The System Information (SYSINFO) report presents an overview of the system, its workload, the average response time for a transaction in a specific service class, report class, or workload group, and the total number of jobs using resources or delayed for resources.

### How to request this report

To request the System Information report, select **1** from the Primary Menu, then select **2** from the Overview Report menu (shown in Figure 6 on page 25) or enter the following command:

```
SYSINFO workload_group | service_class | report_class
```

## Contents of the report

```

RMF V2R2 System Information                               Line 1 of 28
Command ==>>>                                         Scroll ==>> HALF

Samples: 100      System: MVS3 Date: 09/28/16 Time: 10.03.20 Range: 100 Sec

Partition:  MVS1   9672 Model RX4           Appl%:    63 Policy: STANDARD
CPs Online:   4     Avg CPU Util%: 73       EAppl%:   65 Date: 09/28/16
AAPs Online:  0     Avg MVS Util%: 84       Appl% AAP: 0 Time: 14.05.07
IIPs Online:  0

Group  T WFL --Users-- RESP TRANS -AVG USG- -Average Number Delayed For -
      %  TOT ACT  Time  /SEC PROC DEV  PROC DEV STOR SUBS OPER ENQ

*SYSTEM  31 669 26      13.95 5.3 5.0   5.6 4.1 7.0 2.6 2.0 2.0
*TSO     50 534 8       13.95 2.6 2.1   0.4 1.5 2.0 0.8 0.0 0.0
*BATCH   26 11 10      0.00 1.5 1.4   1.4 1.7 0.5 1.8 1.0 2.0
*STC     27 115 8      0.00 1.1 1.5   0.1 1.0 4.5 0.1 1.0 0.0
*ASCH    3 0 0        0.00 0.0 0.0   0.0 0.0 0.0 0.0 0.0 0.0
*OMVS    2 0 0        0.00 0.0 0.0   0.0 0.0 0.0 0.0 0.0 0.0
*ENCLAVE 5 4 N/A      N/A 0.2 N/A   3.7 N/A 0.0 N/A N/A N/A
PRIMEBAT W 26 11 10  46.0 0.06 1.5 1.4 1.4 1.7 0.5 1.8 1.0 2.0
NRPRIME  S 26 11 10  46.0 0.06 1.5 1.4 1.4 1.7 0.5 1.8 1.0 2.0
          1 23 9 9    27.9 0.06 0.9 1.4 0.8 1.6 0.5 1.8 1.0 2.0
          2 29 0 0    54.2 0.02 0.1 0.0 0.2 0.0 0.0 0.0 0.0 0.0
          3 59 1 1    .000 0.00 0.6 0.0 0.4 0.0 0.0 0.0 0.0 0.0
PRIMTSO  W 50 527 8   .759 13.98 2.6 2.1 0.4 1.5 2.0 0.8 0.0 0.0
TSOPRIME S 50 527 8   .759 13.98 2.6 2.1 0.4 1.5 2.0 0.8 0.0 0.0
          1 48 526 8   .403 13.98 2.1 1.9 0.3 1.3 2.0 0.8 0.0 0.0
          2 75 1 1    30.6 0.08 0.3 0.2 0.1 0.1 0.0 0.0 0.0 0.0
          3 75 0 0    126 0.02 0.1 0.0 0.0 0.0 0.0 0.0 0.0 0.0
    
```

Figure 104. SYSINFO Report

The SYSINFO report has two sections. The top section provides you with an overview of the system. It identifies the measured system, the policy name, the policy activation date and time. It also includes information about processor usage during the report interval. In an LPAR environment, the header contains an extra line showing the z/OS view of CPU utilization and the LPAR partition name in which the Monitor III data gatherer is running. For the different aspects of CPU utilization refer to “CPU - CPU Activity report” on page 340.

The bottom section summarizes information about the total system (\*SYSTEM), job classes (\*TSO, \*BATCH, \*STC, \*ASCH, or \*OMVS), enclaves (\*ENCLAVE), workload groups, service classes, and report classes. The \*SYSTEM summary line represents the system total values as summarized information from all other summary lines.

**Note:** It might be possible that there is enclave activity in your system (for example, indicated by EAppl% > Appl% in the SYSINFO report), but the ENCLAVE report issues the message 'Enclave data is not currently available'. The reason is that only those enclaves are shown in the report that have been sampled at least twice and that are active or inactive at the end of the Monitor III MINTIME. Therefore, short-running enclaves will not appear in the report.

When the report interval spans more than one Monitor III MINTIME, the above criteria must match for the last MINTIME in the report interval.

The proportion of the active users in each using or delay category indicates the proportion of the average response time that is spent in that category. The graphic form of this report shows the average number of active users for each type of delay.

**Note:**

1. Report class data lines contain information for the transaction response time (RESP Time) and transaction rate (TRANS/SEC) fields. The rest of the fields are blank.
2. The transaction response time (RESP Time) field in all summary data lines is also blank.
3. There is no graphic support for report class lines.

**Field descriptions**

Table 79. Fields in the SYSINFO Report

Field Heading	Meaning
Partition	Partition name.
CPs Online	The number of general purpose processors (standard CPs) online during the range period.
AAPs Online	The number of zAAPs online during the range period. If the LOADxx PROCVIEW CORE parameter is in effect, the reported value designates the number of online threads.
IIPs Online	The number of zIIPs online during the range period. If the LOADxx PROCVIEW CORE parameter is in effect, the reported value designates the number of online threads.
Processor	Processor family and model (or N/A — if model information is not available). If the processor does not support the Concurrent Processor Upgrade function, the hexadecimal version number is displayed.
Avg CPU Util%	<p>The average utilization percentage for all general purpose processors (CPs) during the report interval (LPAR view of the CPU utilization):</p> $\text{Avg CPU Util\%} = \frac{\text{Sum of LPAR CPU Times}}{\text{Sum of Online Times}} * 100$ <p>The LPAR CPU Time for one general purpose processor is calculated depending on the status of the logical processor:</p> <p><b>Wait Completion NO</b> LPAR CPU Time = PR/SM Dispatch Time</p> <p><b>Wait Completion YES</b> LPAR CPU Time = PR/SM Dispatch Time - Wait Time</p> <p><b>Dedicated</b> LPAR CPU Time = Online Time - Wait Time</p> <p>*** indicates missing or invalid data.</p>
Avg MVS Util%	<p>z/OS view of CPU utilization which is the percentage of the time that the general purpose processors (CPs) were busy:</p> $\text{Avg MVS Util\%} = \frac{\text{Time Range} - \text{Sum of Wait Times}}{\text{Time Range}} * 100$ <p>The time range is the sum of the times the general purpose processors were online. With HiperDispatch mode active, it is the sum of the times the processors were online but not parked.</p> <p>For more information about the z/OS view of CPU utilization refer to "CPU - CPU Activity report" on page 340.</p>
Appl%	Percentage of the maximum general purpose processor capacity used by all address spaces during the report interval. This value is divided by the number of logical processors or cores that have been active during this interval.
EAppl%	Percentage of the maximum general purpose processor capacity used by all address spaces and enclaves during the report interval. This value is divided by the number of logical processors or cores that have been active during this interval.

Table 79. Fields in the SYSINFO Report (continued)

Field Heading	Meaning
Appl% AAP	Percentage of the maximum zAAP capacity used by all address spaces during the report interval. This value is divided by the number of logical zAAP processors or cores that have been active during this interval.
Appl% IIP	Percentage of the maximum zIIP capacity used by all address spaces during the report interval. This value is divided by the number of logical zIIP processors or cores that have been active during this interval.
Policy Date Time	The name and the activation date and time of the service policy in effect during collection of the reported data. This, however, does not imply that the complete policy definition is shown on this report.
Group	The name of a class (*SYSTEM, *TSO, *BATCH, *STC, *ASCH, or *OMVS), an enclave (*ENCLAVE), or a group, including: <ul style="list-style-type: none"> <li>• Workload group names</li> <li>• Service class names</li> <li>• Report class names</li> </ul>
T	Type of workload manager group: <b>W</b> Workload group name <b>S</b> Service class name <b>R</b> Report class name <b>n</b> Service/report class period
WFL %	The workflow percentage of that particular group. A value of 100% indicates no workload contention, while a value of 0% indicates that all requests for system resources are delayed.
Users	The number of users within the group. This category includes the following headings: <b>TOT</b> The total number of users. <b>ACT</b> The average number of active users.  See the definition of Users/Active under "WFEX - Workflow/Exceptions Report" on page 210 for more details.
RESP Time	The average response time (in seconds) for all transactions that ended during the report interval. The response time value is the sum of the queued time and the active time for an average ended transaction.  More than 999 seconds are shown with <ul style="list-style-type: none"> <li>• K - times one thousand seconds</li> <li>• M - times one million seconds</li> </ul> If the RESP Time field is shown in dark blue, the data reported can be statistically insignificant. This can happen if the transaction rate is low or the response time is short compared to the value of the report interval. Use the GROUP Response Time report to get additional information related to 'statistical error' for a specific group.  To increase the accuracy of your data, try increasing the value of the report interval to a value higher than or equal to the response time.
TRANS /SEC	The number of transactions per second.  When used with the number of active users in the report, this field gives you an overview of how fast the system can handle the amount of work for a given group. The number of completed transactions between cycles is accumulated for each sample.  $\text{TRANS /SEC} = \frac{\text{Completed Transaction Count}}{\text{Range Time}}$
AVG USG	The average number of users is summarized for each group. RMF takes the sum of using samples for the address space(s) associated with the group and divides by the number of samples.  The average number of users is reported for the following categories: <b>PROC</b> Average number of users using the processor during the report interval. <b>DEV</b> Average number of users using devices during the report interval.

## Mon III - SYSINFO

Table 79. Fields in the SYSINFO Report (continued)

Field Heading	Meaning
Average Number Delayed For	<p>The average number of delayed users is summarized for each group. RMF takes the sum of delay samples for the address space(s) associated with the group and divides by the number of samples in the range.</p> <p>The average number delayed for is reported for the following categories:</p> <p><b>PROC</b> Number of users experiencing delay because of contention for the processor during the report interval.</p> <p><b>DEV</b> Number of users experiencing delay because of contention for the devices during the report interval.</p> <p><b>STOR</b> Number of users experiencing delay because of contention for storage during the report interval.</p> <p><b>SUBS</b> Number of users experiencing delay because of contention for JES, HSM, or XCF during the report interval.</p> <p><b>OPER</b> Number of users experiencing delay because of a message request, a mount request, or a quiesce during the report interval. Quiesce means that the operator has quiesced the address space.</p> <p><b>ENQ</b> Number of users experiencing delay because of contention for an enqueued resource during the report interval.</p>

### Monitor III Utility fields

You can use the Monitor III Utility to customize the SYSINFO report. In addition to the information shown previously, you can use the Utility to have the following values shown.

Table 80. Additional Fields in the SYSINFO Report

Field Heading	Meaning
SYSAFVCV	The percentage of central storage frames the job used during the report interval.
SYSADJVC	The average number of users experiencing delay when requesting service from JES.
SYSADHVC	The average number of users experiencing delay when requesting service from HSM.
SYSADXVC	The average number of users experiencing delay when requesting service from XCF.
SYSADNVC	The average number of users experiencing delay because of an operator mount request.
SYSADMVC	The average number of users experiencing delay because of an operator message request.
SYSCPUVC	Percentage of the maximum general purpose processor capacity spent on behalf of a class or group.
SYSEAPVC	Percentage of the maximum general purpose processor capacity consumed within a class or group (including enclave time).
SYSSRBVC	Percentage of the maximum general purpose processor capacity spent by SRB work on behalf of a class or group.
SYSTCBVC	Percentage of the maximum general purpose processor capacity used by non-enclave TCB work that executed within a class or group.
SYSIFAVC	Percentage of the maximum zAAP processor capacity used within a class or group.
SYSSUPVC	Percentage of the maximum zIIP processor capacity used within a class or group.
SYSCPVC	Percentage of the maximum general purpose processor capacity used by non-enclave TCB work that executed within a class or group.
SYSIFCVC	Percentage of the maximum general purpose processor capacity used by zAAP eligible work that executed within a class or group.
SYSSUCVC	Percentage of the maximum general purpose processor capacity used by zIIP eligible work that executed within a class or group.
SYSVELVC	Execution velocity. This value is calculated as CPU using, divided by the sum of CPU using and total delays gathered by WLM. The delays gathered by WLM include CPU delay and storage delay only.
SYSPDPVC	CPU time, in seconds, that transactions of a class or group were running at a promoted dispatching priority during the report interval.
SYSCVAVC	Whether CPU reconfiguration changes occurred during the reporting interval (YES or NO).

### Report options

The Report Options panel is exactly the same as for the SYSSUM report and STORS report, shown in Figure 102 on page 173. Selections made on either options panel affect all three reports.

## SYSRTD - Response Time Distribution Report

The Response Time Distribution (SYSRTD) report enables the service administrator and performance analyst to analyze the distribution of response time to see whether a response time goal was met and, if not, how close it came to failing. This report can also be used to "fine-tune" response time goals.

Use the bottom part of the report to see bottlenecks related to a specific system. Then you can use single-system reports for more detailed analysis.

### How to request this report

To request the Response Time Distribution report, select **S** from the Primary Menu, then a **2** on the Sysplex Report menu, (shown in Figure 5 on page 24) or enter one of the following commands:

SYSRTD service\_class, period

SYSRTD report\_class, period

For example, to get a Response Time Distribution report for the service class POSMULTI and service class period 1, enter:

SYSRTD POSMULTI, 1

### Contents of the report




Figure 105. SYSRTD Report - With Response Time Data

The SYSRTD report shows how the response time for a specific service or report class is distributed. Two levels of detail are shown:

- A character graphic shows the distribution of response time for all systems in a sysplex which have data available in the selected period.
- A table shows how each system contributed to the overall response time.

Depending on the goal of the period, different data may be available. A report for a service class period with a response time goal is shown in the previous figure. Here, the response time distribution is shown in the graph at the top of the screen.

A report for a service class period without response time goal is shown in the next figure. Here, no response time distribution graph is shown, and the top section of the report only shows the service class name and the service class period.

```

RMF V2R2 Response Time - RMFPLEX1 Line 1 of 5
Command ==> Scroll ==> HALF
WLM Samples: 100 Systems: 7 Date: 09/28/16 Time: 10.31.40 Range: 100 sec
Class: ALLBATCH
Period: 2

No response time distribution values available,
because there was no response time goal specified
for this service class period.

System Data --Avg. Resp. Time-- Trx --Subsystem Data-- --Exec Data--
WAIT EXECUT ACTUAL Rate Actv Ready Delay Ex Vel Delay
*ALL 32.9M 28.9M 61.8M 39.2 45 38
MVS6 all 5.381 8.03M 8.12M 11.8 82 16
MVS7 all 3.62M 1.48H 1.54H 0.1 91 1
MVS8 part 4.46M 49.5M 54.0M 3.4 86 45
MVS9 all 0.391 11.87 12.27 1.9 79 3
    
```

Figure 106. SYSRTD Report - Without Response Time Data

The response time distribution graph is not shown for heterogeneous report class periods. Please refer to “Performance data” on page 188 for an explanation of homogeneous and heterogeneous report class periods.

### How to read the graph

The horizontal axis shows response time (in seconds) with the response time goal in the middle. The middle section of the graph surrounding the goal shows the distribution of transactions that met between 60% and 150% of the goal.

For example, if the goal is 0.50 seconds, the middle section would include all transactions that completed within 0.30 seconds and 0.75 seconds.

Transactions that completed in less than 60% of the goal are accumulated and shown to the left of the middle section, and transactions exceeding 150% of the goal are accumulated and shown to right of the middle section.



The vertical axis shows the percent of transactions. To provide the best resolution possible, the axis has a variable scale based on the maximum percentage to be shown. The upper limit of the scale can vary from a minimum of 10% to a maximum of 100% in increments of 10%.

The graphical display represents the transactions that completed within a particular time.

In general, the colors mean:

- Fields marked with ■ in green indicates that transactions represented by this area completed within the response-time goal.
- Fields marked with X in red indicates that the transactions represented by this area did not complete within the response-time goal.
- Fields marked with ■ in blue represent the transactions that are not relevant for achieving the goal.
- A red, green, or blue '.' signifies that a small number of transactions took place.

The details of the graph depends on what kind of goal you have defined:

- **Average Response Time Goal**

There are two different cases:

- The goal has been achieved: all columns in the display are shown in *green*.
- The goal has not been achieved: all columns to the left of the average (representing transactions with a response time below the goal) are shown in *green*, the other columns are shown in *red*.

- **Response Time Goal with Percentile**

Now, the picture has this appearance:

- All columns to the left of the goal are shown in *green*, up to the goal percentile.
- All columns to the right of the goal, representing values that did not achieve the goal, are shown in *red* up to the goal percentile.
- All columns representing values that are not relevant for achieving the goal are shown in *blue*.

## Example

Response time goal: 80% of all transactions have a response time less or equal to 1 second.

**Case 1:** 90% are below 1 second

- 80% are shown in *green* (relevant to achieve the goal)
- 20% are shown in *blue* (not relevant to achieve the goal)

**Case 2:** 70% are below 1 second

- 70% are shown in *green* (relevant to achieve the goal)
- 10% are shown in *red* (did not achieve the goal)
- 20% are shown in *blue* (not relevant to achieve the goal)

## Scrollable part of report

The bottom section of this report is scrollable. It shows a list of all systems that have workload activity data gathered for the service class period during the report interval.

The first row in the scrollable area is a summary line. To indicate this, the *System* column displays the word *\*ALL*. The *Data* column remains empty, and all other columns contain the respective time value or delay percentage for the sysplex. The *Response time* columns for example show the same values as the row on the Sysplex Summary report for that service class period.

For each system, a row is shown where important response time data is provided. This information is intended to assist in tracking possible bottlenecks down to a specific system, where the analysis can be continued using the detailed reports for single systems.

### Data reported

The report is for one service class period. Depending on the type of service class, the different parts of the report may contain data or remain empty. Here is a list of what kind of data you can expect under which circumstances:

- Response Time Distribution  
Available only if a response time goal was specified
- Response Time Data  
Almost always available (possibly not for STC)
- Subsystem Data  
Available only for a subsystem transaction class
- Execution Data  
Available only if it is NOT a subsystem transaction class

### Cursor-sensitive control on the SYSRTD Report

In the non-scrollable area on the top of the report, which may show the response time distribution chart or a message that the data for that chart are not available, cursor-sensitive control is not active.

Cursor-sensitive control on the scrollable area on the bottom part of the report showing the system breakdown works as follows:

- In the first row, with *\*ALL* in the *System* column, cursor-sensitive control leads to a Response Time Components Data pop-up panel (see Figure 58 on page 103) which shows a detailed breakdown of the different wait reasons and their average duration.

In all other rows, you get the following:

- Cursor-sensitive control on column *System* leads to the SYSINFO report of the respective system.
- Cursor-sensitive control on column *Data* leads to the Data Index report of the respective system.
- Cursor-sensitive control on columns *Response Time* and *TRX Rate* leads to the GROUP report of the respective system.
- Cursor-sensitive control on any other column leads to the Delay report of the respective system.

## Field descriptions

Table 81. Fields in the SYSRTD Report

Field Heading	Meaning
% of TRX	The vertical axis represents the percentage of transactions that completed within a particular time. The axis has a variable scale in order to allow the best resolution possible. The scale can vary from 10% to 100% in increments of 10%. It is determined based on the maximum percentage to be shown on the graph.
Response time	Response time distribution.  For a description, refer to "How to read the graph" on page 182.
System	The four-character SMF system identifier.
Data	This column indicates whether the system has data for the entire report interval, or only for part of it. <b>all</b> Data could be retrieved that covers the report interval shown in the report header. <b>part</b> Data was retrieved that contains at least one time gap within the report interval. <b>none</b> No data could be retrieved for the report interval.
Avg. Resp. Time	The three columns under this header are the same as on the Sysplex Summary report except that the values here are calculated for a single system (except <b>*ALL</b> ). Refer to "SYSSUM - Sysplex Summary report" for a description.
Trx Rate	The transaction rate is the number of transactions ended per second. It is the same as on the Sysplex Summary report except that the value here is calculated for a single system (except <b>*ALL</b> ).
Subsystem Data	The three subsystem states shown here are the same as on the Work Manager Delay report (Response Time Breakdown) except that the values here are only calculated for a single system (except <b>*ALL</b> ). Refer to "SYSWKM - Work Manager Delays Report" on page 198 for a description.  The difference to the Work Manager Delay report is, that here the begin-to-end and the execution phase are combined within one row.
Execution Data Ex Vel (=Execution Velocity)	This is the same as the <i>Exec Vel - Actual</i> on the Sysplex Summary report with the only difference that the value here is only calculated for a single system (except <b>*ALL</b> ). Refer to "SYSSUM - Sysplex Summary report" for a description.
Execution Data Delay	This is the general execution delay used for the execution velocity calculation.  Note that in a service class more than one transaction can be delayed at the same point of time. For example, if two transactions on average are delayed each time WLM takes a measurement sample, a value of 200 will be displayed.

### Report options

The Report Options panel shows the RMF default options. It is the same as for the Group Response Time (GROUP) report, shown in Figure 59 on page 109, only the header line is different in showing the respective report name.

**Note:** The list of available service classes will be shown only if one of the sysplex reports SYSSUM, SYSRTD, or SYSWKM has been displayed at least once.

---

## SYSSUM - Sysplex Summary report

The Sysplex Summary (SYSSUM) report allows the service administrator and performance analyst to see at a glance whether service goals are being satisfied by:

- Showing a performance status line showing the performance status of the sysplex covering a time range of up to 80 refresh intervals.
- Showing the actual throughput being achieved by all workloads on one report
- Displaying goals not met in red or yellow
- Calculating the Performance Index for each service class period

## Mon III - SYSSUM

The report provides an overview of workload groups, service classes, service class periods, report classes, and report class periods. It allows "summarizing" of actual values for every group using threshold values, and includes a goal versus actual comparison for each period. You can compare different goals by using the performance index.

Furthermore, the response time for all groups is calculated independently of any specified goals, and a transaction rate is provided to enable you to weight the importance of the figures shown.

To facilitate detection of goals that were not met, the line of that service class period, as well as the related workload group and service class, is displayed in red or yellow.

Finally, options are available to select specific groups, or limit the report to groups that have exceeded their goal by a certain amount.

### How to request this report

To request the Sysplex Summary report, select **S** from the Primary Menu, then a **1** on the Sysplex Report menu (shown in Figure 5 on page 24), or enter the following command:

```
SYSSUM workload_group | service_class | report_class
```

## Contents of the report

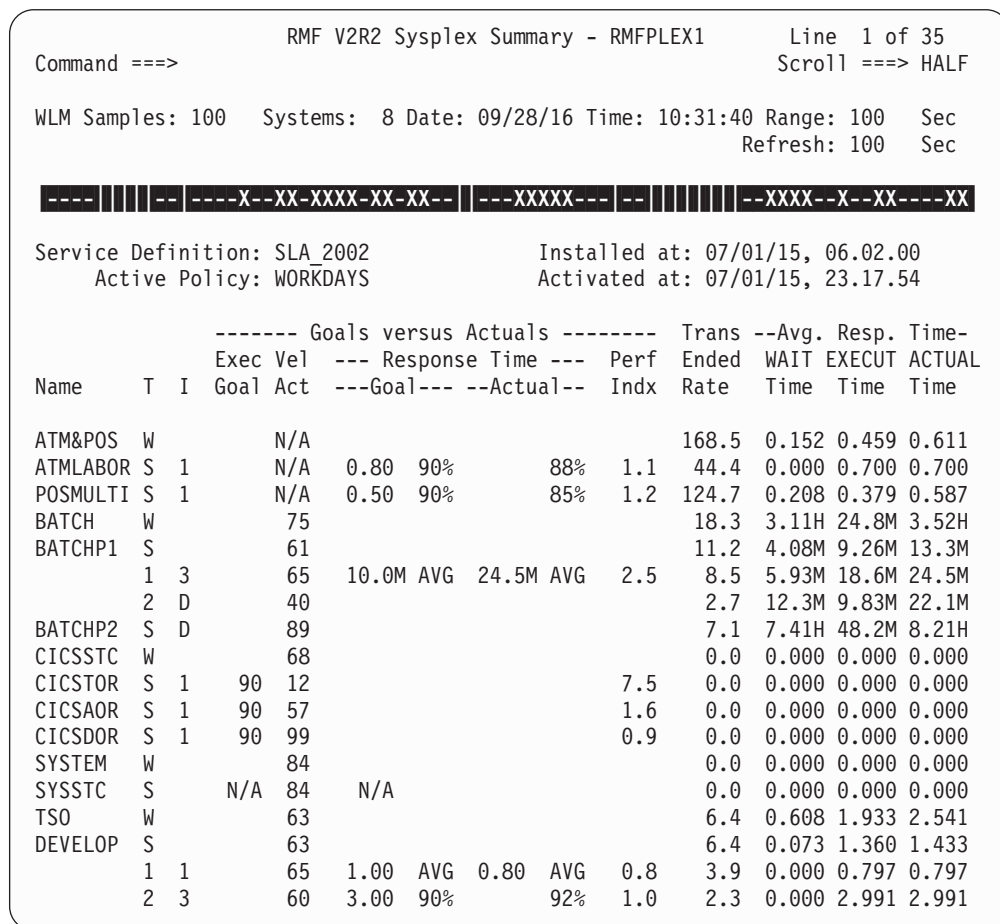


Figure 107. SYSSUM Report - GO Mode

The report can be logically broken into three sections.

### The Performance Status line

In GO mode, a colored status row at the top of the screen gives an overview over the sysplex during the last ranges. For every range, a one-character field will be added to the right-hand side of the performance status line. For each range that has been reported, the one-character field is marked in one of the following ways:

**■ (green)**

If all goals have been met during that interval (the performance index is less than, or equal to, 1 for all periods)

**- (yellow)**

If a warning level was reached during that interval (the performance index was greater than 1 for periods with an importance equal to 3, 4, or 5)

**X (red)**

If goals have been exceeded seriously during that interval (the performance index was greater than 1 for periods with an importance equal to 1 or 2)

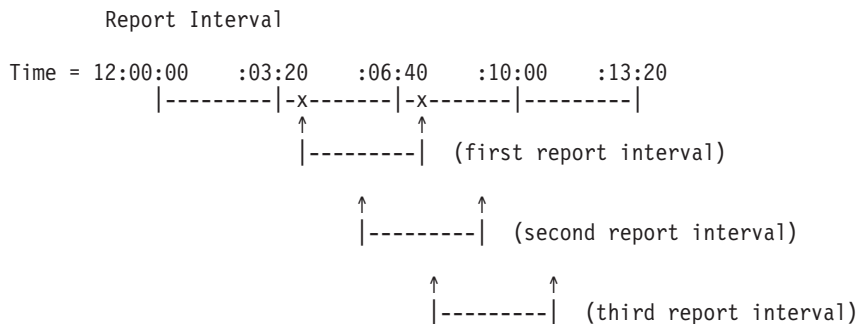
**Blank** If you changed the mode from GO to STOP during some intervals

The **Refresh** value that you can specify on the Session Options panel will define the how often the status line will be updated.

If you define a refresh value that is less than the MINTIME, the refresh value will be reset to equal the MINTIME.

**Example**

For example, if you specify a Refresh value of 100 seconds and a Range of 200 seconds, the status line will be updated every 100 seconds, showing information from the latest 200 seconds.



If you end GO mode, and do not change the Refresh or range values, historical data will be saved, so that when you start GO mode again, the status line will continue where it left off. If you do change the refresh or range value, the status line will start from the beginning again. See Figure 107 on page 187 for an example.

In STOP mode, the row is reduced to a single colored field in the center of the report that shows the overall status of the displayed range. For example:

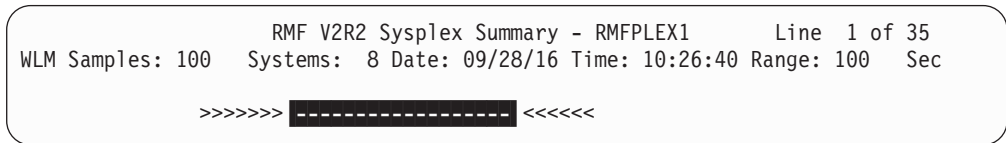


Figure 108. SYSSUM Report - STOP Mode

**Service definition information**

Two sub-header lines show the name of the current service definition together with the installed date and time and the name of the active policy together with the activation date and time.

**Performance data**

The rest of the report shows the execution velocity goals and response time goals versus actual values for each service class period and for each homogeneous report class period. The report also shows the average transaction rate to give you an indication of how significant the actual values are.

**Report Class Periods:**

**Homogeneous report class period:** A report class period is called homogeneous if all its transactions are being assigned to the same service class period.

*Example:* You classify all TSO users to run in service class TSOPROD and distinguish the departments for reporting purposes in report classes TSOEPTA, TSOEPTB, and TSOEPTC. This definition done in the WLM application creates homogeneous report classes.

All other report class periods are called heterogeneous. Reporting for response time distribution and subsystem delays is available only for homogeneous report class periods.

**Heterogeneous report class period:** A report class period is called heterogeneous if its transactions are being assigned to different service class periods.

*Example:* You classify all TSO users by accounting information and assign service classes TSOEPTA and TSOEPTB. There is one common report class TSOREPCL. This definition done in the WLM application creates a heterogeneous report class.

The average response time column shows the average time that a transaction spent waiting in a queue and was active in the system. This gives you an indication of where a possible response time bottleneck may be located.

A performance index is introduced to allow a better comparison between different goals. See Table 82 on page 191 for information on how to calculate the performance index.

The scrollable area is ordered by workload group. Each workload group is followed by a list of its service classes. Every service class is followed by a detailed comparison of actual values versus goals for each service class period. The workload groups, and the service classes below each workload group, are sorted alphabetically.

In detail, the rows show the following:

- For each workload group (indicated by the type W), one line is shown containing the actual values achieved for the whole group. This can be seen as a summary line for that group.
- For each service class (indicated by the type S), one line is shown containing the actual values achieved for the whole class. As for workload groups, it can be seen as a summary line for that class.
- For each service class period (indicated by the period number in the type column), one line is shown containing the defined goals accompanied by the values actually achieved. For a service class with one period, the data of that period is shown.

While the execution velocity goal is a percentage that can easily be compared with an actual value, the WLM response time goals can be specified in two different ways:

- A response time together with a percentile:  
In this case, the actual value is a percentage indicating the percentage of transactions that ended within the time specified in the response time goal.
- An average response time:  
In this case, the average response time value is shown as actual value that can be compared against the goal.

### **Cursor-sensitive control on the SYSSUM Report**

Cursor-sensitive control on this report lets you navigate to detailed reports that offer a possibility to make single system selections.

Cursor-sensitive control of the sysplex field and of the *Systems* field in the report header leads to the Data Index screen.

Using cursor-sensitive control in the *Type* column:

- On a workload group abbreviation, the report is redisplayed containing only workload group entries.
- On service class or report class abbreviation, the processing is analogous to the workload group abbreviation.

Cursor-sensitive control on a field in the *Importance* column gives you a filtered report. What is displayed on the filtered report depends on the **Type** value that you can specify on the Report Options panel.

- If Type is **ALL**, a workload group and all its service classes are displayed if one service class period has the importance you selected using cursor-sensitive control
- If Type is **W**, a workload group is displayed if it contains a service class period with the importance you selected using cursor-sensitive control
- If Type is **S**, a service class is displayed if it contains a service class period with the importance you selected using cursor-sensitive control

Cursor-sensitive control in all other columns:

- If it is a service class period for which subsystem delay data are available, then the Work Manager Delay report is shown.
- Otherwise, the Response Time Distribution report is shown.

Filtering on workload groups and service classes is possible using cursor-sensitive control, and a single workload group or service class can be selected via report option or command parameter.

The default option will be filtering on workload groups. An example of this is shown in Figure 109.

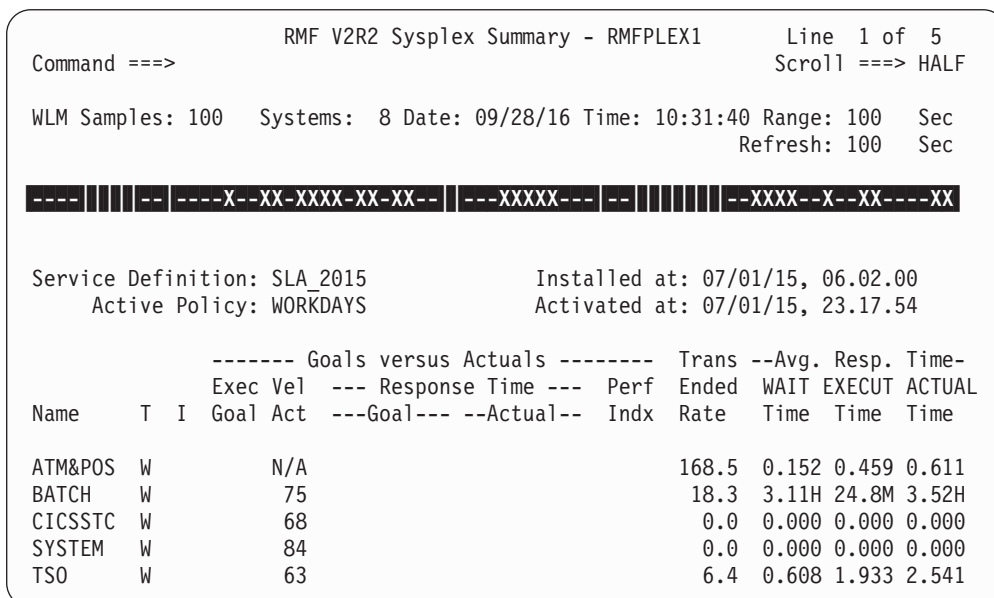


Figure 109. SYSSUM Report for Workload Groups



## Field descriptions

Table 82. Fields in the SYSSUM Report

Field Heading	Meaning
Refresh	The refresh value specified on the Session Options panel. The report will be updated according to this value.
Service Definition	Name of the service definition in effect during collection of the workload activity data.
Installed at	The date and time the active service definition was installed.
Active Policy	Name of the service policy in effect during collection of the workload activity data.
Activated at	The date and time the current service policy was activated.
Name	Name of the group being reported on. This field can contain a workload group name, a service class name, or a report class name. For a line with a period, the field remains empty, because there will always be a summary line further up containing the name of the class to which the period belongs.
T	Type of group being reported on. <b>W</b> Workload group <b>S</b> Service class <b>R</b> Report class <b>n</b> Service/report class period
I	Importance, describes the level of importance assigned to a service class period. Since workload groups, report classes, and system service classes do not have an importance, this column remains empty on those lines. For a service class with multiple periods, this column remains empty as well, whereas for a service class with one period, the importance of that period is shown in the service class row.  If "discretionary" was specified as a goal, this is indicated by a <b>D</b> in this column, since an importance cannot be defined for a discretionary goal. <b>1</b> Highest - describes highest priority service class period for most important work <b>2</b> High <b>3</b> Medium <b>4</b> Low <b>5</b> Lowest <b>D</b> Discretionary
Goals versus Actuals	For a service or report class period, these columns show the goal, if specified, and the actual values corresponding to the goal. That means, if an execution velocity goal was specified, the <i>Response Time</i> goal and actual columns remain empty.  If a response time goal with a percentile was specified, the <i>Actual</i> column shows a percentage which corresponds to the response time specified in the goal. The actual average response time can be found in the columns for <i>Avg. Resp. Time</i> .  For a report class period, this data will be shown only if the period is homogeneous.
Execution Velocity Goal	The target execution velocity for ended transactions that has been in effect for the period during the reported range. This field shows N/A for system service classes, since they do not have a user defined goal.

## Mon III - SYSSUM

Table 82. Fields in the SYSSUM Report (continued)

Field Heading	Meaning
Execution velocity Actual	<p>The execution velocity of the workload group, service or report class, or period being reported on. This value is calculated independent of a specified goal.</p> <p>A high value indicates little workload contention while a low value indicates that the requests for system resources are delayed.</p> <p>See "Execution velocity" on page 15 for details about the execution velocity.</p> <p>Whenever subsystem delays are available for that service class period, <b>N/A</b> is shown in this field. The Work Manager Delays report provides more information for these service class periods.</p> <p>Whenever the service class is a "server", the velocity is calculated, but the field is shown in dark blue. A service class is a "server", when the array of service classes served is not empty.</p> <p>In the <i>Execution Velocity - Actual</i> column, two exceptions may occur:</p> <ul style="list-style-type: none"> <li>• <b>N/A</b> appears for "transaction" service classes, that means, for classes served by "server" service classes. For these classes, the execution velocity cannot be defined meaningfully.</li> <li>• A velocity value in dark blue appears for "server" service classes, that means, for classes containing address spaces that give service to the "transaction" service classes. For those classes some goals may be specified, but they are not used. Instead they are managed based on the goals of the "transaction" service classes.</li> </ul>
Time	<p>The time units shown in the columns <i>Goal versus Actuals - Response Time</i> and <i>Avg. Resp. Time</i> are all seconds, if no unit is marked. Time values bigger than one minute are marked with <b>M</b>, and time values bigger than one hour are marked with <b>H</b>.</p>
Response Time Goal	<p>This field shows two columns which together describe the goal that has been in effect for the service or report class period during the reported range:</p> <ul style="list-style-type: none"> <li>• The average target response time for all ended transactions.</li> <li>• The percentage of transactions that should terminate within the time specified in the goal. This percentage is specified together with the response time value.</li> </ul> <p>For a goal without percentage, <b>AVG</b> is shown in this field. This field shows <b>N/A</b> for system service classes, since they do not have a user defined goal, and it is empty for heterogeneous report class periods.</p>
Response Time Actual	<ul style="list-style-type: none"> <li>• Average response time goal: The value represents the average response time for all ended transactions, followed by <b>AVG</b>.</li> <li>• Response time goal with percentile: The percentage of transactions that actually ended within the time specified in the goal.</li> </ul>

Table 82. Fields in the SYSSUM Report (continued)

Field Heading	Meaning
Performance Index	<p>This index helps to compare goals. If, for example, several execution velocity goals with the same importance are not met, this index helps you decide which group was impacted the most.</p> <p>If RMF cannot calculate the performance index, this field contains N/A:</p> <ul style="list-style-type: none"> <li>For a period with a response time goal: The sum of completed transactions is zero</li> <li>For a period with an execution velocity goal: An actual value is not available (for example, a subsystem service class, for which no actual execution velocity is shown)</li> </ul> <p>This field is blank for heterogeneous report classes periods and report classes. It is also blank for system service classes, since they do not have a user defined goal.</p> <p>The <i>Performance Index</i> field can also be shown in dark blue. This has the same reason as for the <i>Execution Velocity - Actual</i> column, described previously.</p> <p>RMF calculates the performance index depending on the type of goal:</p> <ul style="list-style-type: none"> <li><b>Execution velocity goal</b> <math display="block">\text{Perf Indx} = \frac{\text{Goal \%}}{\text{Actual \%}}</math> </li> <li><b>Average or percentile response time goal</b> <math display="block">\text{Perf Indx} = \frac{\text{Actual (sec)}}{\text{Goal (sec)}}</math> </li> </ul> <p>"Actual" means the maximal response time that actually was reached for the percentage of the goal and is calculated by performing the following three steps:</p> <ol style="list-style-type: none"> <li>Calculate the number of transactions N that correspond to the goal: <math display="block">N = \frac{\sum \text{Transactions} * \text{Goal Percentage}}{100}</math> </li> <li>Add up all transactions until a bucket M is reached where the sum is greater than N.</li> <li>The "actual" response time in the previously shown formula for the performance index is the response time value belonging to the bucket M.</li> </ol> <p><b>Note:</b> Due to this methodology, the maximal value of the performance index for this goal type is 4. If the sum of all transactions belonging to buckets 1 to 13 is below the goal percentile, the performance index is shown as '****'.</p>

The following example shows how to calculate the performance index for a response time goal with a percentile.

**Example**

Calculation of the performance index for a response time goal with percentile:

Example goal: Time = 2.0 sec Percent = 80%

The four lines in this example show:

- Number of buckets
- Response time distribution
- Number of transactions (within that bucket)
- Response time associated with that bucket

```

Bucket:  1  2  3  4  5  6  7  8  9 10 11 12 13 14
Distr :<50% 60% 70% 80% 90% 100% 110% 120% 130% 140% 150% 200% 400% >400%
TRX   : 10 10 20 20 20 20 20 10 10 10 5 5
Time  : 1.0 1.2 1.4 1.6 1.8 2.0 2.2 2.4 2.6 2.8 3.0 4.0 8.0 >8.0
    
```

The three steps of calculation:

**Mon III - SYSSUM**

1. The sum of all transactions is 160, so the number of transactions needed to fulfill the goal is:

$$N = \frac{160 * 80}{100} = 128$$

2. Adding all transactions until sum is greater than N leads to bucket M = 8, because the sum of all transactions including bucket 8 is 130.

3. The response time belonging to bucket 8 is 2.4 seconds, so the result is:

$$\text{Perf Indx} = \frac{2.4}{2} = 1.2$$

**Field descriptions- continuation**

*Table 83. Fields in the SYSSUM Report - Continuation*

Field Heading	Meaning
Trans Ended Rate	The number of transactions ended per second.
Avg. Resp. Time WAIT Time	<p>The average time (in seconds) that a transaction spent waiting because of one of these reasons:</p> <ul style="list-style-type: none"> <li>• Queued Average time a job was delayed for reasons other than the ones mentioned below. This field therefore basically includes the time a job was delayed for initiation. For TSO users, this can be a portion of LOGON processing. For APPC, this is the time the transaction spent on an APPC queue.</li> <li>• R/S Affinity - Resource affinity scheduling delay Average time the job was delayed due to resource or system affinity scheduling. This means that resource(s) required for the job to run were not available at some point while the job was queued to JES2.</li> <li>• Ineligible - Operational or JES scheduling delay Average time a job was delayed due to operational delays or JES scheduling delays, examples are: <ul style="list-style-type: none"> <li>- Job held by operator</li> <li>- Job class or job queue held</li> <li>- Duplicate jobname serialization</li> <li>- Job class execution limits</li> </ul> </li> <li>• Conversion - JCL conversion delay Average time a job was delayed for JCL conversion. Jobs held during conversion (due to affinity, HSM recall, or enqueue contention) contribute only to conversion time, not to ineligible or R/S affinity times. Conversion time is not part of the total response time.</li> </ul> <p>The time a job was delayed due to TYPRUN=HOLD or TYPRUN=JCLHOLD is NOT included in any of the transaction times.</p> <p>In all other cases, this is the average time that transactions spent waiting on a JES or APPC queue. Also note that queue time may not always be meaningful, depending on how the customer schedules work. For example, if a customer submits jobs in hold status and leaves them until they are ready to be run, all of the held time counts as queued time. That time may or may not represent a delay to the job.</p> <p>In the <i>Avg. Resp. Time</i> columns, zeros will show up for "server" service classes in most cases, because their "transactions" are address spaces, and response times are available only for ended transactions. So there are only numbers, when one of the address spaces in that service class ends, or is RESET via operator command.</p>

Table 83. Fields in the SYSSUM Report - Continuation (continued)

Field Heading	Meaning
Avg. Resp. Time EXECUT	<p>For CICS transactions, this includes execution time in AOR and following regions.</p> <p>For IMS transactions, this includes execution time within the MPR.</p> <p>For Batch, TSO, etc., this is the average time that transactions spent in execution.</p> <p>In the <i>Avg. Resp. Time</i> columns, zeros will show up for "server" service classes in most cases, because their "transactions" are address spaces, and response times are available only for ended transactions. So there are only numbers, when one of the address spaces in that service class ends, or is RESET via operator command.</p>
Avg. Resp. Time ACTUAL Time	<p>In general, this is the sum of the previously described wait and execution times, but does not include ineligible time.</p> <p>For CICS transaction service classes, you may see the average EXECUT time greater than the average ACTUAL time, when you would normally expect EXECUT to be less than or equal to ACTUAL. This is because these two fields report on a different set of transactions. EXECUT time can include transactions which originated on a remote system as well as transactions originating locally. ACTUAL time includes response times for only transactions originating locally. If the remote transaction tends to be longer than the local transaction, EXECUT could be greater than ACTUAL.</p> <p>It should be noted that all of these response times are for ended transactions only. Thus, if there is a problem where transactions are completely locked out, either while queued or running, the problem will not be seen on this report until the locked out transactions end.</p> <p>In the <i>Avg. Resp. Time</i> columns, zeros will show up for "server" service classes in most cases, because their "transactions" are address spaces, and response times are available only for ended transactions. So there are only numbers, when one of the address spaces in that service class ends, or is RESET via operator command.</p>

### Monitor III Utility fields

You can use the Monitor III Utility to customize the SYSSUM report. In addition to the values previously described, you can use the Utility to have the following information shown.

Table 84. Additional Fields in the SYSSUM Report

Field Heading	Meaning
Goal type	<p>This goal type indication includes a list of all different types of goals.</p> <ul style="list-style-type: none"> <li>1 Percentile response time goal</li> <li>2 Average response time goal</li> <li>3 Velocity goal</li> <li>4 Discretionary goal</li> </ul>
Duration	<p>Service class period duration in unweighted CPU service units (that means, not multiplied with the service coefficients) per second.</p> <p>A duration is required in all but the last service class period. For single periods or for the last period of multiple periods this value is always zero.</p>
Resource Group	Name of the resource group associated with the work in this service class.
Capacity Min	<p>Minimum service rate, in unweighted CPU service units per second, as defined for that resource group.</p> <p><b>Note:</b> The minimum and maximum values must not be specified by the user.</p>
Capacity Max	Maximum service rate, in unweighted CPU service units per second, as defined for that resource group.
Capacity Actual	Actual service rate, in unweighted CPU service units per second, as consumed by that resource group.

### Report options

The Report Options panel is exactly the same as for the SYSINFO report and STORS report, shown in Figure 102 on page 173. The only difference is, that the group names shown are accumulated from the whole sysplex and not only from a single system.

**Note:** The list of available service classes will be shown only if one of the sysplex reports SYSSUM, SYSRTD, or SYSWKM has been displayed at least once.

---

## SYSTREND - System Trend Report

The System Trend (SYSTREND) report presents the last 20 reporting ranges for the system summary line (\*SYSTEM) or any other selected workload line from the System Information (SYSINFO) report. It can be used:

- To analyze how delay situations develop in the system
- To analyze how long delay situations last in the system
- To understand the change in system utilization
- To identify peak utilization
- As a system summary report

You can also use this report as a starting point to analyze system or workload delay. For a more detailed analysis, select the SYSINFO or respective delay report for any reported range using cursor-sensitive control.

### How to request this report

To request this report, select **U** from the Primary menu, then **ST** from the User menu.

You need to enter a system name on the User menu.

**Note:** If no workload or an invalid workload name is specified on the User Selection menu, the report is created for the \*SYSTEM line as shown on the System Information (SYSINFO) report.

### Contents of the report

The SYSTREND report has two parts.

- The top part provides information about the start date and time of the first and last reported range, the total reported range and the total and average number of samples used to create the report.
- The bottom part shows for each line the CPU utilization for the system and the respective SRB and TCB percentage, total and active users, and average number of jobs using resources or delayed because of resources for the selected workload.

This information is extracted from the SYSINFO reports for the displayed report interval and the meaning is the same as for the SYSINFO report.

The graphic form of the report shows the average number of active users for each type of delay for the selected workload.

**Note:** On the SYSTREND report, the RESP Time, TRANS/SEC and VEC Util columns (shown on the SYSINFO report) have been replaced by the TCB% and SRB% columns contained in the ISPF table of the SYSINFO report.

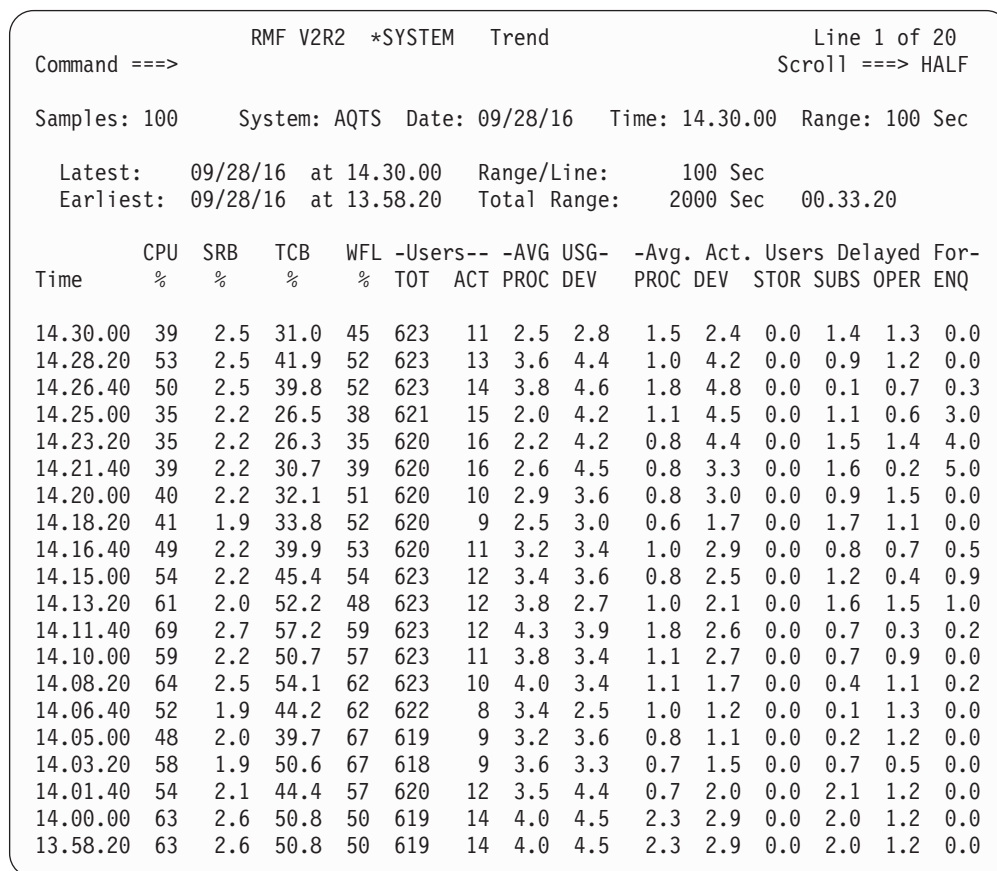


Figure 110. SYSTREND Report

### Field descriptions

Table 85. Fields in the SYSTREND Report

Field Heading	Meaning
Latest:	Begin date and time of the first reported range on the report
Range/Line:	Reported range per displayed line on the report
Earliest:	Begin date and time of the last reported range on the report
Total Range:	Total reported range on the report, expressed in seconds and HH.MM.SS
Time	The begin time of the reported range
CPU %	The average CPU utilization percentage for all processors is also displayed on the SYSINFO report. See Table 79 on page 178 for the calculation.
SRB%	The average percentage of SRB time used by all address spaces per processor during the report interval.
TCB%	The average percentage of TCB time used by all address spaces per processor during the report interval.

All other fields in the SYSTREND report are the same as in the SYSINFO report (see Table 79 on page 178).

### Cursor-sensitive control

Cursor-sensitive control allows you to navigate to the SYSINFO or a delay report for a selected reporting range.

Table 86. SYSTREND Report - Cursor-sensitive Control for Navigation

<i>Report Column where Cursor-sensitive Control is Used</i>	<i>Displayed Report</i>
Time	SYSTREND
CPU%, TCB%, SRB%, WFL%, Users	SYSINFO
-AVG USG- PROC, -Avg. Act. Users Delayed For- PROC	PROC
-AVG USG- DEV, -Avg. Act. Users Delayed For- DEV	DEV
-Avg. Act. Users Delayed For- STOR	STOR
-Avg. Act. Users Delayed For- SUBS	DELAY
-Avg. Act. Users Delayed For- OPER	DELAY
-Avg. Act. Users Delayed For- ENQ	ENQ

**Note:** If the SYSTREND report is recreated using cursor-sensitive control on the Time column, pressing PF3 on the new SYSTREND report will return you to the Primary Menu. In all other cases, pressing PF3 from the report you have selected will return you to the SYSTREND report.

---

## SYSWKM - Work Manager Delays Report

The Work Manager Delays (SYSWKM) report shows details for resource-manager or work-manager oriented subsystems and is intended as a basis on which to start tuning. Using the real-time data, you can use this report to track problems as they happen.

The report shows the average transaction response time and how the various transaction states contribute to it. Furthermore it lists the address spaces that have been used by the transactions. Figure 111 on page 199 and Figure 112 on page 200 show sample reports for CICS and IMS data.

This report allows you to track subsystem problems.

When defining your service definition, you should try to separate short and long transactions into different service classes. This can help in providing more meaningful reports due to internal processing and measurement reasons.

A high value in one or more of the reported states (LOCK, I/O, CONV, DIST, SESS, TIME, PROD, LTCH, MISC, LOC, SYS, or REM) can indicate a problem.

The lower part of the report shows the address spaces serving the reported service class. A high delay value (Capp or Quies) can indicate the cause of a high response time value for the reported service class. This part is empty if you call the report for a report class.

The Proc-Usg and Veloc columns give an indication of how much work is actually being done, and should be as high as possible.

### How to request this report

To request the Work Manager Delays report, select **S** from the Primary Menu, then a **3** on the Sysplex Report menu (shown in Figure 5 on page 24) or enter the following command using the format:

```
SYSWKM service_class,period
SYSWKM report_class,period
```



As this report is available for homogeneous report classes only, you cannot specify a heterogeneous report class with the reportclass parameter.

For example, to get a Work Manager Delays report for the service class POSMULTI and service class period 1, enter:

SYSWKM POSMULTI, 1

## Contents of the report

```

RMF V2R2 Work Manager Delays - RMFPLEX1 Line 1 of 3
Command ==> Scroll ==> HALF
WLM Samples: 100 Systems: 2 Date: 09/28/16 Time: 16.03.00 Range: 100 Sec
Class: POSMULTI Period: 1 Avg. Resp. time: 0.587 sec for 12473 TRX.
Goal: 0.500 sec average Avg. Exec. time: 0.379 sec for 12389 TRX.
Actual: 0.587 sec average Abnormally ended: 0 TRX.

Sub P -----Response time breakdown (in %)------ -Switched--
Type Tot Act Rdy Idle -----Delayed by----- Time (%)
CONV I/O LOCK MISC PROD LOC SYS REM
CICS B 84 9 0 0 65 0 5 5 0 40 25 0
CICS X 43 4 8 0 0 18 6 6 1 0 0 0
IMS X 16 5 0 0 0 11 0 0 0 0 0 0

----- Address Spaces Serving this Service Class POSMULTI -----
Jobname M ASID System Serv-Class Service Proc-Usgr I/O-Usgr Veloc Capp Quies
CICSTOR1 Y 0102 MVS1 CICSTOR 36 6 11 36 0 0
CICSTOR2 0129 MVS2 CICSTOR 64 11 2 54 0 0
CICSAOR1 0258 MVS2 CICSAOR 21 8 3 5 18 0
IMSDBCTL 0091 MVS2 SYSTC 48 7 1 73 0 0

```

Figure 111. SYSWKM Report for Subsystem CICS

The panel shows an example for a CICS system using IMS as database, where all measurement values that theoretically could be provided are actually available. However, in a CICS system with regions spread over several z/OS images and with different CICS releases installed, data may be available or missing in nearly every column, depending on how varied the installation is and how the CICS releases involved differ.

```

RMF V2R2 Work Manager Delays - RMFPLEX1 Line 1 of 6
Command ==> Scroll ==> HALF
WLM Samples: 100 Systems: 2 Date: 09/28/16 Time: 16.03.00 Range: 100 Sec
Class: CRDTAUTH Period: 1 Avg. Resp. time: 0.587 sec for 12473 TRX.
Goal: 0.500 sec for 80% Avg. Exec. time: 0.311 sec for 12389 TRX.
Actual: 0.500 sec for 91% Abnormally ended: 0 TRX.

Sub P -----Response time breakdown (in %)------ -Switched--
Type Tot Act Rdy Idle -----Delayed by----- Time (%)
I/O CONV LOCK LOC SYS REM

IMS X 86 18 7 0 41 15 5 0 0 0

----- Address Spaces Serving this Service Class CRDTAUTH -----
Jobname M ASID System Serv-Class Service Proc-Usg I/O-Usg Veloc Capp Quies
IMSRDRE 0102 MVS1 STC_HIGH 36 6 12 36 0 0
DBRCP1 0129 MVS1 STC_HIGH 64 11 2 54 0 0
DLIRDRE Y 0258 MVS1 STC_HIGH 21 8 3 42 0 0
RDREP001 0033 MVS1 STC_HIGH 53 2 6 33 0 0

```

Figure 112. SYSWKM Report for Subsystem IMS

The report is for one service or report class period. The period number is included in the report because it is possible to define multiple periods, however, in most cases, you will only define one period with one goal for a transaction-oriented subsystem.

The report can be invoked for every service or report class, but data can only be shown when subsystem work manager delays have been gathered for that class in that period.

As mentioned above, such class must contain a group of CICS or IMS transactions. For example, a class can include:

- Transactions that have the same service level objectives, such as response time. In this case, possible service classes could be CICSFAST and CICSSLOW.
- Transactions that relate to one another. In this case, possible service classes could be CICSLOCL for all transactions of local CICS user, and CICSRMOT for all transactions routed from another CICS region.

The report contains three sections:

### Performance information

The top section of the report shows the class name and period together with the goal and response time, execution time, and corresponding completion counts.

The goal shown is one of the following, depending on what was defined for the period:

- A response time goal (either percentage or average value)
- An execution velocity goal with a percentile
- Discretionary

If the goal is not met, it is shown in red.

The third sub-header line shows the actual value corresponding to the goal. It is calculated in the same way as for the Sysplex Summary report.

The **Average Response time** is shown, followed by the number of total transactions that completed normally during the report interval.

Below this, the **Average Execution time** is shown together with the number of transactions that completed their execution phase normally during the report interval.

Finally, the number of **Abnormally ended** transactions are included.

The response time shown is always the average of all transactions, so be careful when comparing this value with a percentile goal.

### Example

Assume that 12434 transactions have an average response time of 0.4 seconds, but the remaining 39 have an response time of 1 minute each. The goal is a response time of 0.5 seconds for 80% of the transactions. Then we get the following:

```

                                Avg. Resp. time: 0.586 sec for 12473 TRX.
Goal: 0.500 sec for 80%   Avg. Exec. time: 0.311 sec for 12389 TRX.

```

Here the goal is met, even though the average response time is 0.586 seconds.

### Response time breakdown

In the middle section, a response time breakdown for the various transaction states is shown, split into total time and execution time.

The unit of the response time values shown can be switched between percentages and seconds using the Report Options panel or use cursor-sensitive control anywhere in the middle section of report.

If you select seconds and the value does not fit, then \*\*\* will be shown in that report field. In this case, changing to percentage will provide a better representation of the figures.

### Address spaces serving

In the bottom section, the **Address Spaces Serving this Service Class** during the report interval are listed in a scrollable area.

- For each address space, the jobname (together with an indication about how WLM is managing a server region), address space id, and system id are shown to allow you to track the address space to the specific z/OS image.
- The service class shown is the class the serving address space belongs to. This helps you to relate this data to the Sysplex Summary report.
- The service percentage (**Service**) shows the percentage of service given to the reported service class. For example, if the address space serves only this class, then 100% is shown. If the address space gives equal service to three different service classes, then 33% is shown.
- Then, for each address space the execution velocity, the processor and I/O using percentage are shown. This gives you a hint of the "health" of the address space.
- Finally, for each address space a capping percentage and a quiesce percentage are shown. The capping column shows the WLM percentage for capping. If the address space was delayed for other reasons as well, the actual capping delay may be much smaller. The actual capping delay is added to the single system Processor Delays (PROC) report.

The quiesce percentage normally shows either 0 or 100, because the address space is either quiesced by the operator with the RESET command or not.

However, for a time range where the quiesce state was changed, a percentage between 0 and 100 is possible, indicating how long the address space was quiesced during the report interval.

This part of the report is empty if you call the report for a report class.

### Cursor-sensitive control on the SYSWKM Report

Cursor-sensitive control used on the response time fields in the sub-header lines in the top portion of the report shows you the Response Time Distribution report for that service class period.

Cursor-sensitive control used on the response time breakdown fields in the middle portion of the report switches between the units that can be selected for the data. This "toggling" does not change the unit selected on the Report Options panel.

- If the current unit is *seconds*, cursor-sensitive control switches the unit to percentage.
- If the current unit is *percentage*, cursor-sensitive control switches the unit to seconds.

Cursor-sensitive control on the server address space section in the scrollable bottom portion of the report is active on the following fields:

- Cursor-sensitive control on column *Jobname* and *ASID* leads to the JOB Delay report of the respective system.
- Cursor-sensitive control on column *System* leads to the SYSINFO report of the respective system.
- Cursor-sensitive control on column *Service Class* and *Service* leads to the GROUP report of the respective system.
- Cursor-sensitive control on column *Proc-Usg* leads to the PROC report of the respective system.
- Cursor-sensitive control on column *I/O-Usg* leads to the DEV report of the respective system.
- Cursor-sensitive control on column *Velocity* leads to the Delay report of the respective system.

### Field descriptions

All fields of the report are described in detail in the following field description table:

Table 87. Fields in the SYSWKM Report

Field Heading	Meaning
Class	The name of the service or report class.
Period	The period number.
Goal	The goal for the reported class as contained in the service policy. The goal can be average response time in seconds, minutes, or hours, the percentage of a response time goal, "Discretionary", "N/A" (not applicable), or blank.
Actual	Depending on the type of goal, this field shows the actual response time, or the field is blank.
Average response time	The average response time of all ended transactions belonging to the period, possibly spread over several systems.
For <i>nnnnn</i> TRX.	<i>nnnnn</i> is the total number of completed transactions.
Average execution time	The average execution time of all ended transactions belonging to that period, possibly spread over several systems.
For <i>nnnnn</i> TRX.	<i>nnnnn</i> is the total number of transactions that completed their execution phase during this report interval.

Table 87. Fields in the SYSWKM Report (continued)

Field Heading	Meaning
Abnormally ended	The number of abnormally ended transactions from all reported systems. This value is not included in the number of total completed transactions.
Subsystem Type	A 4-character identification for the subsystem for which the data was attributed to, as shown for example in the WLM administrative application.
Phase (P)	<b>B</b> This line represents states of the begin-to-end phase of the transactions. <b>X</b> This line represents states of the execution phase of the transactions.
Response time breakdown	<p>Both begin-to-end phase (Phase = B) rows and execution phase (Phase = X) rows show a breakdown of the average response time (B) or average execution time (X). For consistency, all values (both B and X) are related to the average response time (Avg. Resp. Time).</p> <p>If several execution phases (X) are shown, it is not possible to sort them hierarchically. It is only possible to regard the sum of all execution phases as a breakdown of the average response time shown in the sub-header lines.</p> <p><b>Tot</b> Total amount of time that the transactions spent in states that are shown in this report. These states are not a complete breakdown of the response time shown in the sub-header. There is always a gap due to states that are not reported.</p> <p>The value is a sum of all the figures shown in this row in the other "Response time breakdown" columns. <b>Note:</b> Because one transaction can be counted in more than one state during a report interval, this number can be larger than 100.</p> <p><b>Act</b> Time spent in an active state.</p> <p>Besides the time spent in an active subsystem state, this field also contains the time spent in an active application state, if provided by the subsystem (for example, Websphere).</p> <p>Active indicates that, from the work manager's perspective, there is a program executing on behalf of the work request. This does not mean that the program is active from the base control program's perspective.</p> <p><b>Rdy</b> Time spent in a ready state.</p> <p>Ready indicates that there is a program ready to execute on behalf of the work request described by the monitoring environment, but the work manager has given priority to another work request.</p> <p><b>Idle</b> Time spent idle means that no work request (or transaction) is available to be run by the work manager.</p>

Table 87. Fields in the SYSWKM Report (continued)

Field Heading	Meaning
<p>Response time breakdown</p> <p>Delayed by</p>	<p>The report will just present the eight highest non-zero values. These are determined by sorting the sum of the rows for each delay reason.</p> <p><b>LOCK</b> Time spent waiting for a lock.</p> <p><b>I/O</b> Time spent of waiting for I/O.</p> <p>Waiting for I/O indicates that the work manager is waiting on an activity related to an I/O request. This may be an actual I/O operation or some other function associated with the I/O request.</p> <p><b>CONV</b> Time spent waiting for conversation.</p> <p><b>DIST</b> Time spent waiting for distributed request state samples.</p> <p>Waiting for a distributed request indicates that some function or data must be routed prior to resumption of the work request. This is in contrast to 'waiting on conversation', which is a low level view of the precise resource that is needed. A distributed request could involve 'waiting on conversation' as part of its processing.</p> <p><b>SESS</b> Time spent waiting for a session to be established.</p> <p>This is a sum of the time spent waiting for sessions to be established locally (for example, on the current z/OS image), somewhere in the network, or somewhere in the sysplex.</p> <p><b>TIME</b> Time spent waiting for a timer.</p> <p><b>PROD</b> Time spent waiting for another product.</p> <p><b>LTCH</b> Time spent waiting for a latch.</p> <p><b>MISC</b> Time spent waiting for an unidentified resource.</p> <p><b>SSLT</b> Time spent waiting for an SSL thread.</p> <p><b>REGT</b> Time spent waiting for a regular thread.</p> <p><b>WORK</b> Time spent waiting for registration to a work table.</p> <p><b>BPMI</b> Time spent waiting for I/O resulting from a DB2 buffer pool miss.</p>
<p>Switched Time (%)</p>	<p>Percentage of time that transactions spent routed to another region for processing. This percentage also refers to the <i>Average Response Time</i> shown in the sub-header.</p> <p>For a begin-to-end phase, the sum of these percentages should approximately equal the value shown in the <i>CONV</i> column.</p> <p>For an execution phase, these percentages, as well as the figure in the <i>CONV</i> column, are expected to be zero.</p> <p><b>LOC</b> Percentage of time that transactions spent switched on this z/OS image. Subsystems might set this state when they function ship a transaction to another component within the same z/OS image.</p> <p><b>SYS</b> Percentage of time that transactions spent switched to another z/OS image in the sysplex. Subsystems might set this state when they function ship a transaction to another component on another z/OS image within the sysplex.</p> <p><b>REM</b> Percentage of time that transactions spent switched to somewhere within the network. Subsystems might set this state when they function ship a transaction to another component within the network.</p>
<p>Address Spaces Serving this Service Class <i>srvcls</i></p>	<p><i>Srvcls</i> is the name of the <i>Class</i> at the top of the report.</p> <p>The scrollable area below this sub-header line shows a list of address spaces within the sysplex that performed work for the reported service class during the report interval. This list is created for a service class, not for a service class period. If there are reports for several periods of one service class, this section is the same for all reports.</p> <p>An address space is included in this list, when it is a "server" from WLM's point of view, regardless of whether transactions were actually being worked on during the report interval.</p>
<p>Jobname</p>	<p>Jobname of the server.</p>

Table 87. Fields in the SYSWKM Report (continued)

Field Heading	Meaning
M	A Y in this column indicates that WLM managed a server region according to the goals for the transactions being served by the region.
ASID	Address space ID of the server.
System	This is the four character SMF system identifier of the system the server is running on.
Serv-Class	Name of the service class associated with this address space.
Service	Percentage of service that the address space gives to this service class in relation to all of the other service classes it serviced during the report interval.
Proc-Usg	The TCB and SRB using percentage of the address space.
I/O-Usg	The device using percentage of the address space.
Veloc	Execution velocity of the address space.
Capp	The capping delay percentage of the address space.  Using WLM, the minimum and maximum capacity values for a resource group can be used to restrict the amount of processor capacity that a collection of address spaces is allowed to consume, or with discretionary management, if the work for which the job is running is overachieving its goal, this work may be capped in order to divert its resources to run discretionary work (see also section 'Using Discretionary Goals' in <i>z/OS MVS Planning: Workload Management</i> ).
Quies	Percentage of time for which the server address space was quiesced during the report interval.

## Report options

```

RMF Work Manager Delay Report Options           Line 1 of 6
Command ==>                                     Scroll ==> HALF

Change or verify parameters. To exit press END.
Changes will apply to SYSRTD, SYSWKM and GROUP reports.

Type      ==> S      Service or Report class (S R)
Class     ==> SYSTEM Class name
Period    ==> 1      Period number
Unit      ==> P      Unit for Display of Response Time Breakdown
                    (S=Seconds,P=Percentage)
Inactive  ==> NO     Show inactive subsystems (YES NO) in the report

Available Service and Report Classes
APPC      APPCFEED  ASCH      BATCH     BERDFEED  CICSDE    CICSHR
CICSIT    CICSOE        CICSPTS   CICSREGS  CICSSEC   CICSSTRX  DB2
DISCRETN  IMSDE         IMSHR     IMSIT     IMSOE     IMSPS     IMSREGS
IMSSC     IMSTRX        JESSTUFF  MONITORS  MVSSUBSY  ST_TOOLS  ST_USER
STCLOW    SYSOTHER     SYSSTC    SYSTEM     TPNSBATC  TPNSEVEN  TPNSFEED
TPNSODD   VEL3

```

Figure 113. SYSWKM Report Options Panel

The Report Options panel shows the RMF default options.

- Type** Here you specify whether you want to select a service or report class.
- Class** Specification of a service or report class name. If a class is entered for which there is no current data, an empty report is shown.
- Period** Specification of a period number (value between 1 and 8).
- Unit** Specify time unit:

- S The units shown are seconds.
- P The units shown are percentages.

**Inactive**

Specification about display of inactive classes:

- YES** Include subsystem data lines in the report, even if the lines do not contain any data.
- NO** Do not include subsystem data lines, if they do not contain any data.

**Available Service and Report Classes**

A scrollable area is provided containing a list of service and report class names. These names are obtained from the current data.

**Note:** The classes will be shown only if one of the sysplex reports SYSSUM, SYSRTD, or SYSWKM has been displayed at least once.

The RMF *FIND* command works on the scrollable area.

---

## USAGE - Monitor III Job USAGE Report

The Monitor III Job USAGE Report is provided as a complement to the Monitor III Job Delay Report. The USAGE report allows you to identify at a glance the jobs that are consuming the most resources within various resource categories. The report gives you information about job resource consumption in terms of I/O, processor, and storage related key metrics.

Additionally, the report provides a dedicated section with QSCAN usage statistics. This enables you to identify the jobs that have issued either the GQSCAN or the ISGQUERY REQINFO=QSCAN service.

Hence, those jobs that are issuing QSCAN requests improperly or too frequently can be detected more easily.

### How to request this report

To request the USAGE (Job Usage) report, select **1** from the Primary Menu, and then select **4A** on the Overview Report menu (shown in Figure 6 on page 25) or enter the following command using the format:

```
USAGE [ job_class, service_class ]
```

For example, to get a Usage report for TSO service class TSOPRIME, enter:

```
USAGE T, TSOPRIME
```

### Contents of the report

Figure 114 on page 207 shows a sample Job Usage report.



```

RMF V2R2 Job Oriented Usage Line 1 of 14
Command ==> Scroll ==> CSR
Samples: 60 System: SYSF Date: 09/28/16 Time: 15.35.00 Range: 60 Sec
Jobname Service --- I/O --- --- CPU --- - Storage - ----- QScan -----
CX Class Conn EXCP Total TCB Total Fixed Total Resct Time
BHBE T TSODEF 0.399 16.42 0.22 0.22 743 0 0 0 0
XCFAS S SYSTEM 0.264 4.30 0.02 0.02 13443 2943 0 0 0
*MASTER* S SYSTEM 0.041 0.27 0.05 0.01 5472 3643 0 0 0
BJAGHTM BO BTCHDEF 0.037 18.50 0.06 0.06 499 58 8 0 810
CATALOG S SYSTEM 0.037 1.58 0.02 0.02 1418 176 0 0 0
RMFGAT SO SYSSTC 0.033 0.12 0.23 0.23 14505 152 0 0 0
JES2 S SYSSTC 0.019 1.85 0.06 0.06 9390 766 0 0 0
GRS S SYSTEM 0.018 0.00 0.03 0.03 14680 454 3 0 551
CONSOLE S SYSTEM 0.000 0.03 0.01 0.01 3921 157 0 0 0
GPMSSERVE SO GPMSSERVE 0.000 0.02 0.02 0.02 3018 101 0 0 0
WLM S SYSTEM 0.000 0.00 0.10 0.10 29486 272 0 0 0
RRS S STCDEF 0.000 0.00 0.01 0.00 2813 178 0 0 0
TCPIP SO SYSSTC 0.000 0.00 0.16 0.00 8038 150 0 0 0
RMF S SYSSTC 0.000 0.00 0.01 0.01 7281 122 2 0 443

```

Figure 114. Job Usage report

## Field descriptions

Table 88. Fields in the Job Usage report

Field Heading	Meaning
Jobname	Name of the job.
CX	Abbreviation for the job class as follows:  <b>S</b> Started task <b>T</b> TSO <b>B</b> Batch <b>A</b> ASCH <b>O</b> OMVS  An O as second character indicates that the address space is using OMVS services.
Service Class	The name of the service class for this address space.
I/O Conn	Device connect time, in seconds, for this address space in the report interval.
I/O EXCP	Number of EXCP operations per second for this address space in the report interval.
CPU Total	Amount of total processor time, in seconds, for this address space in the report interval as sum of TCB time, global and local SRB time, and preemptable or client SRB time.
CPU TCB	Amount of TCB processor time, in seconds, for this address space in the report interval.
Storage Total	The number of active and idle frames for this address space, averaged over the report interval.

Table 88. Fields in the Job Usage report (continued)

Field Heading	Meaning
Storage Fixed	Number of fixed frames for this address space, averaged over the report interval.
QScan Total	Total number of QScan requests for this address space, including START and RESUME, but not QUIT requests.
QScan Resct	Average number of resources returned by QScan requests for this address space.
QScan Time	Average QScan request time, in microseconds, for this address space.

## Monitor III Utility fields

You can use the Monitor III Utility to customize the Job Usage report. In addition to the information previously shown, you can use the Utility to have the following values shown:

Table 89. Additional fields in the Job Usage report

Field Name	Meaning
JUSPASI	Address space ID (decimal)
JUSPCLA	Abbreviation for the job class as follows: <b>A</b> ASCH <b>B</b> Batch <b>O</b> OMVS <b>S</b> Started task <b>T</b> TSO
JUSPCLP	The number of the service class period for this address space.
JUSPDP	Dispatching priority for this address space.
JUSPTAT	The time that has elapsed since the current transaction in this address space was started, in the form <i>hhhh:mm</i>
JUSPTRT	The time that has elapsed since the current transaction in this address space became resident, in the form <i>hhhh:mm</i> or <i>hh:mm:ss</i>
JUSPTCT	Number of transactions for this address space since address space creation.
JUSPFRXH	Number of fixed frames above 2GB for this address space averaged over the report interval.
JUSPFRXA	Number of fixed frames between 16MB and 2GB for this address space averaged over the report interval.
JUSPFRXB	Number of fixed frames below 16MB for this address space averaged over the report interval.
JUSPDCTT	Accumulated device connect time, in seconds, for this address space since address space creation.
JUSPEXCT	Number of EXCP operations total since address space creation.
JUSPEXCD	Number of EXCP operations in this interval.

Table 89. Additional fields in the Job Usage report (continued)

Field Name	Meaning
JUSPCPUT	Accumulated amount of processor time, in seconds, for this address space since address space creation as sum of TCB time, global and local SRB time, and preemptable or client SRB time.
JUSPTCBT	Accumulated amount of TCB processor time, in seconds, for this address space since address space creation.
JUSPQSPR	Number of specific QSCAN requests for this address space, that are either GQSCAN requests specified by QNAME and RNAME or ISGQUERY requests specifying a search by ENQTOKEN.
JUSPQRSD	Standard deviation for the number of resources returned by QSCAN requests for this address space.
JUSPQTSD	Standard deviation for the QSCAN request time for this address space.

### Cursor-sensitive control on the USAGE Report

Table 90 shows which report is displayed for each cursor-sensitive field.

Table 90. Cursor sensitivity on Job Usage report

Report column	Displayed report
Jobname	DELAYJ Report for selected job
CX	USAGE Report filtered by job class
Service Class	USAGE Report filtered by service class
I/O Conn	DEV Report
I/O EXCP	DEV Report
CPU Total	PROCU Report
CPU TCB	PROCU Report
Storage Total	STORF Report
Storage Fixed	STORF Report

---

## WFEX - Workflow/Exceptions Report

The Workflow/Exceptions (WFEX) report presents information about system activity and system resources.

The top part of the report shows you speedometers in graphic mode, or workflow indicators in tabular mode.

The color of a workflow indicator tells you how well the jobs are performing depending on the exception criteria specified. (Usually, red indicates a problem, yellow indicates caution, and turquoise indicates that a job or volume is missing from the system configuration.)

The speedometer needle points to the relative speed of the job or resource in the system. The shaded part to the left of the needle represents the proportion of a user's time spent doing useful work. The part to the right of the needle represents the proportion of a user's time spent delayed.

You can define the exception criteria on the Definition and Criteria panel.

A line in the Exceptions section of the report corresponds to each colored speedometer or workload indicator. The line has the same color and the same Name as the speedometer or workload indicator, and gives details about the exception.

### Speed (Workflow)

Under **Speed (Workflow)**, a high workflow percentage or speed indicates that a job has the resources it needs to process, and that it is moving through the system at a relatively high speed.

If the Criteria Set you defined for a workflow indicator is met, it will be displayed in the corresponding color.

A low value under Speed indicates that a job has few of the resources it needs and is contending with other jobs for system resources and may indicate a problem.

For resources (for example \*PROC and \*DEV), a high value under Speed indicates that jobs are moving through the system with little resource contention. A low Speed represents a large queue of work requests from users and may indicate a problem.

### Exceptions

In the **Exceptions** section of the report, a line that is displayed in yellow or red indicates a job or resource exceeds the exception criteria defined by you on the **WFEX Report Options: Definition and Criteria** panel. Use cursor-sensitive control to investigate exceptions further.

If **Not avail** appears on your report, the job you selected on the Definition panel was not running during the report interval. If **No work** appears, the job or group was idle (not requesting system resources) during the report interval.

## How to request this report

To request the Workflow/Exceptions, select **1** from the Primary Menu, then select **1** from the Overview Report menu (shown in Figure 6 on page 25) or enter the following command:

WFEX

## Contents of the report

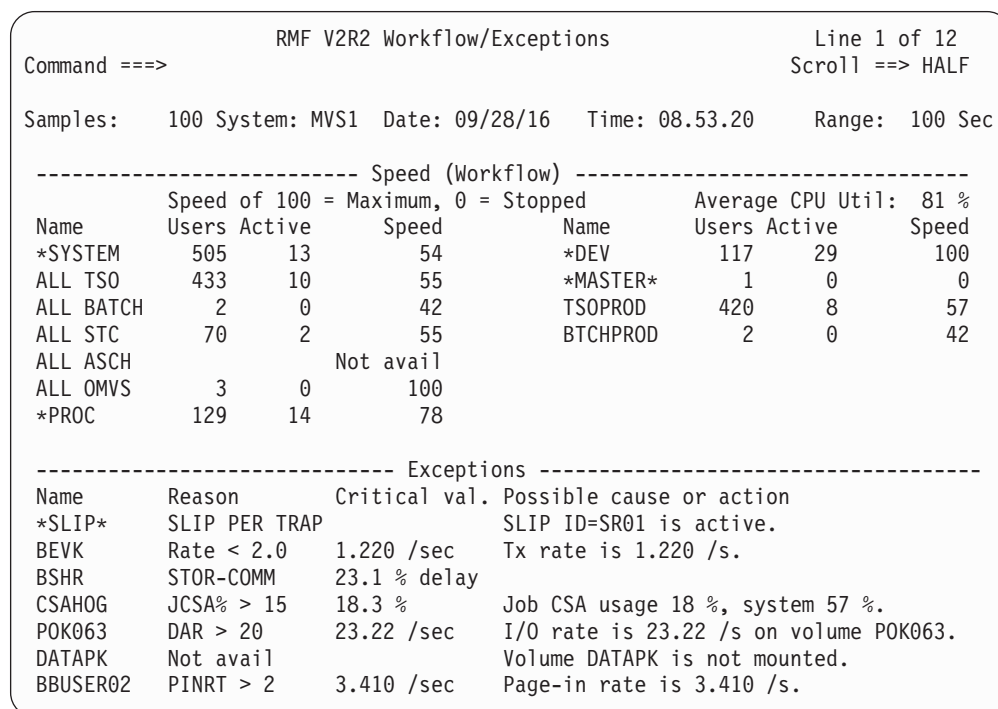


Figure 115. WFEX Report

The report has two parts:

- On the top **Speed** section, RMF reports the workflow of jobs and resources as speed relative to the maximum speed with which they could move through the system.
- On the bottom **Exceptions** section, RMF lists jobs, job groups, or system resources that meet exception criteria.

The workflow and exception lines are color coded according to severity. Usually, red indicates a problem, yellow indicates caution, and turquoise indicates that a job or volume is missing from the system configuration. You can specify exception criteria on the Workflow/Exceptions Report Options panels, or you can use automatic customization.

Figure 115 is an example of the Workflow/Exceptions report. For information about the WFEX Report Options panels, see "Report options" on page 216. For a Workflow/Exceptions report based on your installation's requirements for workload, you can use automatic customization. For information about automatic customization, see "Automatic customization" on page 224.

**Workflow of jobs or job groups** is a measure of the speed at which jobs are moving through the system in relation to the maximum speed at which the jobs

could move through the system. These workflow formulas are described in “Common Monitor III report measurements” on page 12.

A low workflow percentage indicates that a job has few of the resources it needs and is contending with other jobs for system resources. A high workflow percentage indicates that a job has the resources it needs to execute, and that it is moving through the system at a relatively high speed.

For example, a job that could execute in one minute, if all the resources that it needed were available, would have a workflow of 25%, if it took four minutes to execute.

Workflow of resources (processors or devices) represents how well the system is serving users. The speed at which each resource performs the work of all user's is expressed as a value from 0% to 100%. A low resource workflow percentage represents a large queue of work requests from users. A high workflow percentage represents little resource contention.

If **Not avail** appears on your report, the job that you selected on the Definition and Criteria panel was not running during the report interval. If **No work** appears, the job or job group was idle (not requesting system resources) during the report interval.

Exceptions are suppressed without notification when using criteria on historic RMF gatherer records that do not provide the corresponding data. However, if there is at least one valid criteria set defined, and the thresholds are met, the exception is displayed.

### Field descriptions — Speed Section

Table 91. Fields in the Speed Section of the WFEX Report

Field Heading	Meaning
Average CPU Util	The average utilization percentage for all general purpose processors (CPs) during the report interval. For details, see Table 79 on page 178.
Name	The one to ten character identifier of a workflow indicator. It can be a job, job group, or resource (processor or device). You can specify Name on the Label field of the Definition and Criteria panel or leave it blank and use the default name generated by RMF. If an indicator changes color, there is a corresponding line in the Exceptions section of the report with the same name and color giving more information about the exception.
Users/ Active	The average number of users and the average number of active users in an address space or group of address spaces.  The average number of active users is a measure of system workload. A user in a system is either ACTIVE, IDLE, or unknown during a report interval. An ACTIVE user is using a resource or is delayed by a resource. An IDLE user is in terminal wait, timer wait, or is waiting for JES job selection. A user that is not in either of these states is unknown.
Speed	For jobs and job groups, Speed is a measure of the speed at which jobs are moving through the system in relation to the maximum speed at which the jobs could move through the system.  A low workflow percentage indicates that a job has few of the resources it needs and is contending with other jobs for system resources. A high workflow percentage indicates that a job has the resources it needs to execute, and that it is moving through the system at a relatively high speed.  For resources (processors or devices), Speed represents how well the system is serving the users. A low resource workflow percentage represents a large queue of work requests from users. A high workflow percentage represents little resource contention.

## Field descriptions — Exceptions Section

The Exceptions section of the report shows the exceptional situations a job or job group encountered in relation to the hardware and software resources. The exceptions are those specified on the Workflow/Exceptions Report Options panels.

Table 92. Fields in the Exceptions Section of the WFEX Report

Field Heading	Meaning
Name	<p>The one to ten character identifier of a workflow indicator. It can be a job, job group, or resource (processor or device). You can specify Name on the Label field of the Definition and Criteria panel or leave it blank and use the default name generated by RMF. If a threshold from the <b>Definition and Criteria</b> options panel is exceeded, one or more lines in the Exceptions section are shown with a name from the Label field, a specific job name, or resource name.</p>
Reason	<p>Reason gives the explanation for the exception condition that was defined either on the WFEX Report Options panel or by automatic customization. The reason field can show:</p> <ul style="list-style-type: none"> <li>• A resource contributing most to the delay of a job or job group (main delay) <p>A main delay is displayed when either an exception is met that includes a mixture of criteria, or a general using or delay exception is met (i.e. AAU, ADU, USG%, USGD%, USGP%, or WFL%).</p> </li> <li>• A reason meeting a specified exception criterion <p>A specific reason is displayed when a single delay exception is met, even if another resource contributed more to the overall delay.</p> <p>For example -</p> <p style="padding-left: 40px;">If you specify ENQ% &gt; 10 for jobname xxx, then the Reason field displays <b>ENQ-majorname</b> where majorname is the name of the enqueue resource that is causing the delay.</p> <p style="padding-left: 40px;">If you specify COMM% &gt; 20 for jobname xxx, then the Reason field displays <b>STOR-COMM</b>.</p> </li> <li>• An exception <p>The exception statement is displayed when a STOR class exception or another type of single exception that is not delay-oriented is met.</p> <p>For example -</p> <p style="padding-left: 40px;">If you specify CPUS% &gt; 60, then the Reason field displays <b>CPUS% &gt; 60</b>.</p> </li> <li>• SLIP PER TRAP <p>This exception is displayed if a SLIP PER trap is active on your system. The exception line is always reported first and is displayed in yellow. The yellow color is a warning that an active SLIP PER trap can cause performance degradation and should be removed. Note that you cannot exclude or change the color of the SLIP PER trap exception line.</p> </li> </ul> <p>If the <b>Reason</b> field displays an exception statement or the SLIP PER TRAP exception, the field is not split by a hyphen.</p> <p>However, the <b>Reason</b> column is split into two at the hyphen when a main delay or a specific reason for the exception is displayed.</p> <p>The left part of the column depicts the resource contributing to the exception condition.</p> <p>The resource displayed is either the resource contributing most to the overall delay, or the resource that has been specifically defined (single exception).</p>

Table 92. Fields in the Exceptions Section of the WFEX Report (continued)

Field Heading	Meaning
Reason (continued)	<p>The right part of the column depicts the reason for the exception condition. The reason depends on which resource caused it.</p> <p>If the resource under <b>Reason</b> is:</p> <p><b>PROC</b> then the name of the job using the processor most often when the delayed job wanted to use it, appears to the right of the hyphen. If one or more enclaves contributed most to the processor delay, then *ENCLAVE will be shown under <b>Reason</b>.</p> <p><b>DEV</b> then the volume serial number of the device that the job was most often delayed for, appears to the right of the hyphen.</p> <p><b>ENQ</b> then the major name of the serially reusable resource causing the greatest percentage of delay, appears to the right of the hyphen.</p> <p><b>STOR</b> then either COMM, LOCL, VIO, SWAP, or OUTR appears to the right of the hyphen.</p> <p>If you requested a COMM%, LOCL%, VIO%, SWAP%, or OUTR% single exception, the possible causes are:</p> <p><b>COMM</b> Common storage paging  <b>LOCL</b> Local storage paging  <b>VIO</b> Virtual I/O paging  <b>SWAP</b> Swap-in delay  <b>OUTR</b> Swapped out and ready.  <b>XMEM</b> Cross memory address space paging  <b>HIPR</b> Standard hiperspace paging</p> <p>If STOR is the main delay, or a single STOR% exception is met, only <b>STOR</b> appears under <b>Reason</b>. The storage delay reason appears as an informational message under <b>Possible cause or action</b>.</p> <p>Note that the informational message can be overwritten by a user specified message.</p> <p>For single jobs, 'Main reason SSSS causes xxx % delay' appears under <b>Possible cause or action</b>.</p> <p>For job groups, 'Main reason SSSS delays xxx users' appears under <b>Possible cause or action</b>.</p> <p><b>Note:</b> In both cases, SSSS is either COMM, LOCL, VIO, SWAP, XMEM, HIPR, or OUTR.</p> <p><b>OPER</b> then either Message or Mount or Quiesce appears to the right of the hyphen. Message indicates that the operator did not respond to a message. Mount indicates that the operator did not mount a tape. Quiesce indicates that the address space was quiesced by the operator.</p> <p><b>SUBS</b> then either JES, HSM, or XCF appears to the right of the hyphen.</p> <p>Cursor-sensitive control is split at the hyphen of the <b>Reason</b> column. If you press ENTER with the cursor positioned either under the left or under the right part of the hyphen, you can get more information about the exception condition.</p>



Table 92. Fields in the Exceptions Section of the WFEX Report (continued)

Field Heading	Meaning
Critical val.	<p>The critical value that caused the exception condition.</p> <p>The following values can appear under <b>Critical val</b>:</p> <p><b>% delay</b> Percentage of delay caused by the delay category in the <b>Reason</b> field.</p> <p><b>users</b> Average number of delayed users for the group or resource, or maximum number of users.</p> <p><b>%</b> Using percentage in the case of single STOR class exceptions, single common storage exceptions, device class exceptions, CPU% exceptions, or CPUS% exceptions.</p> <p><b>frames</b> Number of frames online in the case of ONLF and ONLXF exceptions.</p> <p><b>bytes</b> Number of bytes of storage used in the case of TSQAO exceptions.</p> <p><b>/sec</b> Rate per second in the case of DAR, RATE, PINRT, ESMRT, or ESPRT exceptions.</p> <p><b>sec</b> Amount of time in seconds in the case of time-related exceptions (i.e. AT, DRT, QT, RT, TET, and ESMAG).</p> <p><b>replies</b> Number of outstanding replies.</p> <p>If the exception resource or user is unavailable on the system, the Critical val. field remains blank.</p> <p>For the formula used to calculate the delay of an address space or group of address spaces see "Address space delay (%)" on page 14.</p>
Possible Cause or Action	<p>Possible Cause or Action describes what might be causing the delay and what you can do about it. You can enter the text for Possible Cause or Action in the Text field on the Definition and Criteria panel of the Workflow/Exceptions Report Options panels, or you can let RMF fill in the text according to some analysis of what the problem was.</p> <p>If the Possible Cause or Action field is blank, use cursor-sensitive control on the Name or Reason field for more information about the delay.</p>

## Report options

```

RMF WFEX Report Options: Action Panel                               Line 1 of 23
Command ==>>>                                                    Scroll ==>> HALF

Enter Action Code in the Action Column. To exit press END.
Action Codes: Select (S) Copy (C) Move (M) Before (B)
              Add (AD) Delete (D) Move Block (MM) After (A)

Action  Class  Qualifier  Indicator  Label          Row  Position
-----  ---
Only Add (AD) and After (A) are valid on this line.
-----  ---
SYSTEM                WF              1      1
TSO                   WF              1      2
STC                   WF              1      3
BATCH  ALL            WF              1      4
ASCH                  WF              1      5
OMVS                  WF              1      6
PROC                  WF              1      7
DEV  ALL              WF              2      1
JOB  *MASTER*         WF              2      2
SRVCLS TSOPROD       WF              2      3
SRVCLS BTCHPROD      WF              2      4
SYSTEM                EX-ANY
JOB                   EX-UNAVAIL
STOR                   EX-AVG
STOR                   EX-AVG
STOR                   EX-AVG      *STOR
    
```

Figure 116. WFEX Report Options Action Panel

To set up your workflow indicators and exceptions, you can either use automatic customization, or you can use the Report Options panels. When the displayed value in the WFEX report meets the specified threshold value, the exception or workflow indicator is highlighted to your specifications.

Automatic customization sets up workflow and exception indicators, threshold values, and highlighting criteria based on your installation's specifications for workload. To use automatic customization, see "Automatic customization" on page 224.

On the Report Options panels, you can specify workflow indicators, exception conditions, or a combination of both. However, on the **Speed** section of the report, a limit of 14 workflow indicators can be displayed. You can also specify the threshold values and color highlighting criteria for the indicators.

RMF displays the Action panel. The Action panel shows a list of the Report Options currently in effect. On the Action panel, you can enter codes in the **Action** column to change, add, and delete workflow and exception indicators, and vary where the indicators will appear in the report.

To add, view, or change criteria, enter the Add (**AD**) or Select (**S**) code in the **Action** column. RMF then displays the Definition and Criteria panel.

| To reset the Report Options to the RMF default values, if customization is set to  
 | NO in your Session Options, enter the RESET command.

Table 93. Fields in the WFEX Action Panel

Field Heading	Meaning
Action	In the Action column, you can enter the following commands: <b>S</b> Access the Definition and Criteria panel <b>AD</b> Add a workflow or exception indicator <b>C</b> Copy a workflow or exception indicator <b>D</b> Delete a workflow or exception indicator <b>M</b> Move a line <b>MM</b> Move a block <b>B</b> Place the moved block or line or copied line on the preceding line <b>A</b> Place the moved block or line or copied line on the following line.
Class	Class is the class of resources for exception.
Qualifier	Qualifier further identifies the class and can be a service class name, job name, volume, or JES initiator class.
Indicator	Indicator is the type of indicator you want RMF to check for. There can be either workflow (WF) or exception (EX-ANY, EX-AVG, EX-GROUP, or EX-UNAVAIL) indicators. You can specify 14 workflow indicators for display in the Speed section of the report. If you specify more than 14, RMF checks only the exception condition specified together with the workflow indicator.
Label	Label is a 10 character identifier of a job or job group, or a resource that you want to appear as <b>Name</b> on the Workflow/Exceptions (WFEX) report.  For an EX-ANY option, the <b>Name</b> field in the WFEX report always contains the name of the job being delayed, rather than the label specified on the option panel.
Row and Position	Row and Position identify the location of the workflow indicators on the Speed section of the Workflow/Exceptions report.  Row and Position are arranged differently on the graphic and tabular WFEX reports. If <b>Not Displayed</b> appears under <b>Row</b> and <b>Position</b> , that workflow indicator will not appear in the Speed section of the report, but an exception condition that is specified together with the workflow indicator can appear in the Exceptions section of the report. Use the HELP key (PF1) for information about how to change where the workflow indicators appear on the WFEX report.

### Definition and Criteria panel

```

RMF WFEX Report Options: Definition and Criteria
Command ==>                               Scroll ==> HALF

Enter or edit information below. To view a list of criteria name values,
place the cursor in a blank "Name" field and press ENTER.
Exception will be displayed if all criteria of one color in a set are met.

Class      ==> _____ For example: SYSTEM, BATCH, JOB, DEV, STC, SRVCLS
Qualifier  ==> _____ For example: Jobname, volume serial, job class
Indicator  ==> _____ WF, EX-ANY, EX-AVG, EX-GROUP or EX-UNAVAIL
Label      ==> _____ Label for workflow monitor or exception line
Alert      ==> _____ Alerting signal: BLINK, BEEP, BOTH, NONE
Text       ==> _____ Leave blank for default

      Criteria set 1           Criteria set 2           Criteria set 3
Name  <>  Yel  Red           Name  <>  Yel  Red           Name  <>  Yel  Red
_____ < > _____ _____ or _____ < > _____ _____ or _____ < > _____ _____
_____ < > _____ _____ or _____ < > _____ _____ or _____ < > _____ _____
_____ < > _____ _____ or _____ < > _____ _____ or _____ < > _____ _____
_____ < > _____ _____ or _____ < > _____ _____ or _____ < > _____ _____
_____ < > _____ _____ or _____ < > _____ _____ or _____ < > _____ _____
_____ < > _____ _____ or _____ < > _____ _____ or _____ < > _____ _____
_____ < > _____ _____ or _____ < > _____ _____ or _____ < > _____ _____

```

Figure 117. WFEX Definition and Criteria Panel

## Mon III - WFEX

On this panel, you modify the report by defining or changing workflow indicators and exception conditions.

In the top half of the panel, provide information about the job or job group, or resource.

In the bottom half of the panel, fill in exception values and highlighting criteria, or choose volumes or job names. You can use cursor-sensitive control on the **Name** field. The corresponding Criteria Names Selection panel is displayed.

To exit this panel, you must either:

- Specify a complete workflow indicator or exception condition
- Use the CANCEL command to cancel any input.

Table 94. Fields in the WFEX Definition and Criteria Panel

Field Heading	Meaning
Class	<p>You specify the class of resources for exception in the Class field. Class can be:</p> <p><b>SYSTEM</b> All jobs in the system</p> <p><b>TSO</b> All TSO/E users</p> <p><b>BATCH</b> All batch jobs</p> <p><b>STC</b> All started tasks</p> <p><b>JOB</b> Single job by name</p> <p><b>PROC</b> Processor</p> <p><b>DEV</b> Device</p> <p><b>STOR</b> Storage</p> <p><b>ASCH</b> ASCH address space</p> <p><b>OMVS</b> OMVS address space</p> <p><b>SRVCLS</b> Service class</p>
Qualifier	<p>In the Qualifier field, you specify qualifiers for the following exception classes:</p> <p><b>BATCH</b> Batch jobs - specify the job class (JES initiator class). <b>BATCH</b> does not require a qualifier, but if you leave <b>Qualifier</b> blank, RMF fills in the default ALL.</p> <p><b>JOB</b> Single job by name - specify a jobname. If indicator EX-UNAVAIL is specified, no qualifier is allowed.</p> <p><b>DEV</b> Device - specify a device number. <b>DEV</b> does not require a qualifier, but if you leave <b>Qualifier</b> blank, RMF fills in the default ALL. If indicator EX-UNAVAIL is specified, no qualifier is allowed.</p> <p><b>SRVCLS</b> Jobs grouped by service class - specify a valid service class name.</p> <p><b>Wildcard Support:</b> You can specify a wildcard character at the end of the input string for the qualifier of the classes JOB and DEV.</p> <p>An asterisk "*" in the last position is not treated as part of the name, instead each name that matches the input string up to the position of the asterisk is treated as if it had been specified in that input field. This allows the specification of a WFEX exception that gives an exception line for each DASD, TAPE or JOB belonging to a group with names starting with identical characters.</p>

Table 94. Fields in the WFEX Definition and Criteria Panel (continued)

Field Heading	Meaning
Indicator	<p>In the Indicator field, you specify the workflow or exception indicator type.</p> <p>For Indicator, you can define the following:</p> <p><b>WF</b> Indicates a workflow indicator which is a permanent line or speedometer in the top part of the WFEX report. If you add exception conditions to indicator WF, they are treated as if the indicator was EX-AVG.</p> <p><b>EX-ANY</b> Indicates that RMF reports one line per address space that meets the condition you specify under Criteria Set.</p> <p><b>EX-AVG</b> Indicates that RMF reports an exception when the average value of the specified group fulfills the conditions you specify under Criteria Set.</p> <p><b>EX-GROUP</b> Indicates that RMF checks for an exception for each job or resource in the specified group but reports only one line in the exception report containing the main delay reason and the number of users.</p> <p>If you need information about resource problems that are causing significant delays, but do not need to know exactly which users are affected, use GROUP.</p> <p><b>EX-UNAVAIL</b> Indicates that RMF reports an exception when the volume or jobname is unavailable. When you specify EX-UNAVAIL, RMF displays the Definition of UNAVAIL panel. See Figure 118 on page 220 for a description of the panel.</p>
Label	<p>In the Label field, you specify a label for the workflow and exception indicators.</p> <p>Label is a 10-character identifier of a job or job group, or a resource that you want to appear as Name on the Workflow/Exceptions (WFEX) report.</p> <p>Label is optional. If you do not enter a Label, the field remains blank on the Definition and Criteria panel and on the Action Panel, but RMF dynamically fills it in on the report.</p>
Alert	<p>In the Alert field, you specify the warning signal for the workflow indicator or the exception line.</p> <p>When the threshold values you specified in the &lt;&gt;, <b>Yel</b>, and <b>Red</b> columns are met, RMF informs you through the alert signal. For Alert, you can specify:</p> <p><b>BLINK</b> The workflow indicator on the tabular report or the exception line on both the tabular and graphic reports blinks</p> <p><b>BEEP</b> Your workstation beeps</p> <p><b>BOTH</b> Your workstation beeps, and the workflow indicator on the tabular report or exception line on both the graphic and tabular reports blinks</p> <p><b>NONE</b> No alerting signal.</p>
Text	<p>In the Text field, you can specify the text for the <b>Possible Cause or Action</b> field of the Exceptions section of the Workflow/Exceptions report.</p> <p><b>Text</b> is optional. If you leave it blank, RMF either:</p> <ul style="list-style-type: none"> <li>• Dynamically fills it in with additional information or with a suggestion of what to do</li> <li>• Leaves it blank when no information is available.</li> </ul>
Criteria Set	<p>Use the <b>Criteria Set</b> to specify the exception conditions and color highlighting that RMF should check.</p> <p>If you are defining a workflow indicator, the <b>Criteria Sets</b> are optional.</p> <p>If you are defining an exception (EX-ANY, EX-AVG, EX-GROUP), you must enter at least one <b>Name</b>, one comparison operand (&lt;&gt;), and one threshold value (<b>Yel</b> or <b>Red</b>) in one criteria set.</p>

Table 94. Fields in the WFEX Definition and Criteria Panel (continued)

Field Heading	Meaning
Name	<p>You enter the conditions that RMF is to check for under <b>Name</b>. You can fill in Name or select the criteria names from a panel. To access the selection panel, use cursor-sensitive control on a blank <b>Name</b> field. To specify a criteria name, place an S in the column next to your choice and return to the Definition and Criteria panel. See "Criteria Names Selection panel" on page 221 for more information.</p> <p>You can specify or select up to seven criteria names for each criteria set.</p> <p>All of the conditions (Names) defined within a Criteria Set must be met in order for the color highlighting to appear on the Workflow/Exceptions report.</p> <p>To delete a criteria name, on the selection panel, you can either:</p> <ul style="list-style-type: none"> <li>• Blank out the S next to the criteria name you want to delete on the Criteria Names Selection panel.</li> <li>• Blank out the threshold values in the <b>Yel</b> and <b>Red</b> fields on the Definition and Criteria panel. When you press END, RMF removes the whole criteria line.</li> </ul>
<>	<p>Specify a comparative operator in the &lt;&gt; field. You can specify one of the following:</p> <p>&lt; &gt; &lt;= &gt;= = != &gt;= &lt;= &lt;&gt; GT LT GE LE EQ NE NG NL</p>
Yel Red	<p>Specify a numeric threshold value in the <b>Yel</b> and <b>Red</b> fields. The values are color coded according to severity. The red value that you specify indicates a problem. The yellow value indicates caution (there may be a problem). If you define Yel and Red to be the same value, red has precedence. You can specify zero or any positive integer with or without a decimal point. For criteria names that represent a percentage, you must specify a value from 0 to 100.</p> <p>When the condition meets the value that you specify, exception lines appear in the Exceptions section of the report in the corresponding color. If the exception condition is also associated with a workflow indicator, the permanent line (in the tabular report) or the speedometer (in the graphic report) changes to the corresponding color in the Speed (Workflow) section.</p>

### Definition of UNAVAIL panel

```

RMF WFEX Report Options: Definition of UNAVAIL           Line 1 of 2

Edit information below. Use action characters to specify color of exception
message. Exceptions are displayed when specified volumes are not available.
Action characters: Turquoise (T)   Yellow (Y)   Red (R)

Class          DEV          Devices in the system
Qualifier
Indicator      EX-UNAVAIL  Exception when any selected volume not available
Label         ==>>> _____  Comment for identification
Alert         ==>>> NONE_    Alerting signal: BLINK, BEEP, BOTH, NONE
Text          ==>>> _____  Leave blank for default

A Volume      A Volume      A Volume      A Volume      A Volume      A Volume
Y PAGE08      T SPOOL1      T DRV021      T D13IOG      T D13MCP      T D94RM1
T 410PRM      T 410SRO
    
```

Figure 118. WFEX Definition of UNAVAIL Panel

On this panel you modify the report by defining or changing exception conditions. The above sample shows a panel for the class DEV.

Specify Label, Alert, and Text information on the top half of the panel.

**Note:** Class, Qualifier, and Indicator cannot be changed on this panel.

Select or type in the volume and color highlighting in the bottom half of the panel. Whenever the selected volume is not mounted, it will appear as an exception in the report highlighted in your chosen color.

To exit this panel, you must either:

- Select or type in at least one volume and press the END key
- Use the CANCEL command to cancel any volume selections.

On this panel, you can enter one of the following codes in the **A** column to color highlight exception lines.

**T** Turquoise highlighting  
**Y** Yellow highlighting  
**R** Red highlighting

You can also type in any volume on a blank line at the top of the list of volumes.

If you want to remove a volume from this list, just specify a **blank** in column **A**.

### Criteria Names Selection panel

On any of the six Criteria Names Selection panels you can modify the report by selecting or changing exception conditions.

Each selection panel shows a complete list of exceptions allowed for the specified Class, Qualifier, and Indicator. You can select up to 7 criteria names. The selected criteria names are shown in the corresponding criteria set when you return to the Definition and Criteria panel.

The six Criteria Names Selection panels are:

- Criteria selection panel for jobs
- Criteria selection panel for service classes, ASCH, and OMVS
- Criteria selection panel for processors
- Criteria selection panel for report performance groups
- Criteria selection panel for storage
- Criteria selection panel for devices

To get to the Criteria Names Selection panel, use cursor-sensitive control on a blank Name field on the Definition and Criteria panel.

All of the conditions (Names) defined within a Criteria Set must be met in order for the color highlighting to appear on the Workflow/Exceptions report. The following figures show sample Criteria Names Selection panels, all panels are scrollable.

```

RMF WFEX Criteria Names for Class: JOB
Command ==>
Select (S) a maximum of 7 items to use in a criteria set. Press END.

More: +
_ AAU Average active users in group _ MSG% Operator message delay %
_ AAUS Avg active users in system _ ONLF Online real storage frames
_ ADU Average delayed users _ ONLXF Online expanded storage frames
_ ASTO% Active storage % _ OPER% Operator delay %
_ AUU Average using users _ OREPL Outstanding replies
_ AVAIL Number of jobs in group _ OUTR% Out/ready delay %
_ COMM% Common storage delay % _ PINRT Page-in rate
_ CPU% TCB+SRB % _ PROC% Processor delay %
_ CPUS% CPU utilization % _ QUI% Quiesced delay %
_ CSA% CSA storage % _ RATE Transaction rate
_ DEV% Device delay % _ SCSA% System CSA use %
_ DLY% Overall delay % _ SECS% System ECSA use %
_ ENQ% Enqueue delay % _ SESQ% System ESQA use %
_ ESMAG Exp storage migration age _ SQA% SQA storage %
_ ESMRT Exp storage migration rate _ SQA0% SQA overflow %
_ ESPRT Exp storage page movement rate _ SSQA% System SQA use %
_ ESQO% ESQA overflow % _ STOR% Storage delay %
_ EWSET Expanded storage WSET frames _ SUBS% SUBS delay %
_ HIPR% Hiperspace storage delay % _ SWAP% Swap delay %
_ HSM% HSM delay % _ TET Transaction elapsed time
_ ISTO% Idle storage % _ TSQA0 Total SQA overflow
_ JCSA% Job CSA use % _ UIC Unreferenced interval count
_ JECS% Job ECSA use % _ USG% Overall using %
_ JES% JES delay % _ USGD% Device using %
_ JESQ% Job ESQA use % _ USGP% Processor using %
_ JSQA% Job SQA use % _ VIO% VIO storage delay %
_ LOCL% Local storage delay % _ WFL% Workflow/Speed %
_ LPA% LPA storage % _ WFLG% Group workflow %
_ MAXU% Maximum allowed user % _ XCF% XCF delay %
_ MNT% Tape mount delay % _ XMEM% Cross-memory storage delay %

```

Figure 119. WFEX Criteria Names for Class: Job



```

RMF WFEX Criteria Names for Class: SRVCLS
Command ==>

Select (S) a maximum of 7 items to use in a criteria set. Press END.

More: +
- AAU Average active users in group - ONLF Online real storage frames
- AAUS Avg active users in system - ONLXF Online expanded storage frames
- ADU Average delayed users - OPER% Operator delay %
- ASTO% Active storage % - OREPL Outstanding replies
- AT Execute response time - OUTR% Out/ready delay %
- AUU Average using users - PINRT Page-in rate
- AVAIL Number of jobs in group - PROC% Processor delay %
- COMM% Common storage delay % - QT Queue time (ASCH)
- CPU% TCB+SRB % - QUI% Quiesced delay %
- CPUS% CPU utilization % - RATE Transaction rate
- CSA% CSA storage % - RT Response time total
- DEV% Device delay % - SCSA% System CSA use %
- DLY% Overall delay % - SECS% System ECSA use %
- ENQ% Enqueue delay % - SESQ% System ESQA use %
- ESMAG Exp storage migration age - SQA% SQA storage %
- ESMRT Exp storage migration rate - SQA0% SQA overflow %
- ESPRT Exp storage page movement rate - SSQA% System SQA use %
- ESQO% ESQA overflow % - STOR% Storage delay %
- EWSET Expanded storage WSET frames - SUBS% SUBS delay %
- HIPR% Hiperspace storage delay % - SWAP% Swap delay %
- HSM% HSM delay % - TET Transaction elapsed time
- ISTO% Idle storage % - TSQA0 Total SQA overflow
- JCSA% Job CSA use % - UIC Unreferenced interval count
- JECS% Job ECSA use % - USG% Overall using %
- JES% JES delay % - USGD% Device using %
- JESQ% Job ESQA use % - USGP% Processor using %
- JSQA% Job SQA use % - VIO% VIO storage delay %
- LOCL% Local storage delay % - WFL% Workflow/Speed %
- LPA% LPA storage % - WFLG% Group workflow %
- MAXU% Maximum allowed user % - XCF% XCF delay %
- MNT% Tape mount delay % - XMEM% Cross-memory storage delay %
- MSG% Operator message delay %

```

Figure 120. WFEX Criteria Names for Class: Service Class

```

RMF WFEX Criteria Names for Class: PROC
Command ==>

Select (S) a maximum of 7 items to use in a criteria set. Press END.

- AAU Average active users in group - LPA% LPA storage %
- AAUS Avg active users in system - ONLF Online real storage frames
- ADU Average delayed users - ONLXF Online expanded storage frames
- ASTO% Active storage % - SCSA% System CSA use %
- AUU Average using users - SECS% System ECSA use %
- AVAIL Number of CPUs - SESQ% System ESQA use %
- CPUS% CPU utilization % - SQA% SQA storage %
- CSA% CSA storage % - SQA0% SQA overflow %
- ESMAG Exp storage migration age - SSQA% System SQA use %
- ESMRT Exp storage migration rate - TSQA0 Total SQA overflow
- ESPRT Exp storage page movement rate - UIC Unreferenced interval count
- ESQO% ESQA overflow % - WFL% Workflow/Speed %
- ISTO% Idle storage %

```

Figure 121. WFEX Criteria Names for Class: Processor

```

RMF WFEX Criteria Names for Class: STOR
Command ==>
Select (S) a maximum of 7 items to use in a criteria set. Press END.

_ ASTO% Active storage %           _ ONLXF Online expanded storage frames
_ CSA%  CSA storage %             _ SCSA% System CSA use %
_ ESMAG Exp storage migration age  _ SECS% System ECSA use %
_ ESMRT Exp storage migration rate _ SESQ% System ESQA use %
_ ESPRT Exp storage page movement rate _ SQA% SQA storage %
_ ESQO% ESQA overflow %          _ SQA0% SQA overflow %
_ ISTO% Idle storage %           _ SSQA% System SQA use %
_ LPA%  LPA storage %            _ TSQA0 Total SQA overflow
_ ONLF  Online real storage frames _ UIC  Unreferenced interval count

```

Figure 122. WFEX Criteria Names for Class: Storage

```

RMF WFEX Criteria Names for Class: DEV
Command ==>
Select (S) a maximum of 7 items to use in a criteria set. Press END.

_ AAU  Average active users in group _ ESQO% ESQA overflow %
_ AAUS Avg active users in system    _ ISTO% Idle storage %
_ ACT% Device active time %          _ LPA%  LPA storage %
_ ADU  Average delayed users         _ ONLF  Online real storage frames
_ ASTO% Active storage %             _ ONLXF Online expanded storage frames
_ AUU  Average using users           _ PND%  Device pending time %
_ AVAIL Number of devices            _ SCSA% System CSA use %
_ CON% Device connect time %         _ SECS% System ECSA use %
_ CPUS% CPU utilization %            _ SESQ% System ESQA use %
_ CSA%  CSA storage %                _ SQA%  SQA storage %
_ DAR  Device activity rate           _ SQA0% SQA overflow %
_ DRT  Device response time          _ SSQA% System SQA use %
_ DSC% Device disconnect time %      _ TSQA0 Total SQA overflow
_ ESMAG Exp storage migration age    _ UIC  Unreferenced interval count
_ ESMRT Exp storage migration rate   _ WFL%  Workflow/Speed %
_ ESPRT Exp storage page movement rate

```

Figure 123. WFEX Criteria Names for Class: Device

To select a criterion name, type an **S** next to the **Name(s)** you want to use in a criteria set and press the END key.

Each **Name** represents a condition that RMF checks for as an exceptional value.

You can select a maximum of seven names for a criteria set.

To delete a criterion name, you can either:

- Blank out the **S** next to the criterion name you want to delete on the Criteria Names Selection panel.
- Blank out the threshold values in the **Yel** and **Red** fields on the Definition and Criteria panel. When you press END, RMF removes the whole criterion line.

### Automatic customization

To automatically customize RMF option sets, specify YES for Customization on the Session Options panel.

Automatic customization ensures that the option set used matches the service policy name and system ID that was in effect when the data was gathered. If the service policy name or the system ID changes between reports, an option set with

the same name is made active. The new option set is listed under Current Option Set on the Option Set Selection Menu (invoked via command OPTSET).

If no option set exists with the same name as the service policy name and system ID associated with the data, a new option set with that name is generated from the current option set and made active.

Automatic customization is turned off if you enter NO in the Customization field on the Session Options panel or if you change the Current Option Set on the Option Set Selection menu to an option set that does not match the service policy name and system ID of the data at which you are looking.

Every time a new option set is selected or created, RMF saves all the changes from the old option set before making the new option set current.

## Workflow/Exceptions graphic Report

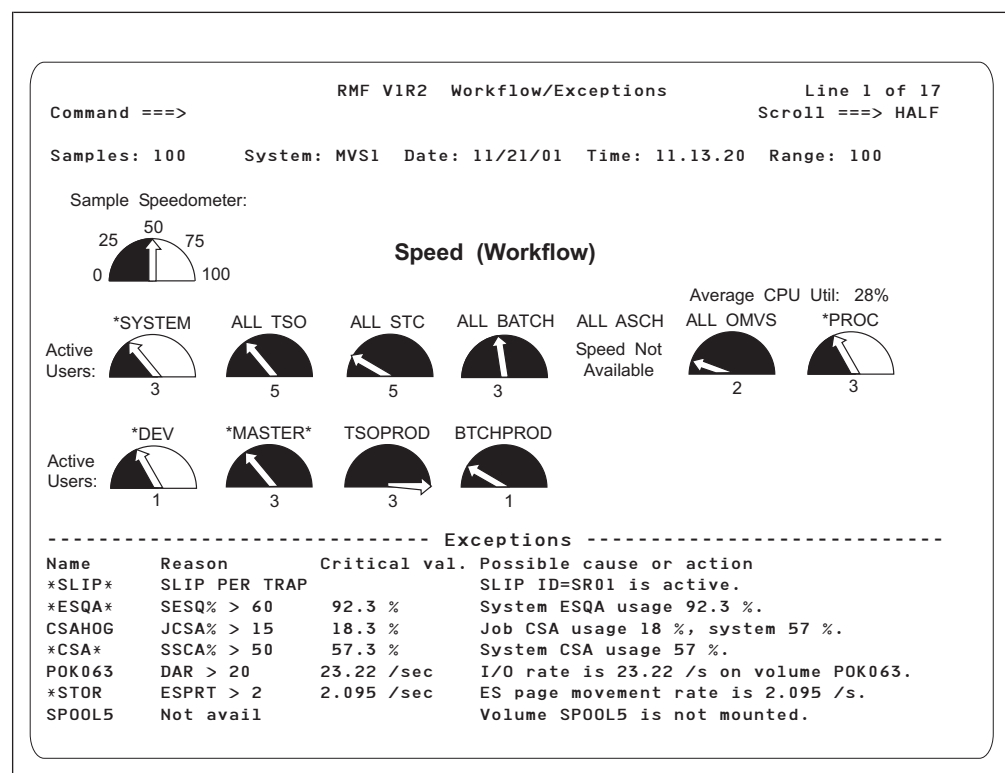


Figure 124. WFEX Graphic Report

The Workflow/Exception graphic report illustrates workflow in speedometers.

The speedometer needle points to the relative speed of the job or resource in the system. The solid/colored part to the left of the needle represents the proportion of a user's time spent doing useful work. The part to the right of the needle represents the proportion of a user's time spent delayed. If the part to the right of the needle is colored either yellow or red, then one or more exception criteria were met. A line in the Exceptions section of the report corresponds to each yellow or red speedometer. The line has the same name and the same color as the speedometer, and gives details about the exception.

## Field descriptions — Graphic WFEX Report

Table 95. Fields in the Graphic WFEX Report

Field Heading	Meaning
Name	The label for the speedometer appearing right above the speedometer. The one to ten character identifier of a workflow indicator. It can be a job, job group, or resource (processor or device). You can specify Name on the Label field of the Definition and Criteria panel or leave it blank and use the default name generated by RMF. If an indicator changes color, there is a corresponding line in the Exceptions section of the report with the same name and color giving more information about the exception.
Speedometer	<p>How to read a speedometer:</p> <p>In the upper left-hand corner of the Workflow/Exceptions graphic report is a sample speedometer with a scale showing how the workflow values are represented.</p> <p>Each speedometer is made up of 2 parts, separated by the speedometer needle:</p> <ol style="list-style-type: none"> <li>1. A <b>left</b> part, which is always solid. The solid left part represents the proportion of a user's time spent doing useful work.</li> <li>2. A <b>right</b> part, which is either hollow or solid (colored). The right part, whether it is hollow or solid represents the proportion of a user's time spent delayed.</li> </ol> <p>If the right part of the speedometer is solid (colored), one or more exception criteria were met. The solid part is colored according to the criteria set on the Definition and Criteria panel of the Workflow/Exceptions Report Options panels.</p> <p>A line in the Exceptions section of the report corresponds to the speedometers with solid right parts and gives details about the exception.</p> <p>The speedometer needle points to the relative speed of the job in the system, from 0 to 100.</p>

## XCF - Cross-System Coupling Facility Delays Report

The XCF Delays report lets you investigate situations where executing jobs are delayed when requesting service from XCF.

### How to request this report

To request the XCF Delays report, select **4** from the Primary Menu, then select **3** from the Subsystem Report menu (shown in Figure 9 on page 27) or enter the following command:

```
XCF [job_class,service_class]
```

### Contents of the report

```

RMF V2R2  XCF Delays
Command ==>>>                                     Line 1 of 1
                                                    Scroll ==>> HALF
Samples: 100    System: MVS1  Date: 09/28/16  Time: 10.03.20  Range: 100  Sec
Jobname  C  Service  DLY  ----- Main Delay Path(s) -----
          C  Class   %    % Path  % Path  % Path  % Path
GRS      S  SYSSTC   8     3 0CA0  2  0C80  2  0EA0
    
```

Figure 125. XCF Report

The graphic form of this report shows the percentage of each user's time spent waiting for XCF services.

RMF reports the overall delay (DLY %) and the four paths contributing most to delay (Main Delay Paths) due to XCF signalling traffic. RMF lists all delayed jobs by descending delay percentages.

**None** appears as the path number for pending jobs without an associated device number.

Possible causes for high XCF delay value might be caused by one or more of the following:

- Path capacity exceeded.
- Other applications are stressing the path.
- XCF delays on the receiving system.
- Some data paths are unavailable or offline.

**Note:** Any delay value shown in the report represents a delay of a message being sent. All messages are sent asynchronously. Whether the application can truly be considered to be delayed will depend on the particular application and how it is implemented. Some applications send signals and go on to do other useful work, others may need to wait for a response to come back.

### Field descriptions

Table 96. Fields in the XCF Report

Field Heading	Meaning
Jobname	Name of the job delayed when requesting service from XCF. The XCF delay report does not summarize data by job groups; all jobs within a job group are reported individually.
C	A one-character abbreviation for the job class as follows: <b>S</b> Started task <b>T</b> TSO <b>B</b> Batch <b>A</b> ASCH <b>O</b> OMVS
Service Class	The name of the service class that a specified job has been running in.
DLY %	Delay the waiting job (address space) is experiencing because of delay for XCF during the report interval, expressed as a percentage.  $DLY \% = \frac{\# \text{ Delay Samples}}{\# \text{ Samples}} * 100$ <p><b>Delay samples</b>                      The single state count of samples being delayed for XCF. RMF increments this count only once for each sample when one or more units of work (TCBs, SRBs, interrupted ready task or asynchronous exit) associated with the address space are delayed for XCF.</p> <p><b>Note:</b> This DLY% value is also found in the XCF field on the Job Delay report.</p>
Main Delay Path	The path number of the path contributing most to the delay due to XCF signalling traffic. The four paths with the highest percentages are displayed. If the job is pending and has no associated device number, NONE is displayed as the path number.

### Report options

The XCF Report Options panel is similar to the Device Report Options panel. See Figure 39 on page 72 for an example. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed. See Figure 37 on page 70 for an example.

## ZFSFS - zFS File System

The zFS File System Report measures zFS activity on the basis of single file systems. With this information, you can monitor DASD performance to ensure that there are no volumes or channels working near the limit of their capacity (space and workload, for example, I/O rates and response times).

### How to request this report

To request the ZFS File System report, select **S** on the Primary Menu, and then select **14** on the Sysplex Report menu (shown in Figure 5 on page 24), or enter one of these commands:

```
ZFSFS
ZFF
```

### Report Options

The Report Options panel for the zFS File System report allows you to specify options for this report.

```

RMF zFS File System Report Options          Line 1 of 285
Command ==>                               Scroll ==> CSR

Change or verify parameters. To exit press END.
Changes will apply to the ZFSFS report.

Name   ==> OMVS.CB8B.JAVATEST.OUTPUT.ZFS
        ALL or one of the available zFS file systems below
Detail ==> YES      Show single system data (YES or NO) in ZFSFS report

                Available zFS File Systems
APIRWW.DB2
NETVIEW.V6R1M0C.ZFS
OMVS.CB8A.JAVATEST.OUTPUT.ZFS
OMVS.CB8A.JAVATEST.ZFS
OMVS.CB8B.JAVATEST.OUTPUT.ZFS
OMVS.CB8B.JAVATEST.TESTSUIT.ZFS
OMVS.CB8B.JAVATEST.ZFS
OMVS.CB8C.JAVATEST.OUTPUT.ZFS
OMVS.CB8C.JAVATEST.TESTSUIT.ZFS
OMVS.CB8C.JAVATEST.ZFS
OMVS.CB8D.JAVATEST.OUTPUT.ZFS
OMVS.CB8D.JAVATEST.ZFS
OMVS.CB8E.JAVATEST.OUTPUT.ZFS
OMVS.CB8E.JAVATEST.TESTSUIT.ZFS
OMVS.CB8E.JAVATEST.ZFS
    
```

Figure 126. ZFSFS - Report Options

#### Name

Specify either ALL or the name of one of the zFS file systems available in the sysplex as shown in the field **Available File Systems**, which provides a list of all zFS file systems that are currently defined to the sysplex.

You can use an asterisk (\*) as the last character of the file system name as a wild card. When a wild card is used, all file systems whose names start with the specified character sequence before the asterisk are reported on, no matter which characters follow.

#### Detail

Specify the desired level of detail in the zFS File System report:

**NO** The report contains summary data for the sysplex only.

YES

The report contains data for the sysplex and all single systems.

If the list of file names is too long to fit on the first page, this report options panel can be scrolled up and down using function keys F7 and F8.

### Contents of the report

```

RMF V2R2  zFS File System - UTCPLXCB          Line 1 of 570
Command ==>                               Scroll ==> CSR
Samples: 120  Systems: 8  Date: 04/23/15  Time: 09.04.00  Range: 120  Sec
----- File System Name -----
System      Owner      Mode      Size Usq%  I/O  Resp Read  XCF
              Rate      Time %      Rate
APIRWW.DB2
  *ALL      CB86      RW S      1320M  1.0  <0.01 0.422  100 <0.01
NETVIEW.V6R1M0C.ZFS
  *ALL      CB86      RW S       12M 49.8  <0.01 0.501  100 <0.01
OMVS.CB8A.JAVATEST.OUTPUT.ZFS
  *ALL      CB8A      RW S       14G  0.2   1250 0.005 99.9 <0.01
OMVS.CB8A.JAVATEST.ZFS
  *ALL      CB8A      RW S     4922M  0.7   21.98 0.004  100 0.025
OMVS.CB8B.JAVATEST.OUTPUT.ZFS
  *ALL      CB8B      RW S       24G 15.6   1794 0.002  100 <0.01
OMVS.CB8B.JAVATEST.TESTSUIT.ZFS
  *ALL      CB8B      RW S     989M 30.6  <0.01 1.418  100 <0.01
OMVS.CB8B.JAVATEST.ZFS
  *ALL      CB8B      RW S     5868M 20.3   28.54 0.003  100 0.100
OMVS.CB8C.JAVATEST.OUTPUT.ZFS
  *ALL      CB8C      RW S     9849M  0.3  <0.01 0.787  100 <0.01
OMVS.CB8C.JAVATEST.TESTSUIT.ZFS
  *ALL      CB8C      RW S     989M  1.0  <0.01 1.102  100 <0.01
OMVS.CB8C.JAVATEST.ZFS
  
```

Figure 127. ZFSFS Report

Table 97 on page 230 describes the fields in this report.

If a file system name is specified and Detail is set to YES, then the performance values returned by each single system in the sysplex are reported underneath the \*ALL summary row for this file system. A name value of ALL will report details for all file systems.

## Mon III - ZFSFS

```

RMF V2R2   zFS File System - UTCPLXCB           Line 1 of 10
Command ==>                               Scroll ==> CSR

Samples: 120   Systems: 8   Date: 04/23/15   Time: 09.04.00   Range: 120   Sec

----- File System Name -----
      System   Owner   Mode   Size Us%   I/O   Resp   Read   XCF
                                Rate   Time   %     Rate
OMVS.CB8B.JAVATEST.OUTPUT.ZFS
  *ALL      CB8B    RW S    24G 15.6   1794 0.002  100 <0.01
  CB8A      CB8B    RW S         0.0 0.000 0.000  0.0 0.000
  CB8B      CB8B    RW S    24G 15.6   1794 0.002  100 0.000
  CB8C      CB8B    RW S         0.0 0.000 0.000  0.0 0.000
  CB8D      CB8B    RW S         0.0 0.000 0.000  0.0 0.000
  CB8E      CB8B    RW S         0.0 0.000 0.000  0.0 0.000
  CB86      CB8B    RW S         0.0 0.000 0.000  0.0 0.000
  CB88      CB8B    RW S         0.0 <0.01 0.000  100 <0.01
  CB89      CB8B    RW S         0.0 0.000 0.000  0.0 0.000
  
```

Figure 128. ZFSFS Report - File system details

From the zFS File System Report, you can navigate to a variety of detail information using cursor-sensitive control. If you place the cursor on any of the lines with file system values, a pop-up window appears showing the details for this file system. Figure 129 shows an example of a pop-up panel with File System Details.

```

zFS File System Details

File System Name : OMVS.CB8B.JAVATEST.OUTPUT.ZFS
Mount
Point : /CB8B/javatest/output

System : CB8B           Owner : CB8B           Mode : RW S

----- Read -----
--- Appl --- XCF --- Aggr
Rate Resp Rate Resp Aggr
      Time   Time
1793 0.002 0.000 0.000 170.7

----- Write -----
--- Appl --- XCF --- Aggr
Rate Resp Rate Resp Aggr
      Time   Time
0.758 0.047 0.000 0.000 682.7

Vnodes      : 33      USS held vnodes      : 5
Open objects : 0      Tokens               : 4
User cache 4k pages : 1  Metadata cache 8k pages : 90

ENOSPC errors : 0      Disk I/O error       : 0
XCF comm. failures : 0  Cancelled operations : 0

Press Enter to return to the Report panel.
  
```

Figure 129. ZFSFS Report - zFS File System Details

Table 98 on page 231 describes the fields in this report.

### Field descriptions

This topic describes the fields in the zFS File System report.

Table 97. zFS File System report field descriptions

Field Heading	Meaning
File System Name	File system name.



Table 97. zFS File System report field descriptions (continued)

Field Heading	Meaning
System	Name of the system connected to the file system.  In the first data line for a file system, the name is '*ALL' to indicate that this line shows the SYSPLEX view of the data rather than a single system view.
Owner	Name of owning system.
Mode	Mount mode of the file system. Possible values are: <b>NM</b> Not mounted. <b>QS</b> Not available because the aggregate is quiesced. <b>RO</b> Mounted in read-only mode. <b>RW</b> Mounted in read-write mode.  The mount mode is followed by an <b>S</b> if the file system is using zFS sysplex sharing (RWSHARE).
Size	Maximum logical size of the file system (in Bytes).
Usg%	Percentage of currently used space by the file system.
I/O Rate	The rate of read and write requests per second (directory and file) made by applications to this file system.
Resp Time	Average response time in milliseconds for read and write requests made by applications to this file system.
Read%	Percentage of read operations contained in 'I/O Rate'.
XCF Rate	The rate of read and write XCF calls per second to the server.

Table 98. zFS File System report field descriptions - zFS File Systems Details panel

Field Heading	Meaning
File System Name	File system name.
Mount Point	Mount point of the file system.
System	Name of the system connected to the file system. If the name is *ALL, this indicates that this line shows the SYSPLEX view of the data rather than a single system view.
Owner	Name of owning system.
Mode	Mount mode of the file system. Possible values are: <b>NM</b> Not mounted. <b>QS</b> Not available because the aggregate is quiesced. <b>RO</b> Mounted in read-write mode. <b>RW</b> Mounted in read-write mode.  The mount mode is followed by an <b>S</b> if the file system is using zFS sysplex sharing (RWSHARE).
Application read rate	The rate of read requests per second (directory and file) made by applications to this file system.
Application read resp time	The average response time for read requests made by applications to this file system (in milliseconds).
XCF read rate	The rate of read XCF calls per second to the server
XCF read resp time	The average response time required for a read XCF call to the server (in milliseconds).
Aggregate read rate	Read data transfer rate in bytes/second for the aggregate.

Table 98. zFS File System report field descriptions - zFS File Systems Details panel (continued)

Field Heading	Meaning
Application write rate	The rate of write requests per second (directory and file) made by applications to this file system.
Application write resp time	The average response time for write requests made by applications to this file system (in milliseconds).
XCF write rate	The rate of write XCF calls per second to the server.
XCF write resp time	The average response time required for a write XCF call to the server (in milliseconds).
Aggregate write rate	Write data transfer rate in bytes/second for the aggregate. The write data transfer rate also includes write activity from zFS daemons for file systems mounted in R/W mode.
Vnodes	Number of vnodes in memory for the file system.
USS held vnodes	Number of vnodes that the logical file system layer, also known as USS, holds for the file system.
Open objects	The number of files and directories currently open.
Tokens	The number of tokens held for objects in the file system from the token manager.
User cache 4K pages	The number of 4K pages in the user file cache for this file system.
Metadata cache 8K pages	The number of 8K pages in the metadata cache for this file system.
ENOSPC errors	The number of ENOSPC errors seen by applications for this file system on this system.
Disk I/O errors	The number of disk I/O errors for disk I/Os performed by this system.
XCF comm. failures	The number of XCF communication timeouts or errors seen by XCF requests made for this file system on this system.
Cancelled operations	The number of times a task was asynchronously ABENDED (CANCEL) or EOMd (FORCE) while running an operation in this file system.

## ZFSKN - zFS Kernel report

The zFS Kernel report provides a variety of measurements counting the calls made to zFS from z/OS UNIX and the average response time of zFS requests. This information gives the basic measure of zFS performance and can be used to determine the appropriate tuning options needed to make best use of the z/OS File System (zFS).

### How to request this report

To request the ZFS Kernel report, select **S** on the Primary Menu, and then select **15** on the Sysplex Report menu (shown in Figure 5 on page 24), or enter one of these commands:

```
ZFSKN
ZFK
```

## Contents of the report

```

RMF V2R2   zFS Kernel   - UTCPLXCB           Line 1 of 8
Command ==>                               Scroll ==> CSR

Samples: 120   Systems: 8   Date: 04/23/15   Time: 09.04.00   Range: 120   Sec

System      - Request Rate -   --- XCF Rate ---   - Response Time -
Name        Local   Remote   Local   Remote   Local   Remote
CB8A        1356   1305   0.025  82.57   8.000   955.0
CB8B        2552   120.8   0.000   0.400  13057   1027
CB8C        11.33   0.175   0.000   0.000   22.00   106.0
CB8D        22315  1245   0.000   0.592  148.1   3115
CB8E        1747   2099   0.000   75.46  186.0   1214
CB86        29605  418.1   0.000   34.98  354.0   4921
CB88        1921   260.7   0.000   43.13  183.0   1738
CB89        517.3  2590   0.000   75.33  97.00   1184

```

Figure 130. zFS Kernel Report

### Field descriptions

This topic describes the fields in the zFS Kernel report.

Table 99. zFS Kernel report field descriptions

Field Heading	Meaning
System Name	Name of the system running zFS.  In the context of requests against file systems, this is the name of the requesting system.
Request Rate	Rate of zFS requests during the report interval for file systems which are locally and remotely owned.  A file system is locally owned if the requesting system is also the owner of the file system. It is remotely owned if the owner of the file system is not the requesting system.
XCF Rate	Rate of zFS requests during the report interval requiring data from another system by XCF, both for locally and remotely owned file systems.
Response Time	Average time in milliseconds required for the completion of the zFS requests during the report interval for locally and remotely owned file systems.

## ZFSOVW - zFS Overview Report

To use a zFS file system within a z/OS UNIX file system hierarchy to its full capacity, it is necessary to apply appropriate tuning options. The zFS performance especially depends on a suitable tailoring of its cache sizes to reduce I/O rates and path lengths. The performance can also be improved by adapting available disk space.

This report provides a summary of zFS activity, request and DASD statistics on the current system and thus helps to control and tune the zFS environment. For example, you can use the HIT% values in the Cache Activity section as an indication of whether the current cache sizes are sufficient.

## How to request this report

To request the zFS Overview report, select **S** from the Primary Menu, then select **13** from the Sysplex Report Selection Menu (shown in Figure 5 on page 24) or enter one of the following commands:

```
ZFSOVW
ZF0
```

## Contents of the report

This topic shows the content of the zFS Overview report.

```

RMF V2R2   zFS Overview - SYSDPLEX   Line 1 of 8
Command ==>                               Scroll ==> CSR
Samples: 120   Systems: 8   Date: 09/28/16   Time: 09.04.00   Range: 120   Sec

System      -----Wait%-----      ----- Cache Activity -----
            I/O  Lock Sleep      ---User---      --Vnode---      -Metadata-
                               Rate Hit%      Rate Hit%      Rate Hit%

CB8A        100  1.4  100   954.6 99.5   1350 94.8   236.5 96.0
CB8B         5.2  0.2  0.0   1226 97.8   1967 86.7   4427 96.1
CB8C         0.0  ----  0.1   61.32 100    649.7 100    62.88 98.3
CB8D         0.1  0.1  0.2   893.1 100    15393 100    1102 97.1
CB8E        100  0.0  7.2   1205 100    2287 97.2   195.1 96.6
CB86         9.4  0.1  0.1   2326 34.5   18660 91.8   42145 91.8
CB88         5.5  0.0  2.5   1242 100    676.6 97.0   96.36 97.8
CB89        100  0.3 46.0   386.7 100    2373 97.4   273.9 97.3

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

Figure 131. zFS Overview Report

## Field descriptions

Table 100. zFS Overview report field descriptions

Field Heading	Meaning
System	Name of the system running zFS.
Wait%	<p>These Wait percentages are reported:</p> <p><b>I/O</b> Percentage of time that zFS requests had to wait for I/O completion.</p> <p><b>Lock</b> Percentage of time that zFS requests had to wait for locks.</p> <p><b>Sleep</b> Percentage of time that zFS requests had to wait for events.</p> <p>Dashes (----) in these fields indicate that RMF is unable to calculate a reasonable value.</p>
Cache Activity section	

Table 100. zFS Overview report field descriptions (continued)

Field Heading	Meaning
User	<p>The user file cache is for caching regular user files that are larger than 7K. The measured statistics have these meanings:</p> <p><b>Rate</b> Total number of read and write requests per second made to the user file cache.</p> <p><b>Hit%</b> Percentage of read and write requests to the user file cache that completed without accessing the DASDs.</p>
Vnode	<p>The vnode cache is used to hold virtual inodes. An inode is a data structure related to a file in the file system, holding information about the file's user and group ownership, access mode and type. The measured statistics have these meanings:</p> <p><b>Rate</b> Number of read and write requests per second made to the vnode cache.</p> <p><b>Hit%</b> Percentage of read and write requests to the vnode cache that completed without accessing the DASDs.</p>
Metadata	<p>The metadata cache is used for file system metadata and for files smaller than 7K. It resides in the primary z/FS address space. The measured statistics have these meanings:</p> <p><b>Rate</b> Number of read and write requests per second made to the metadata cache.</p> <p><b>Hit%</b> Percentage of read and write requests to the metadata cache that completed without accessing the DASDs.</p>

## Navigating to details in the zFS Overview report

From the **zFS Overview report**, you can navigate to a variety of detail information using cursor-sensitive control.

The detail information is provided in panels shown in Figure 132 on page 236 through Figure 135 on page 239. You reach these panels as follows:

- From the *Wait% -I/O* field, you can reach the *I/O Details by Type* panel (Figure 132 on page 236).
- From any value in the *Cache Activity - User* section, you can reach the *User Cache Details* panel (Figure 133 on page 237).
- From any value in the *Cache Activity - Vnode* section, you can reach the *Vnode Cache Details* panel (Figure 134 on page 238).
- From any value in the *Cache Activity - Metadata* section, you can reach the *Metadata Cache Details* panel (Figure 135 on page 239).

### zFS Overview - I/O Details by Type

The *zFS Overview - I/O Details by Type* report displays a breakdown of I/O requests into the following types:

- I/O for file system metadata
- I/O for log data
- I/O for user file data

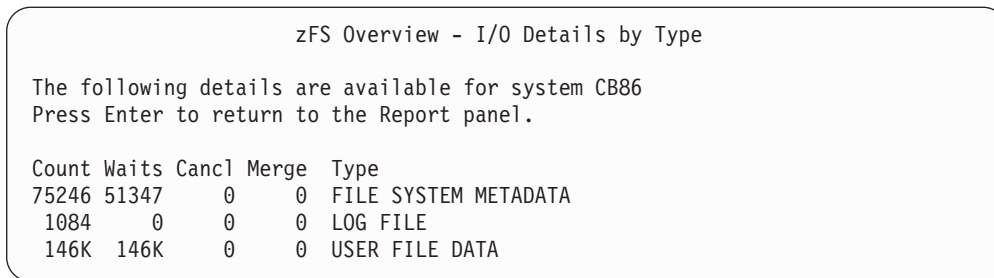


Figure 132. I/O Details by Type

Table 101. Fields in the zFS Overview report - I/O details

Field Heading	Meaning
Count	Total number of I/O requests of the indicated type.
Waits	Number of zFS requests waiting for an I/O completion of the indicated I/O type.
Canc1	Number of cancelled zFS requests during an I/O request of the indicated type, for example, a user tried to delete a file during a pending I/O to this file's metadata.
Merge	Number of merges of two I/O requests into a single request because of better performance.
Type	Type of the I/O request (I/O for metadata, log data or user file data).

### zFS Overview - User Cache Details

The user file cache is for caching regular user files that are larger than 7K. The *zFS Overview - User Cache Details* report displays the following details of the user file cache activity:

```

zFS Overview - User Cache Details

zFS Overview - User Cache Details

The following details are available for system CB88
Press Enter to return to the Report panel.

Size      :      2048M      Storage fixed : NO
Total Pages :      524K
Free Pages :      281K
Segments  :      47K

----- Read -----      ----- Write -----
Rate Hit% Dly% Async      Rate Hit% Dly% Sched      Read% Dly%
      Rate
1241  100  0.0  1236      155  100  0.0  0.150      87.5  0.0

----- Misc -----
Page Reclaim Writes :      0
Fsyncs               :      0
    
```

Figure 133. User Cache Details

Table 102. Fields in the zFS Overview Report - User Cache Details

Field Heading	Meaning
System	Name of the system running zFS.
Size	Total size of the user file cache.
Total Pages	Total number of pages in the user file cache.
Free Pages	Total number of free pages in the user file cache.
Segments	Total number of allocated segments in the user file cache.
Storage fixed	Shows whether the size of the user file cache storage is fixed. If the zFS parameter user_cache_size is set to 'fixed', then zFS reserves real storage for use by zFS only.  The 'fixed option' helps to improve performance during data access and can be applied if you have enough real storage available.
Read Rate	Number of read requests per second made to the user file cache.
Read Hit%	Percentage of read requests to the user file cache that completed without accessing the DASD.
Read Dly%	Percentage of delayed read requests to the user file cache. A read request is delayed if it must wait for pending I/O, for example, because the file is in a pending read state due to asynchronous read ahead from DASD to the user file cache.
Async Read Rate	Number of read aheads per second.
Write Rate	Number of write requests per second made to the user file cache.
Write Hit%	Percentage of write requests to the user file cache that completed without accessing the DASD.
Write Dly%	Percentage of delayed write requests to the user file cache.  The following reasons are counted as write request delays: <b>Write wait</b> Write must wait for pending I/O. <b>Write faulted</b> Write to a file needs to perform a read from DASD. If a write-only updates a part of a file's page, and this page is not in the user file cache, then the page must be read from DASD before the new data is written to the cache.
Scheduled Write Rate	Number of scheduled writes per second.
Read%	Percentage of read requests, based on the sum of read and write requests.

## Mon III - ZFSOVW

Table 102. Fields in the zFS Overview Report - User Cache Details (continued)

Field Heading	Meaning
Dly%	Percentage of delayed requests, with the following events counted as delays: <ul style="list-style-type: none"> <li>• Read wait: a read request must wait for a pending I/O operation.</li> <li>• Write wait: a write request must wait because of a pending I/O operation.</li> <li>• Write faulted: a write request to a file in the user file cache needs to perform a read operation from DASD before writing, because the required page of that file is currently not in the cache.</li> </ul>
Page Reclaim Writes	Total number of page reclaim writes. A page reclaim write action writes one segment of a file from the user file cache to DASD. Page reclaim writes are performed to reclaim space in the user file cache. If page reclaim writes occur too often in relation to the write rate, then the user file cache may be too small.
Fsyncs	Total number of requests for file synchronization (fsync) between user file cache and DASD.

### zFS Overview - Vnode Cache Details

The vnode cache is used to hold virtual inodes. An inode is a data structure related to a file in the file system, holding information about the file's user and group ownership, access mode, and type. The *zFS Overview - Vnode Cache Details* report displays the following details of the vnode cache activity:

```

zFS Overview - Vnode Cache Details

The following details are available for system CB88
Press Enter to return to the Report panel.

Size :      32768

----- Vnodes -----
      Total  Size  Ext.#  Ext.Size  Open  Held
      45183  224   32768    816    118  1024

----- Requests -----
      Total  Rate  Hit%  Alloc  Delete
      81187  676.6  97.0    0    2510
    
```

Figure 134. Vnode Cache Details

Table 103. Fields in the zFS Overview Report - Vnode Cache Details

Field Heading	Meaning
System	Name of the system running zFS.
Size	Number of vnodes that will be initially cached by zFS.
Vnodes Total	Number of currently allocated vnodes in the vnode cache. If more vnodes are requested than are currently available, then zFS dynamically allocates more vnodes.
Vnodes Size	Size of a vnode data structure in bytes.
Vnodes Ext.#	Number of extended vnodes.
Vnodes Ext. Size	Size of an extended vnode data structure in bytes.
Vnodes Open	Number of currently open vnodes.
Vnodes Held	Number of vnodes currently held in zFS by USS.
Requests Total	Number of requests to the vnode cache.
Requests Rate	Number of requests per second made to the vnode cache.
Requests Hit%	Percentage of requests to the vnode data that found the target vnode data structures in the vnode cache. High hit rates indicate a favorable zFS environment, because each miss involves initialization of vnode data structures in the vnode cache.



Table 103. Fields in the zFS Overview Report - Vnode Cache Details (continued)

Field Heading	Meaning
Requests Alloc	Number of requests to create new vnodes (for operations such as create or mkdir).
Requests Delete	Number of requests to delete vnodes (for operations such as remove or failed creates or mkdirs).

### zFS Overview - Metadata Cache Details

The metadata cache is used to contain all file system metadata; this metadata includes all directory contents, file status information, and file system structures. Additionally, it also caches data for files smaller than 7 K. It resides in the primary zFS address space.

The *zFS Overview - Metadata Cache Details* report displays these details of the metadata cache:

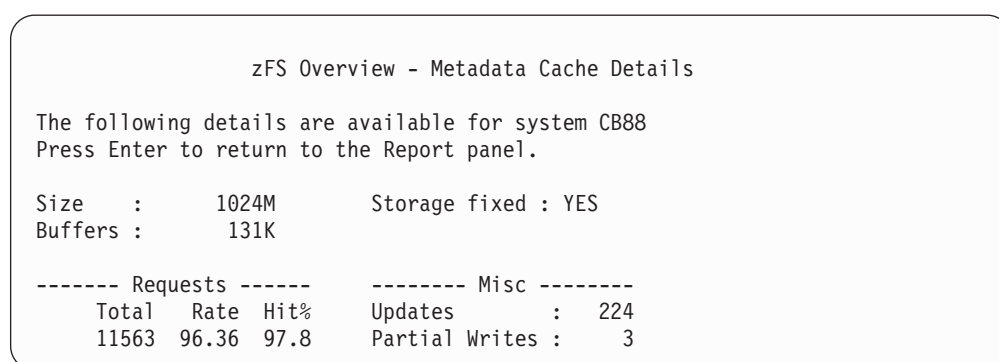


Figure 135. Metadata Cache Details

Table 104. Fields in the zFS Overview Report - Metadata Cache/Backing Cache Details

Field Heading	Meaning
System	Name of the system running zFS.
Size	Total size of the metadata cache.
Buffers	Total number of buffers in the metadata cache. The buffer size is 8K.
Storage fixed	Shows whether the size of the metadata cache storage is fixed. If the zFS parameter meta_cache_size is set to 'fixed', then zFS reserves real storage for use by zFS only. The 'fixed option' helps to improve performance during data access and can be applied if there is enough real memory available.
Total	Number of requests made to the metadata cache.
Rate	Number of requests per second made to the metadata cache.
Hit%	Percentage of requests to the metadata cache completing without accessing the DASD.
Updates	Number of updates made to buffers in the metadata cache.
Partial writes	Number of times that only half of an 8K metadata block needed to be written.



---

## Chapter 3. Snapshot reporting with Monitor II

This information unit describes the format and contents of the following Monitor II reports:

**ARD** Address Space Resource Data report  
**ASD** Address Space State Data report  
**ASRM**  
Address Space SRM Data report  
**CHANNEL**  
Channel Path Activity report  
**DEV** Device Activity report  
**HFS** HFS File System Statistics report  
**ILOCK**  
IRLM Long Lock Detection report  
**IOQUEUE**  
I/O Queuing Activity report  
**LLI** Library List report  
**OPT** OPT Settings report  
**PGSP** Page Data Set Activity report  
**SDS** Sysplex Data Server report  
**SENQ** System Enqueue Contention report  
**SENQR**  
System Enqueue Reserve report  
**SPAG** Paging Activity report  
**SRCS** Central Storage/Processor/SRM report

---

### Monitor II sessions

You can **display** a Monitor II report during:

- An ISPF display session

This session is started with the command

```
RMF
```

This leads to the RMF Primary menu, then you select **2** to get the Monitor II ISPF session.

- A TSO/E display session

This session is started with the TSO/E command

```
RMFMON
```

- A background session

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

```
MODIFY RMF,START AB
```

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

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

In all sessions, you can get the same reports. There is just a small difference in the syntax used to call them:

- Display Session

The reports are called via *commands* according to TSO/E syntax rules:

**Example:** ASD T,A

- Background Session

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

**Example:** ASD(T,A)

This chapter shows report examples from an ISPF session, the report format of a RMFMON session is very similar, and the meaning of all report fields is the same in all versions of a report.

---

## Structure of Monitor II reports

This chapter presents sample reports and the meaning and contents of each field in each report. The sample reports show the display screen contents from an ISPF session for each report.

When the reports are printed, the contents are identical to the report contents shown on the screen with some differences in the layout of the printed output.

### Contents of the Monitor II report header

A Monitor II report header looks different, depending on whether you use the ISPF interface or the TSO/E interface.

#### If you are using an ISPF session

Each report consists of

- A header line identifying the report
- A line for commands and scroll amount field
- A status line for CPU, UIC, and PR. This line also contains the SMF system ID and the current setting of the report mode (Total or Delta).
- A variable number of data lines.

#### If you are using a TSO/E session

Each report consists of

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

Figure 136 shows the different report areas for a TSO/E session. For a description of each area, see Table 105 on page 243.

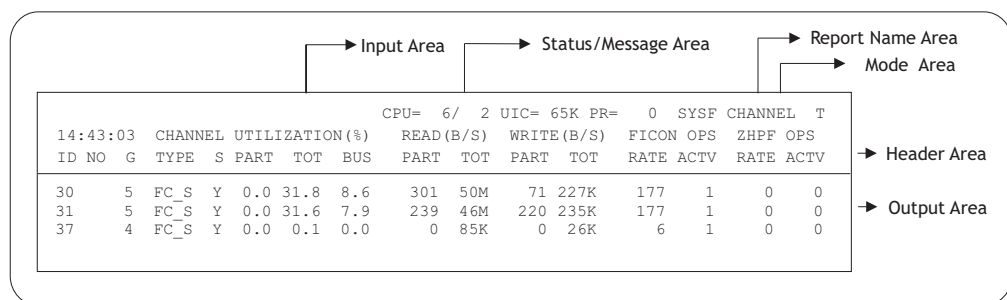


Figure 136. Header of a Monitor II TSO/E session report

When you begin a session, the cursor appears in the **input area**. During the session you issue all display commands from this area. Other areas indicated in the figure are described in Table 105 on page 243.

Table 105. Monitor II Display Session Areas

Area	Function
Report title	The type of measurement data
F	Indicates more pages
Input Area	Issue all commands from here.
Status/Message Area	<p>This area contains:</p> <p><b>CPU</b> Current average processor utilization.</p> <p>This information depends on the activity of Monitor I.</p> <p>If Monitor I CPU gathering is <b>active</b>, the header line shows two views separated by a slash (/):</p> <ul style="list-style-type: none"> <li>• The MVS view of the CPU utilization</li> <li>• The LPAR view of the CPU utilization</li> </ul> <p>If Monitor I CPU gathering is <b>not active</b>, the header line shows:</p> <ul style="list-style-type: none"> <li>• The SRM view of the CPU utilization.</li> <li>• '***' due to missing CPU measurement data for the LPAR view</li> </ul> <p>For more information about the different views of CPU utilization refer to "CPU - CPU Activity report" on page 340.</p> <p><b>UIC</b> The current system unreferenced interval count.</p> <p>Values greater than 9999 are displayed as nnK to indicate a multiple of 1000. The maximum value is 65K.</p> <p><b>PR</b> The rate of page-ins per second excluding swap-ins, VIO (virtual input/output), reclaims, and hiperspaces.</p> <p><b>System</b> The SMF system ID of this system.</p>
Report Name Area	The report name.
Mode Area	The current setting for the report mode (either D for delta or T for total) and hardcopy mode (either H for hardcopy, or blank)
Header Area	Consists of two lines of column headings that identify the data fields included in the report.
Output Area	Contains the report data.

## Different formats of Monitor II reports

Monitor II offers two types of reports:

- **Table Reports** - Example: ASD Report  
Table reports have a variable number of data lines.
- **Row Reports** - Example: ASDJ Report  
Row reports have only one line of data. When you request a row report repeatedly, each request adds one line of data to the display. You can use the repetitive requests to build a table of information.

## Different modes of Monitor II reports

Monitor II offers two modes for the session reports:

- **Total mode**  
A total mode report shows the cumulative total since the beginning of the Monitor I interval.
- **Delta mode**

## Monitor II Reports

A delta report mode shows the change in the activity since the previous request for the report.

### Monitor II display session reports

For a Monitor II display session, RMF creates a single output data set for each session. All **printed output** resulting from either hardcopy mode or the non-ISPF Print command is sent to the same output data set.

You need to allocate this data set before starting the display session:

```
ALLOC F(RMFDMSO) DS(dsname) SHR
```

If you issue the ISPF Print command, the output is stored in data set `userid.SPFx.LIST`, this is the standard way as ISPF handles print output.

See the *z/OS RMF User's Guide* for more details.

Each report printed when the session is in hardcopy mode is delimited by a line of plus signs (+). Each report printed as a result of the print display command is delimited by a line of asterisks (\*).

Fields within the line of delimiters indicate the operands specified on the report request, whether the session is in delta or total mode, and the name of the report.

When there are repetitive requests for the same **row report**, headings appear for the first request; data lines appear for each subsequent request. A field within the line indicates the time of the report.

Figure 137 shows an example of printed output from a display session.

```
RMF Monitor II HARDCOPY LOG
SESSION NAME TSO
z/OS V2R2          SYSTEM ID SYS1      DATE 07/28/2016
                   RPT VERSION V2R2 RMF  TIME 17.12.25
*****
++ MIG= 729K CPU= 41 UIC= 254 PR= 0 SYSTEM=SYS1 ***** TOTAL MODE ***** ARD *****
***** 17.12.28 *****
17:12:28 DEV  FF  FF PRIV LSQA X C SRM TCB  CPU  EXCP SWAP LPA CSA NVI V&H
JOBNAME  CONN 16M 2G  FF  CSF M R ABS  TIME  TIME  RATE RATE  RT  RT  RT  RT
*MASTER* 11432 0 493 1261 116 0.0 7973.9 17944 ---- ---- ---- ----
PCAUTH  0.000 0 94 2 30 X 0.0 0.04 0.05 ---- ---- ---- ----
RASP    0.000 --- 30 --- --- X 0.0 0.02 27.48 ---- ---- ---- ----
TRACE   0.000 0 145 3 49 X 0.0 0.02 0.03 ---- ---- ---- ----
XCFAS   1470 0 2407 2457 2087 X 0.0 245.4 796.20 ---- ---- ---- ----
GRS     29514 17 17 38 33 X 0.0 2135.2 10346 ---- ---- ---- ----
SMXC    0.000 0 47 2 20 0.0 44.81 51.38 ---- ---- ---- ----
SYSBMAS 0.000 0 13 5 19 0.0 0.01 0.02 ---- ---- ---- ----
DUMPSRV 278.1 0 100 2 37 0.0 19.30 61.10 ---- ---- ---- ----
CONSOLE 1148 0 44 2 31 X 0.0 1189.7 1365.2 ---- ---- ---- ----
ANTMAIN 0.085 0 86 2 23 X 0.0 0.02 0.05 ---- ---- ---- ----
ALLOCAS 0.000 0 78 2 21 X 0.0 0.02 0.03 ---- ---- ---- ----
SMF     324.0 0 49 2 36 X 0.0 1.32 43.32 ---- ---- ---- ----
VLF     0.956 1 117 32 34 X 0.0 231.96 235.34 ---- ---- ---- ----
LLA     215.2 10 9 12 56 X 0.0 80.48 96.14 ---- ---- ---- ----
DFRMM   265.3 0 32 2 61 X 0.0 80.95 98.33 ---- ---- ---- ----
FTPSRV01 0.783 1 1 32 S 0.0 7.22 7.64 ---- ---- ---- ----
```

Figure 137. Format of printed reports from a Monitor II display session

The report was printed as a result of a `HARDCOPY ON` or `HARDCOPY` command. This command puts the session in hardcopy mode and causes all reports to be displayed and a printable version to be sent to the output data set.

**Monitor II background session reports**

For a Monitor II background session, RMF creates a single output data set for each report requested. All **interval reports** for that measurement activity are written to a single output data set. Thus, if you request three measurements for a session with five reporting intervals, RMF creates three data sets and writes five reports to each data set.

The printed output generated for each **table report** requested is preceded by a line of plus signs (+). Fields within the line of plus signs indicate the option and any operands specified for the report and whether the report is in delta mode or total mode. Each iteration of the report is separated by a line of plus signs. A field within the line of plus signs indicates the time of the report. The column headings are repeated for each iteration of the report.

The printed output generated for each **row report** is preceded by a single line of plus signs (+). Fields within the line of plus signs indicate the option and any operands specified for the report and whether the session is in delta or total mode. Headings appear only once. There is no delimiter between successive iterations of a row report, and a field within the report line indicates the time of the report.

Figure 138 shows an example of printed output generated during a Monitor II background session.

```

SESSION NAME BB

+++++++ N(274) ++++++ TOTAL MODE ++++++ ASD ++++++
+++++ 14.35.46 ++++++
14:35:46 S C R DP CS ESF ESF TAR WS TX WRK CPU I/O STM
JOBNAME SRVCLASS P L LS PR F +CS WSS IN SC RV RV RV RV
*MASTER* SYSTEM 1 NS FF 133 0 133 0 0 0 0 0 +0 +0 +0
PCAUTH SYSSTC 1 NS 76 33 0 33 0 11 0 0 0 +0 +0 +0
TRACE SYSSTC 1 NS 77 176 0 176 0 11 0 0 0 +0 +0 +0
GRS SYSTEM 1 NS FF 561 0 561 33K 11 0 0 0 +0 +0 +0
CONSOLE SYSTEM 1 NS FF 174 0 174 0 11 1 0 0 +0 +0 +0
ALLOCAS SYSTEM 1 NS 71 986 0 986 0 11 1 0 0 +0 +0 +0
LLA SYSSTC 1 NS 71 170 0 170 0 30 0 150 0 +0 +0 +0
RMF33 SYSSTC 1 NS 71 159 0 159 0 34 0 150 0 +0 +0 +0
+++++ 14.35.54 ++++++
14:35:54 S C R DP CS ESF ESF TAR WS TX WRK CPU I/O STM
JOBNAME SRVCLASS P L LS PR F +CS WSS IN SC RV RV RV RV
DSNDBM1 SYSSTC 1 NS FC 1350 0 1350 0 33 0 150 +0 +0 +0
BOYLEMM BATCH 2 IN 78 85 0 133 0 0 4 101 +0 +0 +0
SMF SYSTEM 1 NS FF 94 0 94 0 11 0 150 +0 +0 +0
DFHSM SYSSTC 1 NS 74 610 0 610 0 527 1 150 +0 +0 +0
VTAM SYSSTC 1 NS FD 678 0 678 0 34 0 0 +0 +0 +0
SOS SYSSTC 1 IN 79 76 0 567 0 0 1.1K 150 +0 +0 +0
AMSAQFT SYSSTC 1 NS 72 54 0 54 0 33 0 150 +0 +0 +0
JES2 SYSSTC 1 NS FE 900 0 900 0 24 0 0 +0 +0 +0
CATALOG SYSTEM 1 NS FF 1552 0 1552 0 11 0 150 +0 +0 +0
+++++ 14.36.02 ++++++
14:36:02 S C R DP CS ESF ESF TAR WS TX WRK CPU I/O STM
JOBNAME SRVCLASS P L LS PR F +CS WSS IN SC RV RV RV RV
HUBERF TSO 2 IN 78 498 0 597 0 0 3 0 +0 +0 +0
ZAPPERD TSO 2 IN 72 110 0 121 0 0 1 150 +0 +0 +0
PTRACYB TSO 2 IN 78 174 0 229 0 0 3 0 +0 +0 +0
IRLMPROC SYSSTC 1 NS FC 61 0 61 0 33 0 0 +0 +0 +0

```

Figure 138. Format of printed reports from a Monitor II background session

---

## ARD/ARDJ - Address Space Resource Data report

The ARD and ARDJ reports give information on the system resources that are used by each address space in the system or each address space that meets the selection criteria that you specify when you requested the report. The information provided in these reports includes, for example, information on processor time, paging, and central storage.

The ARD report enables you to determine which jobs are creating performance problems.

Once a problem job has been identified, you can request an ARDJ report for that particular job. This enables you to focus your reporting on a known problem area.

### How to request this report

Different methods are used to request the ARD and ARDJ reports.

#### How to request an ARD report

- In ISPF, specify **1** on the Address Space Report Selection menu.
- In TSO/E, use **PF1** to select the ARD report.
- Command interface:

##### Display session

```
ARD [class,status]
```

##### Background session

```
ARD [(class,status)]
```

#### How to request an ARDJ report

- In ISPF, specify **4** on the Address Space Report Selection menu.
- Command interface:

##### Display session

```
ARDJ jobname
```

##### Background session

```
ARDJ (jobname)
```

### Contents of the report

The information shown in an ARD and an ARDJ report is identical, except the content of the first column which is:

JOBNAME for the ARD report  
TIME for the ARDJ report

In the ARD report of Figure 139 on page 247, the number of data lines in the report depends on the number of address space identifiers in the system that meet your selection criteria. The shown report is a sample for a system running in z/Architecture.

In the ARDJ report of Figure 140 on page 247, the number of rows depends on your requests to build a table of information for a particular job.



```

RMF - ARD Address Space Resource Data
Command ==>
Line 1 of 85
Scroll ==> HALF

CPU= 9    UIC= 255 PR= 13    System= SYS1 Total

14:51:59 DEV  FF  FF PRIV LSQA X C SRM TCB  CPU  EXCP SWAP LPA CSA NVI V&H
JOBNAME  CONN 16M  2G  FF CSF M R ABS TIME TIME RATE RATE RT  RT  RT  RT

*MASTER* 1476  0 493 1261 110      0.0 149.0 544.5 0.29 0.00 0.0 0.0 0.0 0.0
PCAUTH   0.000 0  94  0 112 X      0.0  0.00  0.00 0.00 0.00 0.0 0.0 0.0 0.0
RASP     0.000 0  30 207  44 X      0.0  0.00  1.90 0.00 0.00 0.0 0.0 0.0 0.0
TRACE    0.067 0 145  1 173 X      0.0  0.00  0.00 0.00 0.00 0.0 0.0 0.0 0.0
DUMPSRV  14.69 0  33  0  76      0.0  1.01  1.73 0.00 0.00 0.0 0.0 0.0 0.0
XCFAS    1470  0 2407 2457 2087 X      0.0 245.4 796.2 2.51 0.00 0.0 0.0 0.0 0.0
GRS      0.000 0  47  52 1047 X S      0.0 235.0 339.8 0.00 0.00 0.0 0.0 0.0 0.0
SMXC     2.400 0  13  0  50      0.0  9.44 11.10 0.00 0.00 0.0 0.0 0.0 0.0
SYBMAS   0.000 0 100 104  31      0.0  4.74  4.83 0.00 0.00 0.0 0.0 0.0 0.0
CONSOLE  25.23 0  44 19 101 X      0.0 61.85 66.37 0.07 0.00 0.0 0.0 0.0 0.0
WLM      0.493 0  86  52 503 X      0.0 2130 2241 0.00 0.00 0.0 0.0 0.0 0.0
ANTMAIN  0.934 0  78  3 163 X X      0.0  2.60  3.05 0.00 0.00 0.0 0.0 0.0 0.0
ANTAS000 0.621 0  49  2 100 X      0.0  0.07  0.08 0.00 0.00 0.0 0.0 0.0 0.0
OMVS    168.5 0 117  85 647 X      0.0 22.57 26.50 0.00 0.00 0.0 0.0 0.0 0.0
IEFSCHAS 0.000 0  9  0  34 X      0.0  0.00  0.00 0.00 0.00 0.0 0.0 0.0 0.0
JESXCF   1.629 0  32  4  78 X      0.0 24.64 42.03 0.00 0.00 0.0 0.0 0.0 0.0
    
```

Figure 139. ARD Report in z/Architecture

```

RMF - ARDJ Address Space Resource Data
Command ==>
Line 1 of 14
Scroll ==> HALF

CPU= 37/ 35 UIC=2540 PR=  0    System= SYS1 Total

BGBO     DEV  FF PRIV LSQA LSQA X C SRM TCB  CPU  EXCP SWAP LPA CSA NVI V&H
TIME     CONN 16M  FF CSF ESF M R ABS TIME TIME RATE RATE RT  RT  RT  RT

13:33:17 10.02 39  4  38  0  --- 20.73 32.26 ---- ---- --- --- --- ---
13:33:50 11.96 35  5  40  0  944 22.33 24.22 39.2 0.00 0.0 0.0 0.0 0.0
13:33:53 12.39 35  5  40  0  962 22.63 24.58 28.7 0.00 0.0 0.0 0.0 0.0
13:34:40 13.87 35  5  40  0  1K 24.42 26.66 34.8 0.00 0.0 0.0 0.0 0.0
13:34:42 14.16 35  5  40  0  1K 24.71 27.00 38.5 0.00 0.0 0.0 0.0 0.0
13:34:44 14.45 36  6  40  0  1K 24.95 27.29 30.5 0.00 0.0 0.0 0.0 0.0
13:34:47 14.79 35  5  40  0  1K 25.22 27.62 24.3 0.00 0.0 0.0 0.0 0.0
13:34:49 15.05 36  6  40  0  1K 25.44 27.88 29.0 0.00 0.0 0.0 0.0 0.0
13:34:52 15.31 36  6  40  0  1K 25.74 28.23 23.7 0.00 0.0 0.0 0.0 0.0
13:34:54 15.69 35  5  40  0  1K 26.06 28.61 42.5 0.00 0.0 0.0 0.0 0.0
13:34:57 16.08 36  6  40  0  1K 26.39 29.01 30.0 0.00 0.0 0.0 0.0 0.0
13:36:17 19.01 39  4  48  0  --- 28.73 31.86 ---- ---- --- --- --- ---
13:36:25 19.01 41  4  50  0  990 28.74 31.86 ---- ---- ---- ---- ----
13:36:32 19.01 40  4  49  0  990 28.74 31.87 0.00 0.00 0.0 0.0 3.7 0.0
    
```

Figure 140. ARDJ Report for a system running in 31-bit mode

### Field descriptions

Table 106. Fields in the ARD and ARDJ reports

Field Heading	Meaning
hh:mm:ss JOBNAME (ARD report)	The time the report was requested, and the one to eight character jobname associated with the address space.

## Mon II - ARD/ARDJ

Table 106. Fields in the ARD and ARDJ reports (continued)

Field Heading	Meaning
jobname TIME (ARDJ report)	The name of the job you requested. This column contains the time the report was requested.
DEV CONN	The device connect time (in seconds) used by the job. If it is greater than 99999 seconds, then it is expressed in hours (a four-digit number with a floating decimal point followed by H). If the device connect time exceeds 76.4 hours, three asterisks will appear in the field.
FF 16M	The number of fixed pages below 16 megabytes for the job.  If the data gathered is not valid, dashes will appear in this and the following fields.
FF 2G	The number of fixed pages for the job between 16 megabytes and 2 gigabytes. <b>Note:</b> This field is only shown for systems running in z/Architecture mode.
PRIV FF	The number of private non-LSQA pages for the job.
LSQA CSF	The number of private LSQA fixed pages for the job in central storage. This value includes fixed frames and private DREF pages.
X M	The cross memory address space indicator. When the field contains X, the line of data describes a cross memory address space; that is, an address space accessed primarily from other address spaces by means of cross memory functions. If it is not a cross memory address space, the field is blank.
C R	An indication whether WLM managed the address space as <i>storage critical</i> and/or <i>CPU critical</i> during the reporting interval. S Storage critical C CPU critical X Both storage and CPU critical
SRM ABS	The total SRM service absorption rate for the job. This field is reported only for address spaces that are currently in central storage. If no data is reported, dashes will appear in the field.
TCB TIME	The number of seconds of TCB processor time used by the current job step.
CPU TIME	The amount of processor (TCB + SRB) time, in seconds, for the current job step. When a valid delta value cannot be computed because the job has changed steps between requests for the report, this field contains dashes when delta mode is in effect.
EXCP RATE	The EXCP rate. This field always contains the rate since the last report request.
The following fields always contain a value that reflects the change since the last report request. They are reported only for address spaces that are currently in central storage.	
SWAP RATE	The page rate (the sum of pages in and pages out) for the job.
LPA RT	The common LPA page-in rate for the current transaction.
CSA RT	The common CSA page-in rate for the current transaction.
NVI RT	The private non-VIO page rate (the sum of pages in and pages out) for the current transaction.

Table 106. Fields in the ARD and ARDJ reports (continued)

Field Heading	Meaning
V&H RT	The hiperspace and private VIO page rate (the sum of pages in and pages out) for the current transaction.

### Report options for ARD, ASD, and ASRM

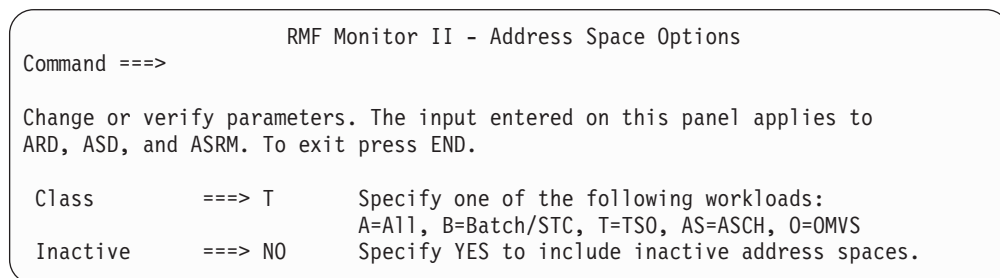


Figure 141. ARD, ASD, and ASRM Report Options Panel

You can specify the class, status, and domain you want the report for.

**Class** Allows you to specify the class. The default value is **A**.

**Inactive**

Allows you to specify the status of the address spaces shown on the report. The default is **NO** to include only active address spaces. **YES** causes all address spaces to be shown.

You can use the following commands on the Report Options panel:

**RESET**

Sets the default options.

**CANCEL**

Ends the option dialog without saving your changes.

Pressing the ENTER key checks the options. If no valid value is found, a message is issued. To resolve the error, either correct the value, or enter the CANCEL command.

**END**

Ends the dialog saving your changes.

### Report options for ARDJ, ASDJ, and ASRMJ

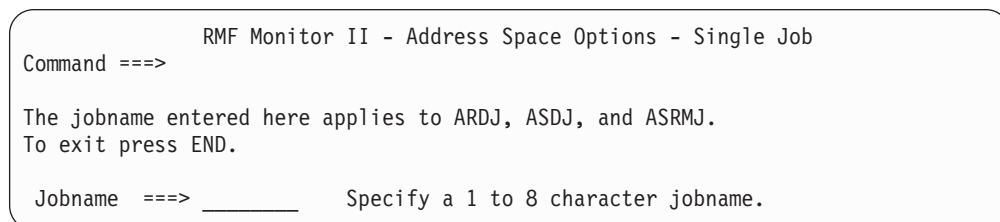


Figure 142. ARDJ, ASDJ, and ASRMJ Report Options Panel

**Jobname**

The jobname must:

- Be one to eight characters long
- Consist of the characters A-Z, 1-9, and the special characters @, #, and \$

The exception to these rules is \*MASTER\*, which is also a valid jobname.

---

## ASD/ASDJ - Address Space State Data report

The ASD/ASDJ report gives an overview of the current state of each address space in the system or each address space that meets the selection criteria that you specify when you request the report. Basically, the report tells you where each address space is and what it is doing.

You can use the ASD report, for example, to determine which jobs are using large amounts of central storage or which jobs are being swapped excessively and why the swapping is occurring.

Once a problem job has been identified, you can request an ASDJ report for that particular job. This enables you to focus your reporting on a known problem area.

If you have a workload delaying your application, you can check the workloads dispatching priority (**DP PR**) on the ASD report, and change it if necessary.

### How to request this report

Different methods are used to request the ASD and ASDJ reports.

#### How to request an ASD report

- In ISPF, specify **2** on the Address Space Report Selection menu.
- In TSO/E, use **PF2** to select the ASD report.
- Command interface:

##### Display session

```
ASD [class,status]
```

##### Background session

```
ASD [(class,status)]
```

#### How to request an ASDJ report

- In ISPF, specify **4** on the Address Space Report Selection menu.
- Command interface:

##### Display session

```
ASDJ jobname
```

##### Background session

```
ASDJ (jobname)
```

### Contents of the report

The information shown in an ASD and an ASDJ report is identical except the heading for the first column which is:

JOBNAME for the ASD report

TIME for the ASDJ report

In the ASD report of Figure 143 on page 251, the number of rows in the report depends on the number of address space identifiers that meet your selection criteria.

In the ASDJ report of Figure 144 on page 251, the number of rows depends on your requests to build a table of information. for a particular job.

**Note:** Information about SRM service is available in the address space SRM data (ASRM) report.

```

Command ==>>          RMF - ASD Address Space State Data          Line 1 of 53
                                                                Scroll ==>> HALF

                          CPU=  3/  3 UIC= 65K PR=  0          System= SYS1 Total

11:08:36              S C R DP  CS      CS TAR X  PIN TX SWAP  WSM
JOBNAME  SRVCLASS P L  LS PR   F      TAR WSS M  RT  SC  RV   RV

*MASTER* SYSTEM  1 NS   FF 6225          0  ---  0  0
PCAUTH  SYSTEM  1 NS   FF 141           0 X ---  0  0
RASP    SYSTEM  1 NS   FF 360           0 X ---  0  0
TRACE   SYSTEM  1 NS   FF 875           0 X ---  0  0
DUMPSRV SYSTEM  1 NS   FF 862           0  ---  0  0
    
```

Figure 143. ASD Report

```

Command ==>>          RMF - ASDJ Address Space State Data          Line 1 of 7
                                                                Scroll ==>> HALF

                          CPU=  2/  2 UIC= 65K PR=  0          System= SYS1 Total

RMFGAT              S C R DP  CS      CS TAR X  PIN TX SWAP  WSM
TIME  SRVCLASS P L  LS PR   F      TAR WSS M  RT  SC  RV   RV

11:03:44 SYSSTC  1 NS   FE 46.5K          0 X ---  0  0
    
```

Figure 144. ASDJ Report

### Field descriptions

Table 107. Fields in the ASD and ASDJ Report

Field Heading	Meaning
hh:mm:ss JOBNAME	The time the ASD report was requested, in the form hh:mm:ss, and the one to eight character jobname associated with the address space.
SRVCLASS	The service class name.
S P	Service class period.
C L	The field showing the current location of the job. The possible contents and their meanings are: <b>IN</b> In storage <b>OT</b> Swapped out and ready <b>LO</b> Logically swapped out <b>NS</b> Non-swappable <b>WM</b> Waiting for a resource (for example, DB2 latch or HSM recall): job is swapped in, is eligible for dispatching, and has accumulated no CPU time after some seconds <b>WL</b> Wait queue: long wait as a result of either WAIT TYPE=LONG or of STIMER for more than 0.5 seconds <b>WT</b> Wait queue: terminal wait <b>WO</b> Wait queue: reasons other than WM, WL, or WT <b>DL</b> TSO user delayed by SRM to meet response time objective <b>PR</b> Privileged <b>&gt;&gt;</b> In the process of being swapped out of storage <b>&lt;&lt;</b> In the process of being swapped into storage

## Mon II - ASD/ASDJ

Table 107. Fields in the ASD and ASDJ Report (continued)

Field Heading	Meaning
R LS	The field showing the reason for the last swap out associated with the job. This field is blank when the current location of the job is NS, IN, or PR. The possible contents and their meanings are: <b>TI</b> Terminal input wait <b>TO</b> Terminal output wait <b>LW</b> Long wait <b>XS</b> Auxiliary storage shortage <b>RS</b> Central storage shortage <b>DW</b> Detected wait <b>RQ</b> Requested swap <b>NQ</b> Enqueue exchange <b>EX</b> Exchange based on recommendation value <b>US</b> Unilateral <b>TS</b> Transition Swap <b>AW</b> APPC wait <b>IC</b> Improve central storage <b>IP</b> Improve system paging rate <b>MR</b> Make room to swap in an out-too-long user <b>IW</b> OMVS input wait <b>OW</b> OMVS output wait <b>SR</b> In-real swap
DP PR	The dispatching priority for the job.
CS F	The number of central storage frames assigned to the job. If a job is swapped out of central storage the number represents the number of central storage frames assigned to the job before the swap out occurred.
CS TAR	The SRM central storage target value. If no special monitoring is requested, this field is blank.
TAR WSS	The target working set size for the job (in number of pages).
X M	The cross memory address space indicator. When the field contains X, the line of data describes a cross memory address space; that is, an address space accessed primarily from other address spaces by means of cross memory functions. If it is not a cross memory address space, this field is blank.
PIN RT	The page-in rate (PIN). RMF calculates the page-in rate in one of two ways: <ul style="list-style-type: none"> <li>For cross-memory address spaces, the calculation is:  <math display="block">\text{PIN} = \frac{\# \text{ Page-ins}}{\text{Transaction Residency Rime}}</math> </li> <li>For all other address spaces, the calculation is:  <math display="block">\text{PIN} = \frac{\# \text{ Page-ins}}{\text{Seconds of CPU Time}}</math> </li> </ul> <p>The number of page-ins is the sum of the following:</p> <ul style="list-style-type: none"> <li>Number of pages brought into central storage one at a time</li> <li>Number of pages brought into central storage in blocks</li> <li>Hiperspace™ read miss count</li> <li>Number of hiperspace pages brought into central storage.</li> <li>Number of shared storage page-ins</li> </ul> <p>This field always contains the rate since the last report request. Dashes (---) in this field indicate that RMF is unable to calculate a value.</p> <p>If the transaction requires storage isolation, the value reported can be used to establish initial threshold values and to evaluate the effectiveness of these values.</p>
TX SC	The swap count for the current transaction.
SWAP RV	The workload manager recommendation value for the job. Note that only integer values are reported. Fractional values appear as zero. The range of possible values is -999 to +999, for details refer to the <i>z/OS MVS Initialization and Tuning Guide</i> .

Table 107. Fields in the ASD and ASDJ Report (continued)

Field Heading	Meaning
WSM RV	The recommended value for address spaces that are being managed by Working Set Management (WSM). The range of possible values is -6000 to +6000. The larger the value, the sooner the address space is likely to be swapped in. If the address space is not managed by WSM, this field is blank.

### Report options

The Report Options panel for the ASD report is the same as for the ARD report. See “Report options for ARD, ASD, and ASRM” on page 249 for a description.

---

## ASRM/ASRMJ - Address Space SRM Data report

The ASRM/ASRMJ report gives an overview of the system resources that are used by each address space in the system or each address space that meets the selection criteria that you specify when you request the report. The report gives, for example, information on processor service, storage service, and I/O service.

The report enables you to determine which jobs are using which services and whether certain jobs are creating performance problems by making excessive use of system services.

### How to request this report

Different methods are used to request the ASRM and ASRMJ reports.

#### How to request an ASRM report

- In ISPF, specify **3** on the Address Space Report Selection menu.
- In TSO/E, use **PF3** to select the ASRM report.
- Command interface:

##### Display session

```
ASRM [class,status]
```

##### Background session

```
ASRM [(class,status)]
```

#### How to request an ASRMJ report

- In ISPF, specify **6** on the Address Space Report Selection menu.
- Command interface:

##### Display session

```
ASRMJ jobname
```

##### Background session

```
ASRMJ (jobname)
```

### Contents of the report

The information shown in an ASRM and an ASRMJ report is identical except the contents of the first column which is:

JOBNAME for the ASRM report

TIME for the ASRMJ report

In the ASRM report shown in Figure 145 on page 254 the number of data lines depends on the number of address space identifiers in the system that meet your

criteria.

```

RMF - ASRM Address Space SRM Data
Command ==>>>
Line 1 of 84
Scroll ==>> HALF

CPU= 37/ 35 UIC=2540 PR= 0
System= SYS1 Total

08:19:48 S TRANS TRANS TX TX TX TX TX TX SESS
JOBNAME SRVCLASS P ACTIVE CUR RES CT SC CPU MSO IOC SRB TOTAL

*MASTER* SYSTEM* 1 441:25 441:25 1 0 126.2M 723.9M 6.889M 720980 857.7M
PCAUTH PROG001 1 441:25 441:25 1 0 1 127 0 0 128
RASP ADMIN005 1 441:25 441:25 1 0 1 82 0 16496 16579
TRACE PROG001 1 441:25 441:25 1 0 1 42 0 0 43
XCFAS ADMIN005 1 441:24 441:24 1 0 1.932M 4.247M 15 0 6.179M
GRS SYSTEM* 1 441:25 441:25 1 0 33.50M 168.4M 15 665060 202.5M
SMXC ADMIN005 1 441:25 441:25 1 0 1 11 0 0 12
SYSBMAS ADMIN005 1 441:25 441:25 1 0 1 54 0 0 55
DUMPSRV SYSTEM* 1 441:24 441:24 1 0 229522 864212 713169 6330 1.813M
CONSOLE PROG001 1 441:25 441:25 1 0 12.60M 28.08M 549981 196 41.23M
ALLOCAS PROG001 1 441:25 441:25 1 2 132 358 15 0 505
SMF SYSTEM* 1 441:24 441:24 1 0 18553 31601 2505 65 52724
VLF ADMIN005 1 441:24 441:24 1 0 2.846M 74.74M 580 0 77.59M
LLA ADMIN005 1 441:24 441:24 1 0 1.157M 14.37M 767327 1023 16.30M
JES3 PROG007 1 412:19 412:19 25 0 21.72M 702.7M 8.492M 12247 733.6M
    
```

Figure 145. ASRM Report

```

RMF - ASRMJ Address Space SRM Data
Command ==>>>
Line 1 of 14
Scroll ==>> HALF

CPU= 37/ 35 UIC=2540 PR= 0
System= SYS1 Total

BARE S TRANS TRANS TX TX TX TX TX TX SESS
TIME SRVCLASS P ACTIVE CUR RES CT SC CPU MSO IOC SRB TOTAL

15:44:27 TSOPRD 1 00:00 00:00:00 --- 0 0 0 0 0 0 -----
15:44:29 TSOPRD 1 00:00 00:00:00 27 0 0 0 0 0 0 606809
15:44:31 TSOPRD 2 00:00 00:00:00 28 0 267 1296 365 0 617346
15:44:32 TSOPRD 3 00:00 00:00:01 28 0 422 2165 745 0 618750
15:44:36 TSOPRD 4 00:00 00:00:06 28 0 6491 48528 2255 0 672692
15:44:49 TSOPRD 1 00:00 00:00:00 --- 0 0 0 0 0 0 -----
15:44:55 TSOPRD 1 00:00 00:00:00 --- 0 0 0 0 0 0 -----
15:45:05 TSOPRD 1 00:00 00:00:00 --- 0 0 0 0 0 0 -----
15:45:18 TSOPRD 3 00:00 00:00:00 34 0 551 2215 130 0 767129
15:45:22 TSOPRD 4 00:00 00:00:03 35 0 2281 15853 1780 0 792687
15:45:26 TSOPRD 4 00:00 00:00:06 35 0 6998 52654 2275 0 834700
15:45:31 TSOPRD 1 00:00 00:00:00 --- 0 0 0 0 0 0 -----
15:45:32 TSOPRD 1 00:00 00:00:00 --- 0 0 0 0 0 0 -----
15:45:33 TSOPRD 1 00:00 00:00:00 --- 0 0 0 0 0 0 -----
    
```

Figure 146. ASRMJ Report

### Field descriptions

Table 108. Fields in the ASRM and ASRMJ Report

Field Heading	Meaning
hh:mm:ss JOBNAME	The time the ASRM report was requested, in the form hh:mm:ss, and the one to eight character jobname associated with the address space.
JOBNAME TIME	The job name for the ASRMJ report. This column contains the time the report was requested.
SRVCLASS	The service class name.



Table 108. Fields in the ASRM and ASRMJ Report (continued)

Field Heading	Meaning	
S P	The service class period.	
TRANS ACTIVE	The transaction elapsed time; that is, the time that has elapsed since the current transaction began, in the form hhhh:mm	
TRANS CUR RES	The time, in the form hhhh:mm or hh:mm:ss, elapsed since the current transaction became resident.	
TX CT	The transaction count for the job. This field is reported only for address spaces that are currently in storage. Dashes indicate no data is reported.	
TX SC	The swap count for the current transaction. Asterisks indicate the number is too large to report.	
TX CPU	The processor service consumed by the current transaction.	An M, indicating millions of units absorbed, can follow the number.
TX MSO	The MSO service consumed by the current transaction.	
TX IOC	The I/O service consumed by the current transaction.	
TX SRB	The SRB service consumed by the current transaction.	
SESS TOTAL	The total SRM services consumed by the entire job.	

## Report options

The Report Options panel for the ASRM report is the same as for the ARD report. See "Report options for ARD, ASD, and ASRM" on page 249 for a description.

---

## CHANNEL - Channel Path Activity report

In general, the CHANNEL report gives you information about channel path activity for all channel paths in the system. The report contains data for every channel path that is online at the time you request the report.

Information about channel path activity, I/O device activity, and I/O request queuing information can be used to identify performance bottlenecks associated with the channel paths.

For all channels that are managed by **Dynamic Channel Path Management (DCM)**, additional information is available. DCM allows an installation to identify channels that they wish to be managed dynamically. These channels are not assigned permanently to a specific control unit, but belong to a pool of channels. Based on workload requirements in the system, these channels will be assigned dynamically by DCM. On top of the report, there is a consolidated data section for managed channels displaying the total number of channel paths for each type and the average activity data. The character **M** as suffix of the acronym for the channel path type is an indicator that the channel is managed by DCM.

## How to request this report

- In ISPF, specify 1 on the I/O Report Selection Menu.
- In TSO/E, use **PF4** to select the CHANNEL report.
- Command interface:

**Display session**

CHANNEL

Background session

CHANNEL

Special considerations of report output

You can obtain the report whether or not a Monitor I session measuring channel path activity is active. However, the channel path type appears only when RMF is active.

Data for total utilization and partition utilization is gathered independently. Because the internal interval used to gather this data is a few seconds, the total utilization and the sum of the partition's utilization sharing that channel might differ if a short RMF interval is specified. If the interval is too small and the appropriate data cannot be gathered, dashes (---) will be reported instead of data.

Contents of the report

```

RMF - CHANNEL Channel Path Activity
Command ==>>
Line 1 of 69
Scroll ==>> HALF

CPU= 37/ 35 UIC=2540 PR= 0 System= CB88 Total

08:01:56 Channel Utilization(%) Read(B/s) Write(B/s) FICON OPS zHPF OPS
ID No G Type S Part Tot Bus Part Tot Part Tot Rate Actv Rate Actv
4 *CNCSM 0.1 0.5
4 *FC_SM 0.0 0.0 0.0 0 0 0 0
12 OSD Y 0.0 0.0 0.0 2K 19K 0 0
14 OSD Y 0.0 0.0 0.0 5K 478K 458K 461K
16 OSD Y 0.4 1.3 0.0 493K 5M 3M 5M
20 CTC_S Y 0.0 0.0
27 CNC_S Y 0.0 0.0
2B CNC_S Y 0.9 3.7
2C CNC_S Y 0.2 0.6
30 5 FC_S Y 0.0 31.8 8.6 301 50M 71 227K 177 1 0 0
31 5 FC_S Y 0.0 31.6 7.9 239 46M 220 235K 177 1 0 0
37 4 FC_S Y 0.0 0.1 0.0 0 85K 0 26K 6 1 0 0
38 4 FC_S Y 0.0 0.1 0.0 0 69K 0 36K 7 1 0 0
39 4 FC_S Y 0.0 0.0 0.0 390 11K 0 15K 2 1 0 0
3A 4 FC_S Y 0.0 0.0 0.0 0 13K 0 16K 3 1 0 0
3E 4 FC_S Y 0.0 0.0 0.0 0 1K 0 568 0 1 0 0
7C CNCSM Y 0.4 1.8
7D CNCSM Y 0.0 0.1
81 3 FC_S Y 2.3 20.2 5.3 3M 30M 671K 2M 998 3 214 1
82 5 FC_S Y 0.1 0.9 0.3 147K 2M 282 89K 30 1 48 1
83 5 FC_S Y 0.1 0.9 0.3 162K 2M 291 86K 30 1 48 1
84 4 FC_S Y 0.0 0.0 0.0 56 223 0 0 1 1 0 0
85 3 FC_S Y 0.3 13.4 1.5 45K 4M 43K 2M 842 1 356 1
8C 3 FC_S Y 0.9 10.9 1.6 585K 7M 80K 1M 718 2 0 0
A6 5 FC_SM Y 0.0 0.0 0.0 0 0 0 0 0 0 0 0
B6 5 FC_SM Y 0.0 0.0 0.0 0 0 0 0 0 0 0 0
E0 IQD Y 0 346K
E1 IQD Y 0 0
E2 IQD Y 0 0
E3 IQD Y 0 0
    
```

Figure 147. CHANNEL Report

### Field descriptions

Table 109. Fields in the CHANNEL Report

Field Heading	Meaning
Channel ID	Hexadecimal channel path identifier (CHPID).
Channel No	For each channel type which is managed by DCM, a summary line is shown with the average values for all channels in this group. These summary lines are characterized by an * preceding the channel path type, and the number of channels of the group is displayed in column No.
Channel G	<p>Generation.</p> <p>The generation is used to differentiate between channels of the same channel type, when one has significant differences from the other. Newer generations with significant differences (for example, the channel throughput) are indicated by a number (1, 2, ...).</p> <p>For example, for a FICON channel, a number 1 indicates that the channel has an auto-negotiated throughput of 1Gbit/sec, or a number 4 indicates a throughput of 2Gbit/sec on a FICON Express4 card or a FICON Express2 card.</p>
Channel Type	<p>Type of channel path.</p> <p>You may issue the console command <code>D M=CHP(xx)</code> to see an explanation of the channel path type.</p> <p>If RMF encounters an error while processing the type, this field is blank. RMF continues to measure channel path activity. Check the operator console for messages.</p>
Channel S	The indication of whether a channel path is defined as shared between one or more logical partitions. Y indicates that the channel path is shared. A blank indicates it is not shared.
<p><b>Note:</b></p> <ol style="list-style-type: none"> <li>1. On a machine running in LPAR mode, but with only one LPAR defined, the <i>Part</i> columns for the <i>Read</i>, <i>Write</i> and <i>Utilization</i> fields display a zero value for channels of type FC (FICON).</li> <li>2. When Channel Path Measurement Facility (CPMF) is not available, for example, on z/OS systems running as z/VM guests, RMF uses sampled data from SRM so that the reported channel utilization is only an approximate value. With increasing channel speed, the channel utilization value becomes more and more inaccurate. Therefore, in such cases, RMF does not provide accurate values of FICON channel utilization.</li> </ol> <p>Beginning with z990 processors, the channel data from SRM is no longer available. As a result, the channel utilization data on a z/OS system running as z/VM guest, is reported as '-----'</p>	
Utilization (%) Part	<p>The channel path utilization percentage for an individual logical partition. RMF uses the values provided by CPMF.</p> <p>The calculation is:</p> $\text{Part Utilization (\%)} = \frac{\text{Channel Path Busy Time}}{\text{Channel Path Elapsed Time}} * 100$ <p>For channels like FICON, OSA Express, or OSA Direct Express, which are running in extended CPMF mode, the calculation is as follows:</p> $\text{Part Utilization (\%)} = \frac{\text{LPAR \# of Channel Work Units}}{\text{Max \# of Channel Work Units} * \text{Channel Path Elapsed Time}} * 100$ <p>For OSAEGbE, the value reflects the microprocessor utilization.</p> <p>For hipersockets, this value is not available.</p>

## Mon II - CHANNEL

Table 109. Fields in the CHANNEL Report (continued)

Field Heading	Meaning
Utilization (%) Tot	<p>The channel path utilization percentage for the CPC during an interval.</p> <p>For processors earlier than z990 and shared channels in LPAR mode, where CPMF is not available, the calculation is:</p> $\text{Utilization (\%)} = \frac{\text{Total \# SRM Observations of Channel Path Busy}}{\text{\# Samples}} * 100$ <p>For unshared channels in LPAR mode, the value for total utilization is the same as partition utilization.</p> <p>For channels like FICON, OSA Express, or OSA Direct Express, which are running in extended CPMF mode, the calculation is as follows:</p> $\text{Utilization (\%)} = \frac{\text{Total \# of Channel Work Units}}{\text{Max \# of Channel Work Units} * \text{Channel Path Elapsed Time}} * 100$ <p>For OSAEGbE, the value reflects the microprocessor utilization.</p> <p>For hipersockets, this value is not available.</p>
Utilization (%) Bus	<p>Percentage of bus cycles, the bus has been found busy for this channel in relation to the theoretical limit.</p> <p>For OSAEGbE, the value reflects the PCI bus utilization.</p> <p>For hipersockets, this value is not available.</p>
Read(B/s)	<p><b>Part</b> Data transfer rates from the control unit to the channel for this partition.</p> <p><b>Total</b> Data transfer rates from the control unit to the channel for the CPC.</p> <p>For hipersockets, this value is not available.</p>
Write(B/s)	<p><b>Part</b> Data transfer rates from the channel to the control unit for this partition.</p> <p><b>Total</b> Data transfer rates from the channel to the control unit for the CPC.</p>
FICON OPS	<p><b>Rate</b> Number of native FICON operations per second.</p> <p><b>Actv</b> The average number of native FICON operations that are concurrently active during the reporting interval.</p>
zHPF OPS	<p><b>Rate</b> Number of zHPF (High Performance FICON) operations per second.</p> <p><b>Actv</b> The average number of zHPF operations that are concurrently active during the reporting interval.</p>

## DEV/DEVV - Device Activity report

The Device Activity report gives information on I/O device use for all online devices you requested either by device class, by device number, or by volume serial number.

The Device Activity report, like the Monitor I session report, can help you to analyze device performance, to identify bottlenecks caused by a particular device, and to overcome obstacles that prevent efficient use of the resource.

Requesting the report during a display session enables you, for example, to track the device use on a real-time basis. You can get a timely picture of device use or track a specific critical device on a real-time basis, thus making it possible to take corrective action immediately.

To evaluate the data, you need to understand what a reporting period is and how it relates to the Monitor I interval.

### Evaluating details of cumulative mode output

The I= field in the header of each report shows the percentage of the Monitor I interval that has elapsed when RMF generates the Monitor II session report.

$$I = \frac{\# \text{ Samples Taken} * \text{ Cycle Time}}{\text{Monitor I Interval Length}}$$

For a report that reflects the total device activity (delta mode is off), the reporting period is the time that has elapsed from the start of the Monitor I interval to the time when you requested the report. The maximum reporting period is one Monitor I interval. When a Monitor II report covers a complete Monitor I interval, the I= field contains an upper-case 'T' (I=T).

### Evaluating details of delta mode output

The I= field in the header of the report equals the percentage of the interval that is represented by the data; thus, for your initial request, the I= field equals the percentage of the interval that expires between your initial request and the time you press the ENTER key. For all subsequent requests, the I= field equals the percentage of the interval that expires.

## How to request this report

Different methods are used to request the DEV and DEVV reports.

### How to request a DEV report

- In ISPF, specify 3 on the I/O Report Selection Menu.
- In TSO/E, use PF6 to select the DEV report.
- Command interface:

#### Display session

```
DEV [type ]
```

#### Background session

```
DEV [(type)]
```

### How to request a DEVV report

- In ISPF, specify 4 on the I/O Report Selection Menu.
- Command interface:

#### Display session

```
DEVV {VOLSER(volid)      }
      {NUMBER(device-number) }
```

#### Background session

```
DEVV {(VOLSER(volid))    }
      {(NUMBER(device-number)) }
```

## Special considerations of report output

The report is based on both hardware measurements and data collected during a Monitor I session. Therefore, a **Monitor I session must be active** when you issue your request.

## Mon II - DEV/DEVV

If no data is available, RMF issues a descriptive message.

Because the data comes from both hardware measurements and Monitor I session measurements, the data required to report some or all of the fields might be invalid or unavailable. A field based on data that is unavailable or invalid contains dashes (---).

The fields that RMF might not be able to report and the possible reasons for the unavailable or invalid data are:

Field Heading	Not reported when
ACTV RATE RESP TIME IOSQ TIME DB DELAY PEND TIME DISC TIME CONN TIME %DEV UTIL	The hardware measurements are not available either because there are not enough CMBs or because the channel measurement facility is disabled.
RESP TIME IOSQ TIME PEND TIME DISC TIME CON TIME %DEV UTIL	The device is attached to a byte multiplexor channel. Byte multiplexor channels collect only activity rate data; that is, the only hardware measurement available is the start subchannel count (SSCH).
LCU	RMF was unable to read the IOCDS.
ALL	The device moved online or offline during the Monitor I interval and total mode was requested.
STG GRP	Reported as <b>**CHGD**</b> when a device is added or deleted during the report interval.

If an \* appears immediately to the right of a field, a hardware measurement timer overflow has occurred. See "DEVICE - Device Activity report" on page 367 for an explanation of this condition.

**Note:** Devices are sorted ascending by an internally composed combination of the subchannel set ID (SSID) the device resides in, (not visible in the report), followed by the four-digit device number.

## Contents of the report

```

Command ==>>>
RMF - DEV Device Activity
Line 1 of 1513
Scroll ==>> CSR

CPU= 8/ 6 UIC=1190 PR= 0 System= SYSF Total

14:42:48 I=85% DEV
STG GRP VOLSER NUM PAV LCU ACTV RESP IOSQ -DELAY- PEND DISC CONN %D %D
RATE TIME TIME CMR DB TIME TIME TIME UT RV

SYC337 C337 1.0H 0094 0.000 .000 .000 .00 .00 .000 .000 .000 0 0
SYC338 C338 1.0H 0094 0.000 .000 .000 .00 .00 .000 .000 .000 0 0
SYC339 C339 1.0H 0094 0.000 .000 .000 .00 .00 .000 .000 .000 0 0
SYC33A C33A 1.0H 0094 0.000 .000 .000 .00 .00 .000 .000 .000 0 0
SYC33B C33B 1.0H 0094 0.000 .000 .000 .00 .00 .000 .000 .000 0 0
SYC33C C33C 1.0H 0094 0.000 .000 .000 .00 .00 .000 .000 .000 0 0
DB2 USC401 C401 1.0H 0095 0.001 .256 .000 .00 .00 .128 .000 .128 0 0
DB2 USC402 C402 1.0H 0095 0.001 .768 .000 .26 .00 .384 .000 .384 0 0
DB2 USC403 C403 1.0H 0095 0.001 .128 .000 .00 .00 .128 .000 .000 0 0
DB2 USC404 C404 1.0H 0095 0.033 1.40 .000 .00 .00 .201 .846 .349 0 0
DB2 USC405 C405 1.0H 0095 0.033 1.33 .000 .00 .00 .177 .881 .275 0 0
DB2 USC406 C406 1.0H 0095 0.033 1.24 .000 .00 .00 .211 .654 .369 0 0
DB2 USC407 C407 1.0H 0095 0.001 .384 .000 .00 .00 .128 .000 .256 0 0
DB2 USC408 C408 1.0H 0095 0.033 1.41 .000 .00 .00 .206 .881 .320 0 0
DB2 USC409 C409 1.0H 0095 0.001 .256 .000 .00 .00 .128 .000 .128 0 0
DB2 USC40A C40A 1.0H 0095 0.033 1.47 .000 .00 .00 .216 .876 .374 0 0
    
```

Figure 148. DEV Report

By default, the DEV report is sorted by LCU, unless you specify the storage group (SG) option. The SG option causes the DEV report to be sorted by device numbers within storage groups.

Type can be either a device class, or one or more volume serial numbers, device numbers, or storage group numbers.

When you request the report during a display session, the data line for any device that is more than 30% utilized is highlighted.

```

COMMAND ==>>>
RMF - DEVV DEVICE ACTIVITY
LINE 1 OF 11
SCROLL ==>> PAGE

CPU= 1/ 1 UIC= 65K PR= 0 SYSTEM= TRX2 TOTAL

I=55% DEV
TIME VOLSER NUM PAV LCU ACTV RESP IOSQ -DELAY- PEND DISC CONN %D %D
RATE TIME TIME CMR DB TIME TIME TIME UT RV

15:26:37 USC401 C401 1 0095 0.038 .592 .000 .00 .00 .156 .000 .436 0 0
15:26:39 USC401 C401 1 0095 0.038 .592 .000 .00 .00 .156 .000 .436 0 0
15:26:49 USC401 C401 1 0095 0.038 .592 .000 .00 .00 .156 .000 .436 0 0
15:27:02 USC401 C401 1 0095 0.040 .591 .000 .00 .00 .154 .000 .436 0 0
15:28:32 USC401 C401 1 0095 0.038 .594 .000 .00 .00 .152 .000 .441 0 0
15:30:58 USC401 C401 1 0095 0.120 .530 .000 .00 .00 .164 .000 .365 0 0
15:33:49 USC401 C401 1 0095 0.048 .570 .000 .00 .00 .162 .000 .407 0 0
15:34:32 USC401 C401 1 0095 0.047 .580 .000 .00 .00 .167 .000 .413 0 0
15:36:58 USC401 C401 1 0095 0.045 .586 .000 .00 .00 .161 .000 .424 0 0
15:37:48 USC401 C401 1 0095 0.040 .586 .000 .00 .00 .161 .000 .424 0 0
15:38:18 USC401 C401 1 0095 0.042 .597 .000 .00 .00 .158 .000 .438 0 0
    
```

Figure 149. DEVV Report

## Mon II - DEV/DEVV

The storage groups appear on the DEVV report only when the device has been assigned to a storage group. SG is not a valid option for DEVV.

### Field descriptions

Table 110. Fields in the DEV and DEVV Report

Field Heading	Meaning
STG GRP	The name that identifies the storage group to which the device belongs. For DEVV, this field is reported only when the volumes specified are members of a storage group. When a device is changed or deleted from a storage group during a report interval, RMF reports <b>**CHGD**</b> in this column.
I%	<p>The percentage of the Monitor I interval that has elapsed when RMF generates the Monitor II session report.</p> $I = \frac{\text{\# Samples Taken} * \text{Cycle Time}}{\text{Monitor I Interval Length}}$ <p>I% may contain values 0 through 99. When I% reaches 100, the field is set to contain an uppercase 'T'. I% continues to be set based on the above calculation.</p> <p>See "Evaluating details of cumulative mode output" on page 259 and "Evaluating details of delta mode output" on page 259 for additional information about the relationship between a Monitor II report period and a Monitor I interval.</p>
VOLSER	The volume serial number (for direct access and magnetic tape reports) of the volume mounted on the device at the end of the reporting interval.
DEV NUM	The hexadecimal device number that identifies a physical I/O device.
PAV	<p>The number of parallel access volumes (base and alias) which were available at the end of the reporting interval. If the number has changed during the reporting interval, it is followed by an '*'. If the device is a HyperPAV base device, the number is followed by an 'H', for example, 5.4H. The value is the average number of HyperPAV volumes (base and alias) in that interval.</p> $\text{Average \# of HPAV devices} = \frac{\text{Accumulated \# of HPAV devices}}{\text{Number of Samples}}$
LCU	<p>The hexadecimal identifier of the logical control unit (LCU) to which the device belongs. The set of devices associated with an LCU measurement are not the same on all processors because the definition of an LCU is model-dependent.</p> <p>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 - System Assist Processor), which is part of the channel subsystem, manages and schedules I/O work requests to the various devices within the LCU.</p> <p>There are two reasons that this field is blank:</p> <ul style="list-style-type: none"> <li>• RMF encountered an error while gathering data, check the operator console for messages.</li> <li>• This is a non-dedicated device in a z/VM guest system environment.</li> </ul>
ACTV RATE	<p>The rate at which start subchannel (SSCH) instructions to the device completed successfully.</p> $\text{ACTV RATE} = \frac{\text{\# Successful SSCH Instructions}}{\text{Interval}}$



Table 110. Fields in the DEV and DEVV Report (continued)

Field Heading	Meaning
RESP TIME	<p>The average number of milliseconds the device required to complete an I/O request. This value reflects the total hardware service time and the front end software queuing time involved for the average I/O request to the device. The channel measures active time, which starts at the acceptance of a SSCH instruction (indicated by a condition code 0) and ends at the acceptance of the channel end (primary status pending). It does not, however, include the time required to process the interruption. The IOS queue length is factored in to reflect the front end queuing time.</p> $\text{ACT TIME} = \frac{\text{Total Device Active Time}}{\text{Measurement Event Count}}$ $\text{RESP TIME} = \text{ACT TIME} + \text{IOSQ TIME}$ <p>The active time is the sum of connect, disconnect, and pending time as described later.</p>
IOSQ TIME	<p>The average number of milliseconds an I/O request must wait on an IOS queue before a SSCH instruction can be issued.</p> $\text{IOSQ TIME} = \frac{\text{IOSQ Count}}{\text{\# Monitor I Samples} \times \text{Device Activity Rate}}$
DELAY CMR	<p>The average number of milliseconds that a successfully initiated start or resume function needs until the first command is indicated as accepted by the device.</p> $\text{DELAY CMR} = \frac{\text{Initial Command Response Time}}{\text{Measurement Event Count}}$
DELAY DB	<p>The average number of milliseconds of delay that I/O requests to this device encountered because the device was busy. Device busy might mean that the volume is in use by another system, the device is reserved by another system, head of string busy condition caused the contention, or some combination of these conditions has occurred.</p> <p>A value is reported every 10 seconds.</p> <p>In a PR/SM environment, this value is updated every 20 seconds.</p> $\text{DELAY DB} = \frac{\text{Device Busy Delay Time}}{\text{Measurement Event Count}}$ <p>If the data is not valid, a dash (-) will be displayed.</p>
PEND TIME	<p>The average number of milliseconds an I/O request remains queued in the channel. This value reflects the time between acceptance of the SSCH function at the subchannel (SSCH-function pending) and acceptance of the first command associated with the SSCH function at the device (subchannel active). This value includes the time waiting for an available channel path and control unit as well as the delay due to shared DASD contention. If the value is high, refer to the device LCU entry in the I/O Queuing Activity report for an indicator of the major cause of the delay.</p> $\text{PEND TIME} = \frac{\text{Device Pending Time}}{\text{Measurement Event Count}}$
DISC TIME	<p>The average number of milliseconds the device was disconnected (not transferring data) while processing an SSCH instruction. Thus, this value reflects the time when the device was in use but not transferring data. It includes the overhead time when a device might disconnect to perform positioning functions such as SEEK/SET SECTOR as well as any reconnection delay.</p> $\text{DISC TIME} = \frac{\text{Device Disconnect Time}}{\text{Measurement Event Count}}$ <p>RMF calculates the total disconnect time by adding the pending time and connect time for the device and subtracting the result from the active time.</p>

## Mon II - DEV/DEVV

Table 110. Fields in the DEV and DEVV Report (continued)

Field Heading	Meaning
CONN TIME	<p>The average number of milliseconds the device was connected to a channel path and actually transferring data between the device and central storage. Typically, this value measures data transfer time but also includes the search time and the time needed to maintain channel path, control unit, and device connection.</p> <p style="text-align: center;">Device Connect Time</p> <p>CONN TIME = ----- Measurement Event Count</p>
% D UT	<p>The percentage of time during the interval when the device was in use. This percentage includes both the time when the device was involved in I/O operations (connect and disconnect time) and the time when it was reserved but not involved in an I/O operation.</p> <p>The percentage reported represents the time during the interval when the device is "tied up". When it could not be used to service a request from another system. Some small portion of device busy (reserved) time is missed when the device is reserved, but the I/O request is pending in the channel.</p> $\% D UT = \left( \frac{CONN + DISC}{INT * PAV} + \frac{RSV}{N} \right) * 100$ <p><b>CONN</b> Device connect time  <b>DISC</b> Device disconnect time  <b>RSV</b> Number of Monitor I samples when the device was reserved but not involved in an I/O operation  <b>INT</b> Monitor I interval time (seconds)  <b>PAV</b> The number of parallel access volumes (or 1 for a non-PAV device)  <b>N</b> Total number of Monitor I samples</p> <p>The % DEV UTIL field on a Device Activity report can exceed 100% for a device that is 100% utilized. This is because the device connect time from the channel measurement block is a longer time period than the RMF measurement interval. Therefore, it is possible that the value can be slightly higher than 100%.</p>
%D RV	<p>The percentage of time during the interval when a shared device was reserved by the processor on which RMF was started. The range of %D RV is 0 through 99%. When the device has reserved 100% of the interval, a T is shown in this field.</p> <p>At each RMF cycle, RMF checks to see if a device is reserved. If so, a counter is kept of all such samples. At the end of the interval, the percentage is computed.</p> $\% D RV = \frac{\# \text{ Device-reserved Samples}}{\# \text{ Monitor I Samples}}$

## Report options

```

RMF Monitor II - Device Activity Options

Command ==>

Specify one of the options below. For Volume, Device Number, and Storage
Group, a single name, a list of names, and a range of names is valid.
To exit press END.

Device Class ==> DASD_      Specify one of the following classes:
                             DASD, TAPE, COMM, CHRDR, UNITR or GRAPH

Volume          ==> _____ Ex: P500002:P50004,P50007

Device Number   ==> _____ Ex: 0580:0584

Storage Group   ==> _____ Ex: MANF13:MANF20
                             Specify SG to display all
                             Storage Groups.

```

Figure 150. DEV Report Options Panel

On the Report Options panel of the DEV report, you can specify one of the four options device class, volume, device number, or storage group.

### Device Class

Allows you to measure all devices in a certain class.

If you leave the panel empty, the device class is the default, and the class DASD will be used.

### Volume

If you want a report on a specific volume or volumes, you can specify volume numbers as a single number (aaaaaa), a range of numbers (aaaaaa:zzzzz), or a list of numbers (aaaaaa,bbbbbb,dddddd).

### Device Number

If you want a report on a specific device, you can specify a single number, a range of numbers, or a list of numbers.

Device numbers are hexadecimal and four characters long.

### Storage Group

If you want a report on a specific storage group or storage groups, you can specify a single storage group, a range of storage groups, or a list of storage groups.

Storage group names are one to eight characters.

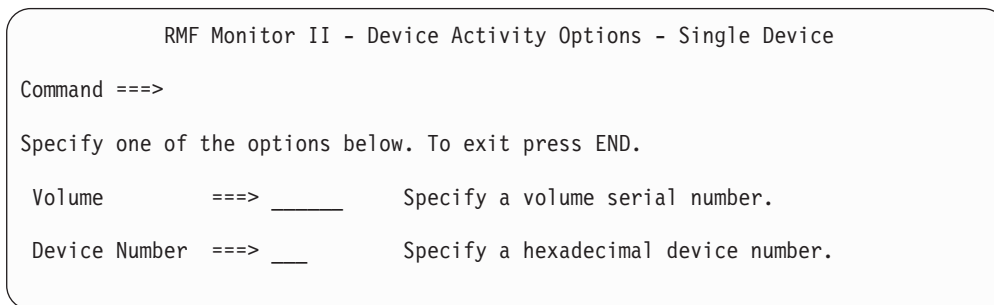


Figure 151. DEVV Report Options Panel

**Volume**

Allows you to specify specific DASD or tape devices in the form aaaaaa. You can specify a single volume, in the format aaaaaa, a list of volumes in the format aaaaaa,bbbbbb,ccccc, ... or a range of volumes in the format aaaaaa:zzzzzz, where aaaaaa is the first volume and zzzzzz is the last. Your entry cannot exceed 32 characters, including commas and colons.

The name is restricted to the characters A-Z, 0-9, @, # and \$.

**Device Number**

Allows you to specify a four-digit hexadecimal device number.

---

## HFS - Hierarchical File System Statistics report

The HFS report provides data for capacity planning and for basic performance analysis and problem determination:

- A general understanding of the throughput recognized and achieved by HFS allows you to optimally use your resources.
- The ability to display performance statistics of HFS enables you to identify potential problems and bottlenecks within the HFS component and to take corrective actions.

### How to request this report

- In ISPF, specify 5 on the I/O Report Selection Menu.
- Command interface:

**Display session**

HFS [hfsname]

## Contents of the report

```

Command ==>          RMF - HFS File System Statistics          Line 1 of 12
                      Scroll ==> PAGE

                      CPU= 37/ 35 UIC=2540 PR=  0           System= SYSA Delta

File System Name: OMVS.SYS1.ROOT
Mount Date: 07/28/2016  Time: 11:29:55           Elapsed Time: 00:01:40

----- Allocation (MB) -----      ----- Index Events -----
System          914  Data          440  New Level      10
Attr. Dir       3.50  Cached          16  Splits         50  Joins          3

          ----- File I/O -----      --- Metadata I/O ---      ---- Index I/O ----
                Count   Rate          Count   Rate          Count   Rate
Cache           1300  13.000           300  3.000           280  2.800
DASD             57  0.570             10  0.100            20  0.200
Hit Ratio       95.80                96.77                93.33
Sequential     1234  12.340
Random          123  1.230
    
```

Figure 152. HFS Report

## Field descriptions

Table 111. Fields in the HFS Report

Field Heading	Meaning
General Section	
File System Name	The name of the file system which has been selected for reporting.
Mount Date	Date when the file system has been mounted (mm/dd/yyyy).
Time	Time when the file system has been mounted (hh:mm:ss).
Elapsed Time	Delta mode: Time between two consecutive clickings on Enter.  Total mode: Time since the file system has been mounted.  The format can be in days and hours (6d 19h) or in hh:mm:ss.
Allocation - All values in megabytes	
System	Amount of storage allocated to this HFS.
Data	Amount of storage internally used within HFS for data files, directories and HFS internal structures like the attribute directory (AD).
Attr. Dir	Amount of storage used for the attribute directory (AD). This number is included in the "Data" field.  The attribute directory is the internal HFS structure (index) which contains attribute information about individual file system objects as well as attributes of the file system itself.
Cached	Amount of data buffer storage cached by this file system.
Index Events	
New Level	Number how often HFS added a new level to its index structure.  The index statistics are relative to all of the indices in the HFS data set. The attribute directory (AD) is one index (the largest) but each directory (including the root) is also an index.
Splits	Number how often an index page was split into two pages because new records were inserted. This gives an idea of how much insertion activity there has been for the index structure.
Joins	Number how often HFS was able to combine two index pages into one, because enough index records had been deleted in the two pages.
File I/O - all values are reported as <b>Count</b> and <b>Rate</b> (counts per second).	
Cache	Number of times the first page of a data file was requested and found in virtual storage (cache).

## Mon II - HFS

Table 111. Fields in the HFS Report (continued)

Field Heading	Meaning
DASD	Number of times the first page of a data file was requested but was not found in the cache, therefore an I/O was necessary.
Hit Ratio	Percentage of cache-found requests based on total number of requests.
Sequential	Number of sequential I/O requests.  A sequential I/O is one of a series of I/Os to read or write a data file, where the first I/O started at the first byte of the file and each subsequent I/O was for the next sequential set of bytes.
Random	Number of random I/O requests.  A random I/O is an I/O that does not read or write the start of a file, and was not preceded by an I/O that read or wrote the immediately preceding set of bytes.
Metadata I/O - all values are reported as <b>Count</b> and <b>Rate</b> (counts per second).	
Cache	Number of times the metadata for a file was found in the cache during file lookup.
DASD	Number of times the metadata for a file was not found in the cache during file lookup and an index call was necessary.
Hit Ratio	Percentage of cache-found requests based on total number of requests.
Index I/O - all values are reported as <b>Count</b> and <b>Rate</b> (counts per second).	
Cache	Number of index page read or write hits.
DASD	Number of index page read or write misses.
Hit Ratio	Percentage of cache-found requests based on total number of requests.

## Report options

```

RMF Monitor II - HFS Report Options           Line 1 of 19
Command ==>>>                               Scroll ==>> HALF

Select (S) or fill-in a file system name. To exit press END.

Selected file system name: OMVS.SYS5.S670CB1.BOOKSRV
Number of mounted file systems: 19           Display: YES   (YES/NO)

You can use FIND to search for a specific HFS file system name.

Sel  HFS File System Name
-    OMVS.SYS5.ROOT
-    OMVS.SYS5.S670CB1.ADSM
-    OMVS.SYS5.S670CB1.BIN
-    OMVS.SYS5.S670CB1.BOOKSRV
-    OMVS.SYS5.S670CB1.DCEAS
-    OMVS.SYS5.S670CB1.DCEBASE
-    OMVS.SYS5.S670CB1.DCEDFS.GLOBAL

```

Figure 153. HFS Report Options Panel

Table 112. Fields in the HFS Report Options Panel

Field Heading	Meaning
Selected file system name	The currently selected file system name to be reported on (up to 44 characters). This field is an input field and can be overtyped. Any data set name which adheres to the MVS rules for data set names is accepted (fully-qualified without enclosing quotes).
Number of mounted file systems	The number shows how many file systems are currently mounted (including HFS and other file systems).

Table 112. Fields in the HFS Report Options Panel (continued)

Field Heading	Meaning
Display	<b>Yes/No</b> specifies whether the names of the currently active file systems should be listed for selection. <b>No</b> is the default. <b>Note:</b> If the number of HFS file systems which are currently mounted is high it may take a while to provide the requested names.
Sel	An "S" can be placed in front of the file system name to be selected. Putting an "S" in this column results in replacing the file system name in the header field "Selected file system name".
HFS File System Name	The name of a file system which was found active. The file system names are sorted in alphabetical order.

## ILOCK - IRLM Long Lock Detection report

Services of the IMS/VS Resource Lock Manager (IRLM) are used by IMS to serialize application program requests for data base records to ensure that two programs do not access the same record for update at the same time.

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

Excessive use of a resource on one instance can suspend the work on other systems. To avoid such locking situations or, in the worst case, a re-IPL, the report provides information for IMS operators to perform the necessary actions to eliminate the problem.

### How to request this report

- In ISPF, specify **9** on the Resource Report Selection menu.
- Command interface:

**Display session**

```
ILOCK [ALL]
```

### Special considerations

There is no data gathering component for this report. Instead, the retrieval of the IRLM data from the RMF SMF data buffer is done by the reporter. To have the data available in the SMF data buffer (SMF record type 79 subtype 15), it is necessary to specify this option explicitly, for example:

```
S RMF,,, (SMFBUF(RECTYPE(70:78,79(15))))
```

For details, please refer to the *z/OS RMF User's Guide*.

Data collection is initiated by the operator who enters at the console the **runtimeo-exit** for one system in the sysplex:

```
F irlmid,RUNTIMEO
```

The command will be propagated automatically to all other systems.

When the SMF records are eventually written by the IRLMs in the data sharing group, the reporter can fetch these SMF records out of the RMF SMF data buffer.

As a consequence, you have to ask the operator to issue this command if you get informed that there is no data available for the report.

**Note:** Access to the SMF data buffer requires appropriate security authorization. Please, refer to chapter *Setting Up RMF - Access Definitions* in the *z/OS RMF User's Guide* for details.

## Contents of the report

```

Command ==>>          RMF - ILOCK IRLM Long Lock Detection          Line 1 of 15
                                                                Scroll ==>> HALF

                          CPU= 37/ 35 UIC=2540 PR=  0              System= RMF5 Total

State  Type  Lock_Name          PSB_Name  Elap_Time  CICS_ID
      IMS_ID Recovery-Token      PST# Trx/Job  Wait_Time  DB/Area
-----
CF Structure ACOXLOCK          at 07/28/2016 13:02:10 Deadlock Cycle 00002EC7
-----
TOP      BMP      09C943CFA7800101D7000000000000 DFSSAMB1  00:06:04
BLOCKER  AC03      AC03      0000000300000000      0006  IRLMTPZ
-----
TOP      BMP      09C3614505800101D7000000000000 DFSSAMB1  00:06:09
BLOCKER  AC01      AC01      0000000600000000      0006  IRLMTPA
-----
WAITER   BMP      09C3614505800101D7000000000000 DFSSAMB2
        AC02      AC02      0000000800000000      0007  IRLMWT1  00:05:52  DI21PART
-----
WAITER   BMP      09C943CFA7800101D70000000000000 DFSSAMB7
        AC02      AC02      0000000900000000      0008  IRLMWT2  00:05:42  DI21PART
-----

```

Figure 154. ILOCK Report

## Field descriptions

Table 113. Fields in the ILOCK Report

Field Heading	Meaning
CF Structure	The name of the coupling facility structure used by IRLM.  The date/time field shows the time when the SMF record for the first displayed lock entry was written.
Deadlock Cycle	The hexadecimal deadlock cycle number generated by IRLM and passed to IMS. As IMS requests longlock data for two cycles, this value is used to uniquely identify when the data was gathered.
State	State distinguishes between a blocker (holder) and a waiter: <b>BLOCKER</b> Transaction holds a lock at the time the exit is driven. <b>TOP BLOCKER</b> A blocker which holds a resource and has waiter(s) waiting on him, but is not found elsewhere as a waiter in another blocker's wait list. This might be the most likely transaction to kill to let everyone else run. <b>WAITER</b> Transaction is waiting for a lock. <b>BLOCKER/WAITER</b> The transaction was found as a blocker and waiter. <b>Note:</b> To display all blocker and waiters, you have to call the ILOCK command with the parameter ALL, otherwise TOP BLOCKERs will be shown, only.
Type	Identifies the region type a transaction can execute in.  Types are DBCTL (DB control), BMP (batch message processing), IFP (fast path), MPP (message processing region), SYPST (fast path system service ITASK), BATCH, and CICS.
IMS_ID	Name given to the IMS region at the time it is brought up.



Table 113. Fields in the ILOCK Report (continued)

Field Heading	Meaning
Lock_Name	The unique identifier used by IMS to obtain a lock on a resource. This name varies between 9-11 characters and contains the DMB#/DCB#/RBA of the resource (data) we are requesting a lock for.
Recovery-Token	Recovery Token - a 16-byte token used to uniquely identify a unit of work.
PST#	Partition Specification Table (PST) Number.  As the IMS region is initially brought up, several PSTs are initialized, and each is assigned a unique ID (or PST number).  The PST block is the primary block used to dispatch transactions in IMS, and the PST number is used to uniquely identify each transaction.
PSB_Name	Partition Specification Block Name - the name given to a PSB (Program Specification Block) at PSBGEN time. This block is used to define which segments a particular application can have access to.
Trx/Job	The name of the transaction in a BMP or MPP region, or the job name for all of the remaining region types.
Elap_Time	The field is available for blockers, it contains the elapsed time between the time the PST was scheduled, or the unit of work (UOW) was created, and the time the 79.15 record was created for this entry.
Wait_Time	The field is available for waiters, it contains the elapsed time between the time that IRLM processed the request for the resource and the time that the 79.15 record was created for this entry.
CICS_ID	CICS task identifier - an 8-digit ID generated by CICS and passed to IMS at the time a CICS application is scheduled.
DB/Area	A name given to a data base (IMS full-function) or an area (IMS fastpath) at DBDGEN time. It is used to uniquely identify the data base or area the lock is held on by this transaction.

## IOQUEUE - I/O Queuing Activity report

The IOQUEUE report provides information, grouped by LCU (logical control unit), on the I/O configuration. The information includes contention rate, queue lengths, and percentages of time when one or more I/O components were busy. Information about the LCU is useful because the LCU is the focus of I/O configuration and path management measurements for a related group of I/O devices.

For all channels that are managed by Dynamic Channel Path Management (DCM), additional information is available. DCM allows an installation to identify channels which they wish to be managed dynamically. These channels are not assigned permanently to a specific control unit, but belong to a pool of channels. Based on workload requirements in the system, these channels will be assigned dynamically by DCM. For each LCU with DCM managed channels, a summary line displays the minimum and maximum number of connected DCM managed channels, the number of defined DCM managed channels and accumulated activity data.

| PAV base mode is the mode when alias devices are assigned to one PAV base  
| device. An I/O for a PAV base device is executed using aliases assigned to that  
| PAV base device.

| HyperPAV mode is the mode when a pool of alias devices is assigned to one LCU.  
| An I/O for a PAV base device can be executed using any alias device of that pool.

| SuperPAV mode is the mode, when a pool of alias devices is assigned to one LCU  
| and multiple LCUs are grouped into one Alias Management Group (AMG). An

I I/O for a PAV base device can be executed using any alias device of these multiple  
I alias pools. The favored way is to use the alias device assigned to the same LCU  
I (home LCU) the PAV base device is assigned to.

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 - System Assist Processor), which is part of the channel subsystem, manages and schedules I/O work requests to the various devices within the LCU.

### Using the information given in the report

This report can tell you about the cause of performance problems associated with channel paths and devices. You could, for example, find the reason for an unusually long pending time reported on the device report. Check the relationship between the percentage of requests deferred for device busy and control unit busy for the LCU on the IOQUEUE report.

To help you determine the best way to fix a performance problem related to an LCU, you can request the report during a display session. This tracks the I/O queuing on a real-time basis.

### Evaluating details of cumulative mode output

The I= field in the heading of each report shows the percentage of the Monitor I interval that has elapsed when RMF generates the Monitor II session report.

$$I = \frac{\# \text{ Samples Taken} * \text{Cycle Time}}{\text{Monitor I Interval Length}} * 100$$

For a report that reflects the total device activity (DELTA mode off), the reporting period is the time that has elapsed from the start of the Monitor I interval to the time when you requested the report. The maximum reporting period is one Monitor I interval. When a Monitor II report covers a complete Monitor I interval, the I= field contains an upper-case "T" (I=T).

### How to request this report

- In ISPF, specify 2 on the I/O Report Selection Menu.
- Command interface:

#### Display session

IOQUEUE [type]

#### Background session

IOQUEUE [(type)]

### Special considerations of report output

The report depends on data that the Monitor I session collects. To get this report, the Monitor I I/O Queuing Activity report must be active. The Monitor I gatherer gets a new set of model dependent data every second or every cycle, whichever time period is greater.

If the hardware measurements are not available, the channel measurement facility is not available. If there is a failure in the diagnose interface, RMF does not provide model-dependent data generated by the hardware for the following fields:

- CONTENTION RATE
- DELAY Q LENGH
- CHPID TAKEN
- %CU BUSY

If the data is not reliable (indicated by a successive invalid sample count greater than zero), RMF does not provide model-dependent data generated by the hardware for the following fields:

- ACTIV RATE
- AVG Q LENGH
- %REQ DEFER

Data items that are not valid are marked by dashes (---) in the output display.

When an LCU has no activity during the interval, it is omitted from the report. If channel paths were brought online or taken offline during the interval, data is still formatted, but only for the channel paths and control units that were online and had some connection to a device or set of devices of the LCU at the time the report was requested appear in the report.

In a z/VM guest system environment, the report for an z/OS system that is authorized via the VM RMCHINFO directory option, shows static configuration data. Measurement data is not available.

## Contents of the report

RMF - IOQUEUE I/O Queuing Activity										Line 92 of 684					
Command ==>										Scroll ==> CSR					
CPU= 3/ 3 UIC=2540 PR= 0										System= S5C Total					
03:20:32	I= 11%	DCM Group			Cont	Del	Q	AVG	CHPID	%DP	%CU	AVG	AVG		
Path	DCM	CTL	Units	MN	MX	DEF	LCU	Rate	Lngh	CSS	Taken	Busy	Busy	CUB	CMR
D6		5F00					0048				0.39	0.0	0.0	0.0	0.2
							0048	0.0	0.00	0.4	0.85	0.0	0.0	0.0	0.2
B0	PF	8000					0069				84.24	0.0	0.0	0.0	0.2
B1	PF	8000					0069				82.79	0.0	0.0	0.0	0.2
B2	NP	8000					0069				0.00	0.0	0.0	---	---
B3	NP	8000					0069				0.00	0.0	0.0	---	---
95	PF	8000					0069				81.48	0.0	0.0	0.0	0.2
							0069	0.0	0.00	0.4	248.52	0.0	0.0	0.0	0.2
B0	NP	8100					006A				0.00	0.0	0.0	---	---
B1	NP	8100					006A				0.00	0.0	0.0	---	---
B2	PF	8100					006A				122.94	0.0	0.0	0.0	0.2
B3	PF	8100					006A				122.42	0.0	0.0	0.0	0.2
95	NP	8100					006A				0.00	0.0	0.0	---	---

Figure 155. IOQUEUE Report

### Field descriptions

Table 114. Fields in the IOQUEUE Report

Field Heading	Meaning
Path	<p>The hexadecimal channel path identifier (CHPID) of the online channel path attached to the physical control units in the LCU. There can be up to eight channel paths in an LCU. Only channel paths identified in the Monitor I report as ONLINE to the system and having connection to a device or group of devices of the LCU appear in the Monitor II report.</p> <p>If applicable, the path attribute is indicated with the CHPID:</p> <p>PF preferred path                      NP non-preferred path                      NS path attribute not specified</p> <p>For devices residing in control units that do not support path attributes, only the CHPID is displayed.</p>
DCM	<p>If the channel path is under control of DCM, this is indicated by a Y in this column. The activities of all DCM channels belonging to the same LCU will be summarized in a separate line.</p>
CTL Units	<p>The hexadecimal identifier of each physical control unit associated with an online channel path in the LCU group.</p>
DCM Group	<p>The values in columns MIN MAX DEF report the minimum and maximum number of DCM managed channels for one LCU (in this interval) as well as the installation-specified definition for this LCU.</p> <p>The line with these values is available only for LCUs with DCM managed channels. It contains in addition the accumulated values of the I/O activity rate, the director port contention, and the control unit contention of all DCM managed channels. These values may include also measurements of managed channels which were partially online.</p>
LCU	<p>The hexadecimal identifier of the logical control unit (LCU).</p> <p>An LCU is the set of devices attached to the same physical control unit or a group of physical control units with one or more devices in common. Each physical control unit and each device can belong to only one LCU. They cannot be shared between LCUs.</p>
Cont Rate	<p>The rate at which the SAP places delayed I/O requests on the CU-HDR for this LCU. The SAP places an I/O request on the CU-HDR when all paths to the subchannel are busy and at least one path to the control unit is busy. For devices with only one path, or for devices where multiple paths exist and the busy condition is immediately resolved, the SAP does not count the condition.</p> $\text{Cont Rate} = \frac{\text{\# Enqueued Requests}}{\text{Monitor I Interval Time}}$
Del Q Lngth	<p>The average number of delayed requests on the control unit header (CU-HDR). Each time a request is enqueued from the CU-HDR, RMF counts the number of requests remaining on the queue and adds that number to the accumulator. At the end of the interval, RMF divides the total number of accumulated queued requests by the number of times a request was enqueued.</p> $\text{Del Q Lngth} = \frac{\text{Accumulated Queue Length} - \text{\# Enqueued Requests}}{\text{\# Enqueued Requests}}$
AVG CSS	<p>The average number of milliseconds of delay that an I/O request encountered after the acceptance of the start or resume function at the subchannel for the LCU, until the channel subsystem's first attempt to initiate the operation.</p> $\text{AVG CSS} = \frac{\text{Channel Subsystem Time}}{\text{\# I/O Operations Accepted}}$
CHPID Taken	<p>The rate at which I/O requests to devices of this LCU are satisfied by each CHPID during the interval. By reviewing the rate at which each channel path of the LCU satisfies I/O requests, you can see how evenly the work requests are distributed among the available paths and how effectively those paths are arranged for the LCU.</p> $\text{CHPID Taken} = \frac{\text{\# I/O Operations on that Path}}{\text{Monitor I Interval Time}}$

Table 114. Fields in the IOQUEUE Report (continued)

Field Heading	Meaning
% DP Busy	<p>This field indicates director port contention. It is the number of times an I/O request was deferred because the director port was busy during the measurement interval.</p> $\% \text{ DP Busy} = \frac{\text{DPB}}{\text{DPB} + \text{CUB} + \text{SUC}} * 100$ <p><b>DPB</b> Number of deferred I/O requests due to director port busy  <b>CUB</b> Number of deferred I/O requests due to control unit busy  <b>SUC</b> Number of successful I/O requests on that path</p>
% CU Busy	<p>This field shows the relationship for each channel path of the LCU, between requests deferred due to control unit busy and total successful requests serviced by that path. Each CHPID of the LCU measures the distribution of control unit contention.</p> $\% \text{ CU Busy} = \frac{\text{CUB}}{\text{DPB} + \text{CUB} + \text{SUC}} * 100$ <p><b>DPB</b> Number of deferred I/O requests due to director port busy  <b>CUB</b> Number of deferred I/O requests due to control unit busy  <b>SUC</b> Number of successful I/O requests on that path</p>
AVG CUB	<p>The average number of milliseconds of delay that an I/O request encountered for the channel path because the control unit was busy.</p> $\text{AVG CUB} = \frac{\text{Control Unit Busy Time}}{\# \text{ I/O Operations Accepted on that Path}}$
AVG CMR	<p>The average number of milliseconds of delay that a successfully initiated start or resume function needs until the first command is indicated as accepted by the device. It allows to distinguish between real H/W errors versus workload spikes (contention in the fabric and at the destination port).</p> $\text{AVG CMR} = \frac{\text{Initial Command Response Time}}{\# \text{ I/O Operations Accepted on that Path}}$

## Report options

```

RMF Monitor II - I/O Queuing Activity Options

Command ==>

Specify one of the options below. For LCU number, a single number,
a list of numbers, and a range of numbers is valid. To exit press END.

Device Class ==> DASD_ Specify one of the following classes:
                    DASD, TAPE, COMM, CHRDR, UNITR OR GRAPH

LCU Number   ==> _____ Ex: D:F,4E,55
    
```

Figure 156. IOQUEUE Report Options Panel

You can specify either a class or a device number.

### Device Class

Allows you to specify the device class. If you leave this field empty, RMF uses DASD.

### LCU Number

Allows you to request specific logical control unit numbers. The numbers must be in three-digit hexadecimal format. You can specify any combination of a single number, a list of numbers, or a range of numbers. Your entry must not exceed 32 characters, including commas and colons.

## LLI - Library List report

The information shown in the LLI report provides the status of the key system libraries that are defined in the following lists:

- Load module link list
- Pageable link pack area list
- List of authorized libraries (APF list)

This information can help you to check whether the status of these libraries is correct for your current environment.

### How to request this report

- In ISPF, specify **L** on the Monitor II Primary Menu. This leads you to the Library List and OPT Settings Selection Menu. Here you can select:

**1 Link list**

LNKLSTxx - Link Library List

**2 LPA list**

LPALSTxx - LPA Library List

**3 APF list**

IEAAPFxx - Authorized Program Library List

- In the command interface of an ISPF or TSO/E display session, specify:

**LLI** for the Link Library List

**LLI LPA**

for the LPA Library List

**LLI APF**

for the Authorized Program List

In addition, an optional parameter **A** can be specified to create a report with more details. Table 115 on page 278 lists the values that are available with this parameter only. Due to performance reasons, the default for each command is just to collect basic information.

### Contents of the report - Link Library List

```

RMF - LLI Program Library Information                               Line 1 of 38
Command ==>>>                                                    Scroll ==>> HALF

                                CPU= 37/ 35 UIC=2540 PR=  0           System= RMF8 Total

----- Link Library List <IPL                                     > -----
DevNum  DevType  Volser APF Ext  Data Set Name
020F    33903    630D14 Y   3  SYS1.LINKLIB
020F    33903    630D14 Y   2  SYS1.MIGLIB
020F    33903    630D14 Y   3  SYS1.CSSLIB
0975    33903    SMSRMF Y   1  RMF530.GRSREP.LINKLIB
020F    33903    630D14 Y   1  SYS1.SHASLINK
020F    33903    630D14 Y   1  SYS1.SHASMIG
020E    33903    MVSTGT Y   1  SYS1.RMF.V630.D05.LINKLIB
031A    9345-2    RMFUSR Y   1  DRIVER.SYS1.LINKLIB
    
```

Figure 157. LLI Report - Link Library List

The link library list is determined by Parmlib members LNKSTxx or PROGxx. It represents the current active link list set. The list contains the names of all link libraries as well as an indication whether the library is authorized (APF) or not.

The header line Link Library List contains the indication about the status of the link list:

**IPL** The system is currently running with the link list set that has been selected during IPL.

**name** Name of the current link list set.

## Contents of the report - LPA Library List

```

RMF - LLI Program Library Information                               Line 1 of 7
Command ==>                                                    Scroll ==> HALF

                        CPU= 37/ 35 UIC=2540 PR=  0                System= RMF8 Total

----- LPA Library List -----
DevNum  DevType  Volser  Data Set Name
020F    33903    630D14  SYS1.LPALIB
0975    33903    SMSRMF  RMF530.GRSREP.LPALIB
020E    33903    MVSTGT  SYS1.RMF.V630.D05.LPALIB
020F    33903    630D14  SYS1.ISAMLPA
020F    33903    630D14  SDSF.ISFLPA
020E    33903    MVSTGT  SYS1.TCP.SEZALPA
020E    33903    MVSTGT  SYS1.REXX.V130.SEAGLPA
    
```

Figure 158. LLI Report - LPA Library List

The pageable link pack area list is determined by Parmlib members LPALSTxx. It is a fixed-size list which is pointed to from the CVT. The list contains the names of all link libraries that reside in the PLPA. These modules are authorized (APF) by default.

## Contents of the report - APF Library List

```

RMF - LLI Program Library Information                               Line 1 of 47
Command ==>                                                    Scroll ==> HALF

                        CPU= 37/ 35 UIC=2540 PR=  0                System= RMF8 Total

----- APF List - Format STATIC -----
DevNum  DevType  Volser  SMS RACF  Data Set Name
020F    33903    630D14  N  N      SYS1.LINKLIB
020F    33903    630D14  N  N      SYS1.SVCLIB
020F    33903    630D14  N  N      SYS1.LINKLIB
020F    33903    630D14  N  N      SYS1.ISAMLPA
020F    33903    630D14  N  N      SYS1.VTAMLIB
020F    33903    630D14  N  Sec=?  SYS1.JES3LIB
020F    33903    630D14  N  N      SYS1.NFSLIB
020F    33903    630D14  N  N      SYS1.SCEERUN
020F    33903    630D14  N  N      SYS1.SEPWMOD1
020F    33903    630D14  N  N      SYS1.SEPWMOD2
    
```

Figure 159. LLI Report - APF Library List

## Mon II - LLI

The list of authorized programs is determined by Parmlib members IEAAPFxx or PROGxx. Link libraries can be authorized by the option LNKAUTH=LNKLST, LPA libraries are always authorized libraries.

### Field descriptions

Table 115. Fields in the LLI Report

Field Heading	Meaning
Fields that belong to each LLI Report:	
DevNum	Device number of the device on which the library is located. '???' is shown if Monitor I is not active, or volume is not mounted.
DevType	Device type, for example, 33903. '???????' is shown if Monitor I is not active, or volume is not mounted.  Shown only if option 'A' has been provided, otherwise blank.
Volser	Volume serial.  For the LPA Library List report, this value is shown only if option 'A' has been provided, otherwise blank.
Data set name	Name of the library
Fields that belong to the Link Library Report:	
APF	<b>Y</b> The link library is APF-authorized. <b>N</b> The link library is not APF-authorized.
Ext	Number of extents allocated for the link library.  Shown only if option 'A' has been provided, otherwise blank.
Fields that belong to the APF Library Report:	
Format	APF list format: <b>STATIC</b> IEALPAXx is used, or <b>STATIC</b> has been selected in PROGxx. <b>DYNAMIC</b> DYNAMIC has been selected in PROGxx.
SMS	<b>Y</b> The library is SMS-managed. <b>N</b> The library is not SMS-managed.
RACF®	<b>Y</b> The library is defined to RACF. <b>N</b> The library is not defined to RACF. <b>Sec=?</b> RACF-indication is not known, this could be an indicator that the library might not be secured correctly.  Shown only if option 'A' has been provided, otherwise blank.

## OPT - OPT Settings report

The OPT Settings report displays information about the currently active OPT member in the PARMLIB and the current settings of OPT parameters.

### How to request this report

- In ISPF, specify **L** on the Monitor II Primary Menu. This leads you to the Library List and OPT Settings Selection Menu. Here you can select:
  - 4 IEAOPTxx - OPT Settings
- In the command interface of an ISPF or TSO/E display session, type the command **OPT**.



## Contents of the report

```

RMF - OPT Settings                               Line 1 of 39
Command ==>>                                   Scroll ==> CSR

CPU= 3/ 2 UIC=1027 PR= 0                        System= T2 Total

OPT: 00          Time: N/A
-- Parameter -- - Default - -- Value -- Unit ----- Description -----
ABNORMALTERM      Yes          Yes Y/N  Abnormal terminations in routing
ABMSUCAPPING      No           No Y/N   Absolute, permanent MSU capping
BLWLINTHD         20          20 sec  Time blocked work waits for help
BLWLTRPCT         5           5 0/00  CPU cap. to promote blocked work
CCCAWMT           3200        3200 usec Alternate wait management time
CCCSIGUR          45          22 msec Min. mean-time-to-wait threshold
CNTCLIST          No           No Y/N   Clist commands count individually
CPENABLE          10,30|0,0   10,30 %  Threshold for TPI (low,high)
DVIO              Yes          Yes Y/N   Directed VIO is active
ERV               500         500/CB SU Enqueue residency CPU Service/DP
FULLPRESYSTEM     No           No Y/N   System AS can preempt other work
HIPERDISPATCH    Yes          Yes Y/N   Hiperdispatch is desired/active
IFAHONORPRIORITY  Yes          Yes Y/N   Allows CPs to help zAAPs
IIPHONORPRIORITY  Yes          Yes Y/N   Allows CPs to help zIIPs
INITIMP           0           0/FE #   INITIMP value/DP for initiators
IRA405I           70,50,50    70,50,50 % Fixed storage of <16M,16M-2G,tot
MANAGENONENCLAVE No           No Y/N   Manage non-enclave work
MAXPROMOTETIME    6           6 *10s  Holder allowed to run promoted
MCCAFCTH          400,800     2848,5696 # Threshold for storage (low,ok)
MCCFXEPR          92          92 %    Fixed storage threshold < 16 MB
MCCFXTPR          80          80 %    Fixed online storage threshold
MT_CP_MODE        1           1 #     MT CP mode
MT_ZIIP_MODE      1           1 #     MT zIIP mode
PROJECTCPU        No           Yes Y/N  CPU projection for zAAPs, zIIPs
RCCFXET           82,88       82,88 %/A Fixed<16 MPL threshold (low,high)
RCCFXTT           66,72       66,72 %/A Fixed MPL threshold (low,high)
RMPPTOM           3000        3000 msec SRM invocation interval
RTPIFACTOR        100         100 %   PI affects server routing weights
STORAGENSWDP      Yes          Yes Y/N  Sets non-swap. ASID non-dispatch.
STORAGESEVERMGT   No           No Y/N   Storage I/O priority management
STORAGEWTOR       Yes          Yes YNA  WTOR to cancel AS in shortage
SUPPSAFINFOMSG   No           No Y/N   Suppress SAF informational msg
TIMESLICES        1           1 #     Time slices for discretionary wrk
VARYCPU           No           No Y/N   VARYCPU is enabled
VARYCPUMIN        1           1 #     VARYCPUMIN value
WASROUTINGLEVEL   0           0 #     WebSphere routing level
WLMIRDSTRUC       4DIGITS     NONE type WLM IRD structure type
ZAAPAWMT          3200        3200 usec AWM time value for zAAPs
ZIIPAWMT          3200        3200 usec AWM time value for zIIPs

```

Figure 160. OPT Settings

### Field descriptions

Table 116. Fields in the OPT Settings report

Field Heading	Meaning
OPT	Suffix xx in the name of the active option member IEAOPTxx. The option member contains parameters that affect system resource manager (SRM) decisions.
Time	Timestamp when the IEAOPTxx member was activated. If the system programmer did not change the active IEAOPTxx member of SYS1.PARMLIB since the last IPL, then 'N/A' is shown.
Parameter	Name of the WLM OPT parameter.
Default	Default value(s) of the parameter. If more than one default exists, the values are separated by a vertical bar ( ).

## Mon II - OPT

Table 116. Fields in the OPT Settings report (continued)

Field Heading	Meaning
Value	Current value(s) of the parameter. This value may differ from the value originally specified. With two values displayed, separated by '/', the second value is provided by SRM. Also parameters that are not set in the IEAOPTxx member are shown with the default value, if not changed otherwise. For information on how SRM handles the settings of OPT parameters, refer to the <i>z/OS MVS Initialization and Tuning Reference</i> .  When RMF cannot obtain any data for a parameter, 'No Data' is shown.
Unit	Unit in which the parameter value is measured.
Description	Basic description of the purpose of the parameter. For detailed information refer to the <i>z/OS MVS Initialization and Tuning Reference</i> .

---

## PGSP - Page Data Set Activity report

The PGSP report provides information on page data set activity. The reporting interval is the period between any two consecutive Monitor II requests.

The PGSP report can help you to determine whether the optimum size has been allocated for each page data set.

### How to request this report

- In ISPF, specify 3 on the Resource Report Selection menu.
- In TSO/E, use PF7 to select the PGSP report.
- Command interface:

#### Display session

```
PGSP PAGE
```

#### Background session

```
PGSP (PAGE)
```

### Special considerations of report output

A Monitor I session measuring page data set activity must be active when you request the report.

Dashes (---) in the data fields indicate that RMF could not provide a value because the page device has been varied online during the reporting interval, or a Monitor I interval ended.

## Contents of the report

```

RMF - PGSP Page Data Set Activity
Command ==>>>
Line 1 of 4
Scroll ==>> PAGE

CPU= 5/ 1 UIC= 65K PR= 0 System= TRX1 Total

S VOLUME DEV DEV %SLOTS PAGE I/O REQ AVG PAGES 10:22:26
T SERIAL NUM TYPE IN USE TRAN TIME RATE PER I/O V DATA SET NAME

P TRX1PP 445D 33903 21.14 0.000 0.000 0.000 PAGE.VTRX1PP.PLPA
C TRX1PP 445D 33903 0.10 0.000 0.000 0.000 PAGE.VTRX1PP.COMMON
L TRX1P1 455D 33909 0.00 0.000 0.000 0.000 Y PAGE.VTRX1P1.LOCAL1
S N/A N/A N/A 0.36 0.000 0.000 0.000 N/A
    
```

Figure 161. PGSP Report

## Field descriptions

Table 117. Fields in the PGSP Report

Field Heading	Meaning
S T	The one-letter identifier of the type of paging space. The identifiers are: <b>P</b> PLPA <b>C</b> Common <b>L</b> Local <b>S</b> SCM (Storage Class Memory)
VOLUME SERIAL	The volume serial number of the volume on which the data set resides. N/A for SCM.
DEV NUM	The device number. N/A for SCM.
DEV TYPE	The device type. N/A for SCM.
% SLOTS IN USE	The percentage of the slots in the page data set that are in use. When you request the report, RMF calculates the percentage from the Monitor I sampling values. $\% \text{ SLOTS IN USE} = \frac{\# \text{ Slots in the Data Set} - \# \text{ Available Slots}}{\# \text{ Slots in the Data Set}} * 100$ In case of SCM, this value is the percentage of 4K SCM blocks in-use by ASM in relation to the number of total blocks available to ASM.
PAGE TRAN TIME	The page transfer time in seconds. When you request the report, RMF calculates the value from the current Monitor I sampling values. $\text{PAGE TRAN TIME} = ((\text{USE} * \text{INT}) / \text{N}) / \text{XFER}$ <b>USE</b> Number of samples when the data set is in use <b>XFER</b> Total number of pages transferred <b>N</b> Total number of samples <b>INT</b> Monitor I interval time (seconds)
I/O REQ RATE	The number of I/O requests per second for the data set made between the beginning of the interval and the time you request the report.
AVG PAGES PER I/O	The average number of pages that were transferred to or from the page data set.

This field always contains a value that reflects the activity since the last report request.

Table 117. Fields in the PGSP Report (continued)

Field Heading	Meaning
V	This field indicates whether or not the local paging data set accepts VIO pages. The symbols are: Y VIO pages are accepted N VIO pages are not accepted
DATA SET NAME	The name of the page data set being monitored. N/A for SCM.  If a page data set name is longer than 23 characters, it is truncated in the report to 22 characters followed by an asterisk (*).  If a data set has bad slots, the data set name is preceded by an asterisk (*).  When the operating system detects errors in a data set that prevent its further use, RMF can no longer monitor the data set. RMF indicates that monitoring is terminated by <i>preceding</i> the data set name with two asterisks (**).

### Report options

Due to the fact that swap data sets are no longer supported, the Report Options panel may not be used because report option SWAP leads to an empty report.

---

## SDS - Sysplex Data Server report

The report provides statistics about the usage of the SMF Data Buffer, and the usage of the Sysplex Data Services.

The numbers presented in the first section aid in finding the optimal size of the SMF buffer to hold as many SMF records as an installation might want to keep for immediate sysplex reporting. The second section contains statistics about the exploitation of the callable services. You can use these numbers to optimize the usage of the callable services within other applications than RMF.

The SDS report requires that the RMF address space has been started. Otherwise, no statistics can be provided, neither from the reporting system, nor from any remote system in the sysplex.

### How to request this report

- In ISPF, specify 8 on the Resource Report Selection menu.
- Command interface:

#### Display session

SDS

## Contents of the report

```

RMF - SDS RMF Sysplex Data Server                               Line 1 of 13
Command ==>                                                    Scroll ==> HALF

CPU= 37/ 35 UIC=2540 PR= 0                                     System= SYS1 Total

RMF Sysplex Data Server Statistics
Report Start 04/29/2016 16:49:24, End 05/06/2016 12:28:21, Duration 6d 19h

SMF Buffer Statistics, Start 04/29/2016 16:49:24, Duration 6d 19h
Buffer   Records   Record   Avg Queue   Avg Rec   Records   Buffer
Size (b) Arrived   Rate (/h) Time (ms)  Length   in Buffer  Wrap Time
    1M      23473      143       51         1K       615      04:23:18

Callable Services Statistics
Service  Requests  Request  Avg Srv   Avg Sys   Avg Amnt
Name     Arrived   Rate (/h) Time (ms) /Req     Data /Req
ERBDSQRY      0         0         0         0         0
ERBDSREC      0         0         0         0         0
ERB2XDGS      0         0         0         0         0
ERB3XDRS    3058      18        2463        3        21K
    
```

Figure 162. SDS Report

## Field descriptions

Table 118. Fields in the SDS Report

Field Heading	Meaning
RMF Sysplex Data Server Statistics	
Report Start End Duration	Start and end time of data collection, duration of data collection. The format for duration can be in days and hours (6d 19h) or in hh:mm:ss.
SMF Buffer Statistics	
Start Duration	Start time for SMF buffer statistics, duration of SMF buffer statistics. These values might have been reset during the report duration by a MODIFY command.
Buffer size	Size in bytes as specified in the SPACE subparameter of the SMFBUF parameter for the RMF address space.
Records Arrived	Number of records that arrived during the buffer statistics duration.
Record Rate (/h)	SMF data arrival rate in records per hour.
Avg Queue Time (ms)	Data server internal processing time for SMF data in milliseconds.
Avg Rec Length	Average record length for the buffer statistics duration in bytes.
Records in Buffer	Number of records currently in the SMF data buffer.
Buffer Wrap Time	Wrap-around time for the SMF data buffer (data residency time); this time is estimated before the second wrap and measured afterwards. The format can be in days and hours (6d 19h) or in hh:mm:ss.
Callable Services Statistics	
Service Name	Name of the sysplex data service module: <b>ERBDSQRY</b> RMF Query Available Sysplex SMF Data Service <b>ERBDSREC</b> RMF Request Sysplex SMF Record Data Service <b>ERB2XDGS</b> RMF Monitor II Sysplex Data Gathering Service <b>ERB3XDRS</b> RMF Monitor III Sysplex Data Retrieval Service
Requests Arrived	Number of calls to the data service.

Table 118. Fields in the SDS Report (continued)

Field Heading	Meaning
Request Rate (/h)	Data services call rate in calls per hour.
Avg Srv Time (ms)	Average response time (in milliseconds) for data services.
Avg Sys /Req	Average number of systems a request was propagated to.
Avg Amnt Data /Req	Average amount of data returned by a request in bytes.

## SENQ - System Enqueue Contention report

The report is a “snapshot” report that describes the contention or ownership at the time RMF processes the request for the report. The SENQ report tracks contention for or ownership of serially-reusable resources. SENQ reports only the contention caused by ENQ and DEQ macro instructions.

**Note:** For information about the contention caused by the RESERVE macro instruction, use the reserve activity report.

Especially when invoked during a display session, the report can help you to determine, on a real-time basis, which resources and jobs are contributing to any bottlenecks caused by resource contention.

For a detailed description of the different parameters see the *z/OS RMF User’s Guide*.

### How to request this report

- In ISPF, specify **1** on the Resource Report Selection menu.  
By default, you get the summary report, other reports can be selected via the Report Options panel.
- In TSO/E, use **PF8** to select the SENQ report.
- Command interface:

#### Display session

```

SENQ {S
      {D
      {A,sysname
      {E,sysname
      {majorname[,minorname]}
    
```

#### Background session

```

SENQ {(S
      {(D
      {(A,sysname
      {(E,sysname
      {(majorname[,minorname]})}
    
```

### Different types of SENQ reports

You can request the SENQ report as:

- Summary report for all resources that had contention
- Detail report for all resources that had contention or for a specific resource identified by name
- Report of resources held by a specific system, whether or not there is a contention

## Summary report

The summary report includes all resources that had a contention. It describes the number of tasks that own each resource and the number of tasks waiting for the resource. Figure 163 shows a sample summary report.

To request a summary report, specify **S** with the SENQ command.

## Detail report

The detail report also includes all resources for which there is contention. In addition, it identifies by jobname, system, and address space identifier the jobs that own the resource, and the jobs that are waiting for the resource. When you request the report for a specific resource by major name or major and minor name, the report includes detail data for the resource or group of resources requested. Figure 165 on page 286 shows a sample detail report.

To request a detail report, specify **D** with the SENQ command.

## Report on resources held by a specific system

The report on resources held by a specific system in a global resource serialization complex identifies either all the resources held or just the exclusively-held resources.

To request all the resources held, specify **A**, **sysname** with the SENQ command.

To request just the exclusively-held resources, specify **E**, **sysname** with the SENQ command.

This report is useful when attempting to recover an inactive system in a global resource serialization complex.

You can request this report from an active system in the complex and determine from the report the resources that the inactive system held.

Figure 164 on page 286 shows a sample resource report for a specific system.

## Contents of the report

```

Command ==>>          RMF - SENQ System Enqueue Contention          Line 1 of 4
                                                                Scroll ==>> HALF

                        CPU= 37/ 35 UIC=2540 PR=   0          System= SYS1 Total

14:52:05 TSK  TSK  TSK MAJOR NAME
          OWN  WTE  WTS  MINOR NAME

                        CLRSHARE
                        1   0   1  CLRVSAM.HBB4420.ILOG7          (SYSS)
                        SYSDSN
                        1   0   1  KEYES.MYSLM.MACLIB            (SYSS)

```

Figure 163. SENQ Summary Report

```

RMF - SENQ System Enqueue Contention
COMMAND ==>>>                                     Line 1 of 14
                                                    Scroll ==>>> HALF

CPU= 37/ 35 UIC=2540 PR=  0                      System= SYS1 Total

11:44:26
JOBNAME  SYSTEM  ASID  REQ  MAJOR NAME
                                MINOR NAME

ADMPRINT M303   185  EO   ADMPRNTQ
                                BACKGROUND DRIVER EXECUTING. (SYSS)
HSM      M303   13   SO   ARCGPA
                                ARCMCDS (SYSS)
D10MCW1  M303   116  SO   BLXDASDS
                                D83INFO.BLGPMLS (SYSS)
D32VJF1  M303   88   SO   CLR12.UN
                                IUSER (SYSS)
D31BAR1  M303   217  EO   GIMSMP
                                C87JPLX.SMPCSI.CSI (SYSS)
D31BAR1  M303   35   EO   SPFDSN
                                D31BAR1.IAREND.P06.PLS (SYSS)
D75DVB1  M303   41   EO   D75DVB1.IATOSDR.W1A.ASM (SYSS)
D96JRF1  M303   163  EO   D96JRF1.IEFDB413.PS1.PLS (SYSS)
    
```

Figure 164. SENQ GRS Report

```

RMF - SENQ System Enqueue Contention
COMMAND ==>>>                                     Line 1 of 4
                                                    Scroll ==>>> HALF

CPU= 37/ 35 UIC=2540 PR=  0                      System= SYS1 Total

11:48:44
JOBNAME  SYSTEM  ASID  REQ  MAJOR NAME
                                MINOR NAME
                                SYSDSN

D75TYT1  AQXI    136  EO   D75TYT1.IATCNNJ.P10.ASM (SYSS)
D75TYT1A AQXA     29  EW
GALER    AQXI     62  EO   GALER.PA21100.PTMPRINT (SYSS)
GALERH   AQXA     30  SW
    
```

Figure 165. SENQ Detail Report

### Field descriptions

Table 119. Fields in the SENQ Report

Field Heading	Meaning	
TSK OWN	The number of tasks that currently own the resource.	Summary report only.
TSK WTE	The number of tasks that currently are waiting for exclusive use of the resource.	
TSK WTS	The number of tasks that currently are waiting for shared use of the resource.	



Table 119. Fields in the SENQ Report (continued)

Field Heading	Meaning	
JOBNAME	The name of the job that has requested use of the resource.	Detail and specific system reports only.
SYSTEM	The identifier of the system on which the job that owns or requests the resource is running.	
ASID	The address space identifier of the job that has requested use of the resource.	
REQ	A two-character field that describes the request.  The first character indicates the type of the request: <b>E</b> The request was for exclusive use of the resource <b>S</b> The request was for shared use of the resource  The second character indicates the status of the request: <b>O</b> The requestor owns the resource <b>W</b> The requestor is waiting for the resource	
MAJORNAME MINORNAME	The name and scope of the resource. The major name, which corresponds to the <i>qname</i> field in the ENQ and DEQ macro instructions, is one to eight characters in length; it is aligned under the MAJORNAME heading. The minor name, which corresponds to the <i>rname</i> field in the ENQ and DEQ macro instructions, can be from 1 to 255 characters in length. However, only 44 characters can appear in the report. When a minor name exceeds 44 characters, it is truncated in the report, and an asterisk (*) following the scope indicates that the name has been truncated. If the minor name contains unprintable characters, RMF reports in the form 'name', where <i>name</i> appears as up to 44 hexadecimal digits. Each minor name is aligned under the MINORNAME heading. RMF recognizes only 44 characters. Therefore if two minor names (both with the same major name) are longer than 44 characters and differ only beyond the forty-fourth character, RMF cannot distinguish between them.  The scope of the resource follows the minor name. A resource with a scope of 'SYSTEMS' is followed by (SYSS). A resource with a scope of 'SYSTEM' is followed by (STEP).  When the major and minor names are blank on the report, the last non-blank values (above) are valid.	

## Report options

RMF Monitor II - System Enqueue Options

Command ==>

Specify one of the options below. To exit press END.

Enqueue contention report:

Summary       ==> YES       Specify YES for a summary or NO for a detailed report.

Enqueue report by system:

All            ==> \_\_\_       Specify YES for all owned resources or NO for exclusively owned resources only.

System ID     ==> \_\_\_\_\_ Specify the system holding the resources.

Enqueue report by major-/minorname:

Majorname    ==> \_\_\_\_\_

Minorname    ==> \_\_\_\_\_

Figure 166. SENQ Report Options Panel

You can specify either a summary report, a report by system, or a report by major-/minorname.

**Summary**

Allows you to specify a summary report.

**YES** Is the default value. A summary report includes the scope of the resource, the number of tasks waiting for exclusive use of the resource, and the number of tasks waiting for shared use of the resource.

**NO** Causes a detailed report to be generated.

**All** Allows you to specify a report that includes all resources that a system holds in a global serialization complex.

**System ID**

The system for which the enqueue report is requested.

To create a report that includes all resources owned by the system, also specify YES for All.

**Majorname, minorname**

Allows you to specify a detailed report for a specific resource that had contention. The **majorname** is a 1 to 8 character major name of a serially-reusable resource. If you specify only a major name, RMF lists all resources grouped under the major name. The optional **minorname** contains the minor name of the resource. The minor name can be 1 to 30 characters.

If you specify **S** or **D** as majorname, you need to specify a minorname.

**A** and **E** cannot be used as majornames.

---

## SENQR - System Enqueue Reserve report

The report is a 'snapshot' report that describes the status of all RESERVE requests outstanding at the time RMF processes the request for the report.

The SENQR report enables you to track RESERVE macro instructions issued to reserve a shared direct access device (shared DASD) for use by a particular system.

Figure 167 on page 289 shows a sample SENQR report requested for all devices.

### How to request this report

- In ISPF, specify **2** on the Resource Report Selection menu.  
By default, you get all volumes, a specific volume can be selected via the Report Options panel.
- In TSO/E, use **PF9** to select the SENQR report.
- Command interface:

**Display session**

```
SENQR {ALLVSER}  
      {volser }
```

**Background session**

```
SENQR {(ALLVSER)}  
      {(volser) }
```

## Contents of the report

```

RMF - SENQR System Enqueue Reserve                               Line 1 of 4
Command ==>>>                                                    Scroll ==>> HALF

                                CPU= 37/ 35 UIC=2540 PR=   0           System= SYS1 Total

14:52:57          SYSTEM ENQUEUE RESERVE REPORT
JOBNAME ASID SYSTEM  REQ VOLUME DEV  RSV MAJOR   MINOR
CATALOG  34 AQTS   S0  TS0020 0AF1 CNV SYSIGGV2 CATALOG.VTS0020
BMORRISP 70 AQTS   E0  RES84Z 0D4E ON SYSVTOC  RES84Z
JES2     31 AQTS   E0  JES2PK 0702 OFF SYSZJES2  JES2PKSYS1.BKUPCKPT
CATALOG  34 AQTS   S0  CAT212 04B1 OFF SYSZVVD5  CAT212
    
```

Figure 167. SENQR Report

### Field descriptions

Table 120. Fields in the SENQR Report

Field Heading	Meaning
JOBNAME	The name of the job that issued the RESERVE macro instruction for the device identified under DEV.
SYSTEM	The identifier of the system on which the job that owns or requests the resource is running.
ASID	The address space identifier of the job that issued the RESERVE macro instruction for the device identified under DEV.
REQ	<p>The two-character field that describes the request. The first character indicates the type of the request:</p> <p><b>E</b>      The request was for exclusive use of the device  <b>S</b>      The request was for shared use of the device</p> <p>The second character indicates the status of the request:</p> <p><b>O</b>      The requestor owns the device  <b>W</b>      The requestor is waiting for the device</p>
VOLUME	The volume serial of the volume mounted on the device identified under DEV. If reserves are issued on systems other than the one on which you request the report, the field is blank.
DEV	The address of the device for which the RESERVE macro instruction was issued. If reserves are issued on systems other than the one on which you request the report, the field is blank.
RSV	<p>The indicator of the reserve status of the device.</p> <p><b>ON</b>      The device is reserved by the processor on which RMF is running.  <b>OFF</b>      The device is being serialized via RESERVE macro instructions, but is currently not reserved.  <b>CNV</b>      The device has been converted to a GRS ENQ.</p> <p>If reserves are issued on systems other than the one on which you request the report, the field is blank.</p>
MAJOR MINOR	<p>The name used to control access to the device by means of RESERVE macro instructions. The major name, which corresponds to the <i>qname</i> field in the RESERVE macro instruction, is one to eight characters in length. It is aligned under the MAJOR heading.</p> <p>The minor name, which corresponds to the <i>rname</i> field in the RESERVE macro instruction, can be from 1 to 255 characters in length. However, only 32 characters can appear in the report. When a minor name exceeds 32 characters, it is truncated in the report, and an asterisk (*) indicates the name has been truncated. If the minor name contains unprintable characters, RMF reports it in the form 'name', where <i>name</i> is shown in hexadecimal digits and is only 29 digits in length. Each minor name is aligned under the MINOR heading. RMF recognizes only 44 characters. Therefore if two minor names (both with the same major name) are longer than 44 characters and differ only beyond the forty-fourth character, RMF cannot distinguish between them.</p>

## Report options

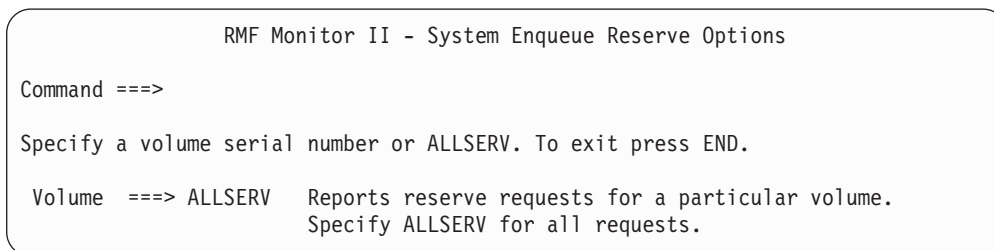


Figure 168. SENQR Report Options Panel

Enter the reserved volume name. For a report of all reserved volumes enter ALLSERV.

---

## SPAG - Paging Activity report

The Paging Activity report presents overview information on system paging activity.

This report enables you to see the paging activity of your system more clearly. This is due to the fact that most of the report fields reflect rates that show the activity since the last request for the report. Exceptions are the following fields: TIME, AFC (length of the available frame queue), HI UIC (highest UIC), and ESF AVL (number of available expanded storage frames).

Each report consists of one line of data that gives a “snapshot” view of system paging activity at the time the report was requested. When you invoke the report repetitively, you can build a table showing the differences over a period of time. Figure 169 on page 291 shows how you can repeat requests for the report to build a table of system paging activity.

### How to request this report

- In ISPF, specify 4 on the Resource Report Selection menu.
- In TSO/E, use **PF10** to select the SPAG report.
- Command interface:

#### Display session

```
SPAG
```

#### Background session

```
SPAG
```

## Contents of the report

```

RMF - SPAG Paging Activity
Command ==>>>
Line 1 of 9
Scroll ==>> HALF

CPU= 37/ 35 UIC=2540 PR= 0
System= SYS1 Total

TIME    LPA  CSA  SWP  PGS-SWPD  PRIV_IN  PRV  V&H  TAR  HI  ES  MIG  ESF  MIG
        IN  IN  OUT  IN  OUT  BLK  NBK  OUT  I+O  CWS  AFC  UIC  RTE  AGE  AVL  RTE
14:06:45 --- --- --- --- --- --- --- --- --- --- 0 251 255 ---- 5.5 735 ----
14:06:48 0.0 0.0 3.3 0.0 0.0 4.6 7.7 12 20 0 184 255 470 5.5 151 19.0
14:06:50 0.0 0.0 1.0 0.0 0.0 18 23 3.0 26 0 264 255 148 5.5 205 25.0
14:06:51 0.0 0.0 2.0 0.0 0.0 0.0 0.0 0.0 0.0 0 246 255 240 5.5 126 57.0
14:06:51 --- --- --- --- --- --- --- --- --- --- 0 234 255 ---- 5.5 93 ----
14:06:52 0.0 0.0 1.0 0.0 0.0 0.0 0.0 0.0 0.0 0 207 255 133 5.5 145 50.0
14:06:53 0.0 0.0 2.0 0.0 0.0 1.0 1.0 0.0 0.0 0 271 255 168 5.5 181 100
14:06:54 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0 277 255 5.0 5.5 184 0.0
14:06:55 0.0 0.0 2.0 0.0 0.0 0.0 0.0 0.0 0.0 0 504 255 206 5.5 251 0.0
    
```

Figure 169. SPAG Report

## Field descriptions

Table 121. Fields in the SPAG Report

Field Heading	Meaning
TIME	The time the report was requested.
LPA IN	The rate of LPA pages paged in. After the initial request, this field represents the rate since the previous report.
CSA IN	The rate of CSA pages paged in. After the initial requests, this field represents the rate since the previous report.
SWP OUT	The rate of successful swap-outs.
PGS-SWPD IN	The rate of pages swapped in. After the initial request, this field represents the rate since the previous report.
PGS-SWPD OUT	The rate of pages swapped out. After the initial request, this field represents the rate since the previous report.
PRIV_IN BLK	The rate of private area (VIO + non-VIO) pages paged in, in blocks, not including the first page. After the initial request, this field represents the rate since the previous report.
PRIV_IN NBK	The rate of private area (VIO + non-VIO) pages paged in. After the initial request, this field represents the rate since the previous report. This rate includes single pages plus the first page of each block.
PRV OUT	The rate of private area (VIO + non-VIO) pages paged out. After the initial request, this field represents the rate since the previous report.
V&H I+O	The rate of hiperspace and VIO pages paged in and paged out. After the initial request, this field represents the rate since the previous report.
TAR CWS	The target working set size for the common area.
AFC	The total number of frames currently available.
HI UIC	The highest unreferenced interval count (UIC). Values greater than 9999 are displayed as nnK to indicate a multiple of 1000. The maximum value is 65K.
ES RTE	The rate of pages sent to expanded storage. After the initial request, this field represents the rate since the previous report. In a system without expanded storage, the field heading appears in the report, but the field is blank.

## Mon II - SPAG

Table 121. Fields in the SPAG Report (continued)

Field Heading	Meaning
MIG AGE	<p>The length of time a page resides on expanded storage before it migrates to auxiliary storage. This field provides a snapshot of the migration age taken at the last sample. In a system without expanded storage, the field heading appears in the report, but the field is blank.</p> <p>If there is no unit specified for MIG AGE, the value is in migration hours. Other values are indicated as follows: <b>M</b> Migration minutes <b>S</b> Migration seconds 1.5 migration seconds are equivalent to 1 real second, this means that the displayed value has to be divided by 1.5 to get real seconds, minutes or hours.</p>
ESF AVL	The number of expanded storage frames currently available and not in use. In a system without expanded storage, the field heading appears in the report, but the field is blank.
MIG RTE	The rate of page migration from expanded storage to auxiliary storage. After the initial request, this field represents the rate since the previous report. In a system without expanded storage, the field heading appears in the report, but the field is blank.

---

## SRCS - Central Storage/Processor/SRM report

The SRCS report is a one-line summary of the current utilization of central storage, the processor, and SRM facilities. When you repeat the requests for the report, you can build a table showing the differences over a period of time. Figure 170 on page 293 is an example of repeated requests.

This report detects abnormal situations when they occur so you can request a more detailed report to further analyze the situation.

### How to request this report

- In ISPF, specify **5** on the Resource Report Selection menu.
- In TSO/E, use **PF11** to select the SRCS report.
- Command interface:

#### Display session

SRCS

#### Background session

SRCS

## Contents of the report

```

RMF - SRCS Central Storage / Processor / SRM           Line 1 of 8
Command ==>>>                                       Scroll ==>> HALF

                CPU= 14/ 6 UIC= 65K PR=  0           System= SYSF Total

      TIME    AFC    HI SQA  LPA  LPA CSA  L+C  PRI  LSQA LSQA CPU  IN  OUT  OUT  OUT
                UIC    F    F    FF    F    FF    FF    CSF  ESF  UTL   Q  LOG  RQ  WQ
16:44:56 185K  65K 0.0M 5.3K  82 5.3K 388 7687 26K      4  53  45  0  45
16:44:59 185K  65K 0.0M 5.3K  82 5.3K 388 7708 26K      4  52  46  0  46
16:44:59 185K  65K 0.0M 5.3K  82 5.3K 388 7708 26K      4  52  46  0  46
16:45:00 185K  65K 0.0M 5.3K  82 5.3K 388 7789 26K      3  52  46  0  46
16:45:00 185K  65K 0.0M 5.3K  82 5.3K 388 7769 26K      3  52  46  0  46
16:45:01 185K  65K 0.0M 5.3K  82 5.3K 388 7769 26K     14  52  46  0  46
    
```

Figure 170. SRCS Report

## Field descriptions

Table 122. Fields in the SRCS Report

Field Heading	Meaning
TIME	The time the report was requested.
AFC	The average number of available frames.
HI UIC	The highest unreferenced interval count (UIC). Values greater than 9999 are displayed as nnK to indicate a multiple of 1000. The maximum value is 65K.
SQA F	The total number of SQA frames, including frames in central storage.
LPA F	The total number of LPA frames.
LPA FF	The total number of LPA fixed frames.
CSA F	The total number of CSA frames.
L+C FF	The total number of fixed LPA and CSA frames.
PRI FF	The total number of private non-LSQA fixed frames. If the number of fixed frames is greater than 9999, asterisks (****) appear in this field.
LSQA CSF	The total number of private LSQA frames in central storage.
LSQA ESF	The total number of private LSQA frames in expanded storage. This column is blank if the system is running in z/Architecture.

If invalid data occurs, dashes appear in this field.

## Mon II - SRCS

Table 122. Fields in the SRCS Report (continued)

Field Heading	Meaning
CPU UTL	<p>The average processor utilization percentage for all general purpose processors (CPs) currently online.</p> <p>If Monitor I CPU gathering is active, it is the MVS view of CPU utilization which is the percentage of the time that the general purpose processors were busy:</p> $\text{CPU UTL} = \frac{\text{Time Range} - \text{Sum of Wait Times}}{\text{Time Range}} * 100$ <p>The time range is the sum of the times the processors were online. With HiperDispatch mode active, it is the sum of the times the processors were online but not parked.</p> <p>If Monitor I CPU gathering is not active, CPU UTL is the SRM view of CPU utilization (CCVTUTILP).  <b>Note:</b> The CPU UTL value is a snapshot of CPU usage over a short period of time and is identical to the first value in the CPU= field in the header. For details, please refer to Table 105 on page 243.</p>
IN Q	The current length of the SRM in queue.
OUT LOG	The current number of address spaces that are logically swapped out.
OUT RQ	The current length of the SRM out ready queue.
OUT WQ	The current length of the SRM out wait queue.



---

## Chapter 4. Real-time reporting with Monitor I

Monitor I produces interval reports that are created at the end of a measurement interval, for example, 30 minutes.

You can obtain Monitor I session interval reports during or at the end of RMF processing, or they can be generated at a later time by the Postprocessor.

“CHAN - Channel Path Activity report” on page 334

“CPU - CPU Activity report” on page 340

“CRYPTO - Crypto Hardware Activity report” on page 361

“DEVICE - Device Activity report” on page 367

“ENQ - Enqueue Activity report” on page 381

“FCD - FICON Director Activity report” on page 392

“IOQ - I/O Queuing Activity report” on page 403

“PAGESP - Page Data Set Activity report” on page 416

“PAGING - Paging Activity report” on page 418

“TRACE - Trace Activity report” on page 450

“VSTOR - Virtual Storage Activity report” on page 454

See Chapter 5, “Long-term overview reporting with the Postprocessor,” on page 297 for a description of these reports.



---

## Chapter 5. Long-term overview reporting with the Postprocessor

Postprocessor reports are based on data gathered as SMF records by RMF (Monitor I, Monitor II, and Monitor III), by web servers, and by Lotus® Domino® servers.

This information unit describes the following report types:

- **Interval and Duration reports**  
Reports are available as single-system and sysplex reports
- **Exception report**  
Presents a summary of values that exceeded installation-defined thresholds.
- **Overview report**  
Provides an improved version of the Exception and Summary report and offers data for further processing in spreadsheet or other applications.
- **Summary report**  
Presents an overview of system activity.

Postprocessor reports are available as either textual reports or XML reports, or both. You can use the XML output format of a report for further processing, for example, with an XML parser.

Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide*. provides or navigates to all required information on how to produce and view XML reports.

Table 123 presents an overview of available formats for all Postprocessor reports.

*Table 123. Available formats for Postprocessor reports*

Report	Text	XML produced by Postprocessor job	XML available in Spreadsheet Reporter
CACHE	yes	yes	yes
CF	yes	yes	yes
CHAN	yes	yes	yes
CPU	yes	yes	yes
CRYPTO	yes	yes	yes
DEVICE	yes	yes	yes
DOMINO	yes		
ENQ	yes	yes	yes
ESS	yes	yes	yes
FCD	yes	yes	yes
HFS	yes	yes	yes
HTTP	yes		
IOQ	yes	yes	yes
OMVS	yes	yes	yes

## Postprocessor reports

Table 123. Available formats for Postprocessor reports (continued)

Report	Text	XML produced by Postprocessor job	XML available in Spreadsheet Reporter
PAGESP	yes	yes	yes
PAGING	yes	yes	yes
PCIE		yes	yes
SCM		yes	yes
SDEVICE	yes	yes	yes
SDELAY		yes	yes
TRACE	yes		
VSTOR	yes	yes	yes
WLMGL	yes	yes	yes
XCF	yes	yes	yes
Exception report	yes		
Overview report	yes	yes	
Summary report	yes		

---

## Interval and duration reports

The Postprocessor can generate interval reports based on data gathered as SMF records by RMF (Monitor I, Monitor II, and Monitor III), by web servers, and by Lotus Domino servers.

Interval reports can be created either as **single-system reports** using the report option:

REPORTS(option)

or as **sysplex reports** with the report option:

SYSRPTS(option)

The Postprocessor can either get its input from data sets with SMF records from all systems in the sysplex, or it can access all current SMF records in the sysplex automatically using the RMF Sysplex Data Server.

For details on how to call the Postprocessor with the different options and capabilities, refer to the *z/OS RMF User's Guide*.

All Monitor I interval reports can be produced as real-time reports during the Monitor I gatherer session. The table of available reports can be found in chapter Chapter 4, "Real-time reporting with Monitor I," on page 295.

Samples of the interval reports printed during a Monitor II session appear in Chapter 3, "Snapshot reporting with Monitor II," on page 241.

In addition to interval reports, the Postprocessor can create duration reports. You can get the reports using the following command:

**Duration report:**

DINTV(hhmm)

A duration report is similar to the interval report for the same system activities. However, it summarizes activities of all the RMF measurement intervals that fall within the duration interval. The duration interval is the period of time covered in the duration report.

Duration reports allow you to measure your system's performance over long periods of time with a minimal amount of system overhead and a minimal volume of printed output.

The fields in the duration report are similar to those in the corresponding interval report. The differences are described in the sections for each report.

Table 124. Interval and Duration Reports

Report Option	Report Name	Gathered by	SMF Record
REPORTS(CACHE)	Cache subsystem activity	Monitor I	74.5
SYSRPTS(CF)	Coupling facility activity	Monitor III	74.4
REPORTS(CHAN)	Channel path activity	Monitor I	73
REPORTS(CPU)	CPU activity	Monitor I	70.1
REPORTS(CRYPTO)	Crypto hardware activity	Monitor I	70.2
REPORTS(DEVICE)	Device activity	Monitor I	74.1
REPORTS(DOMINO)	Lotus Domino server activity	Lotus Domino server	108.1, 108.3
REPORTS(ENQ)	Enqueue activity	Monitor I	77
REPORTS(ESS)	Enterprise Disk Systems activity	Monitor I	74.5, 74.8
REPORTS(FCD)	FICON director activity	Monitor I	74.7
REPORTS(HFS)	HFS statistics	Monitor III	74.6
REPORTS(HTTP)	HTTP server activity	IBM HTTP Server (IHS) powered by Domino	103.1, 103.2
REPORTS(IOQ)	I/O queuing activity	Monitor I	78.3
REPORTS(OMVS)	OMVS kernel activity	Monitor III	74.3
REPORTS(PAGESP)	Page data set activity	Monitor I	75
REPORTS(PAGING)	Paging activity	Monitor I	71
REPORTS(PCIE)	PCIE activity	Monitor III	74.9
REPORTS(SCM)	Storage Class Memory Activity	Monitor III	74.10
REPORTS(SDELAY)	Serialization delay	Monitor III	72.5
SYSRPTS(SDEVICE)	Shared device activity	Monitor I	74.1
REPORTS(TRACE)	Trace activity	Monitor I	76
REPORTS(VSTOR)	Virtual storage activity	Monitor I	78.2
SYSRPTS(WLMGL)	Workload activity	Monitor I	72.3
REPORTS(XCF)	XCF activity	Monitor III	74.2
<b>Note:</b> The ENQ, SDELAY, and TRACE report are only available as interval reports.			

## PP - Interval and duration reports

In addition, the Postprocessor can create the following interval reports based on data collected during a Monitor II background session. Duration reports are not available. You can find a description of these reports in Chapter 3, "Snapshot reporting with Monitor II," on page 241.

Table 125. Monitor II Interval Reports

Report Option	Report Name	SMF Record
REPORTS(ARD/ARDJ)	Address space resource data	79.2
REPORTS(ASD/ASDJ)	Address space state data	79.1
REPORTS(ASRM/ASRMJ)	Address space SRM data	79.5
REPORTS(CHANNEL)	Channel path activity	79.12
REPORTS(DEV/DEVV)	Device activity	79.9
REPORTS(IOQUEUE)	I/O queuing activity	79.14
REPORTS(PGSP)	Page data set activity	79.11
REPORTS(SENQ)	System enqueue contention	79.7
REPORTS(SENQR)	System enqueue reserve	79.6
REPORTS(SPAG)	Paging activity	79.4
REPORTS(SRCS)	Central storage/Processor/SRM	79.3

## Single-system report header

REPORT TITLE				PAGE nnnn
z/OS V2R2	SYSTEM ID cccc	DATE mm/dd/yyyy	INTERVAL mm.ss.ttt	
	RPT VERSION V2R2 RMF	TIME hh.mm.ss	CYCLE s.ttt SECONDS	

Figure 171. Header of a Single-System Report

All report headings contain the following information:

Field Heading	Meaning
Report title	The type of measurement data.
z/OS V2R2	The version of the operating system.
SYSTEM ID cccc	The SMF system ID of this system.
RPT VERSION V2R2 RMF	The version of the RMF Postprocessor.
DATE mm/dd/yyyy	The starting date of the measurement interval where mm is the month, dd is the day, and yyyy is the year.
INTERVAL mm.ss.ttt	The length of the measurement interval during which input is gathered for the report generators, where mm is the minutes, ss is seconds, and ttt is thousandths of seconds.
TIME hh.mm.ss	The time the interval began, where hh is hours, mm is the minutes, and ss is seconds.
CYCLE s.ttt	The length of the cycle at which data is sampled, where s is seconds and ttt is thousandths of seconds. This field appears in the heading for each report that includes sampled data.
PAGE nnnn	The page number of the report (generated by the report program), where nnnn is the page number.

All calculated numeric values in the reports are rounded to the nearest printable value, unless otherwise noted in the report descriptions. All data fields in the

reports are obtained from the corresponding SMF record image unless otherwise indicated. Those data fields that are not obtained directly are calculated from fields in the SMF record image.

## Sysplex report header

REPORT TITLE				
z/OS V2R2	SYSplex sysplex RPT VERSION V2R2 RMF	DATE mm/dd/yyyy TIME hh.mm.ss	INTERVAL mmm.ss.ttt CYCLE ss.ttt SECONDS	PAGE nnnn

Figure 172. Header of a Sysplex report

The difference in this header compared to the single-system report is that the sysplex name is shown instead of the SMF system ID.

## Duration report header

REPORT TITLE				
z/OS V2R2	SYSTEM ID cccc RPT VERSION V2R2 RMF	START mm/dd/yyyy-hh.mm.ss END mm/dd/yyyy-hh.mm.ss	INTERVAL hhh.mm.ss CYCLE s.ttt SECONDS	PAGE nnnn

Figure 173. Header of a Single-System Duration Report

In the heading area of a duration report, the START field shows when the first measurement interval within the duration interval began. The END field shows the date and time when the last interval ended.

## Overview condition names

Data displayed in most Postprocessor interval reports can be used to determine when an Exception or Overview report should be taken, as described in the *z/OS RMF User's Guide*. The overview condition names (also shortly called overview names or overview conditions) that refer to single fields in a report are shown in the rightmost column of the table of spreadsheet range names for the report, so that you can correlate them with a field heading.

For full details of the algorithms used to determine the overview condition, and the SMF record fields they apply to, see the *z/OS RMF User's Guide*.

---

## CACHE - Cache Subsystem Activity report

The Cache Subsystem Activity report provides cache statistics on a subsystem basis as well as on a detailed device-level basis.

With the help of cache control units, access time to data that resides on a DASD can be reduced to the minimum allowed by the speed of the channels, providing significant I/O response time improvements. Important questions for performance analysts are:

- *How many control units do I need?*
- *What is the optimum size of the cache?*
- *How much non-volatile storage (NVS) do I need?*
- *Which devices should or should not be cached?*
- *Is the cache performing effectively?*

The Cache Subsystem Activity report provides answers to these questions.

## How to request this report

Monitor I gathers data for this report with the default option CACHE as SMF record type 74.5. If you want to suppress gathering, you have to specify option NOCACHE.

To produce this report, specify  
REPORTS(CACHE(options))

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

### *Example URLs for the DDS API:*

```
http://ddshost:8803/gpm/rmfpp.xml?reports=CACHE(SUBSYS)
http://ddshost:8803/gpm/rmfpp.xml?reports=CACHE(SUMMARY)
http://ddshost:8803/gpm/rmfpp.xml?reports=CACHE(DEVICE)
```

## Different report levels

The contents of the report depends on the reporting options:

**Cache Summary reporting** - REPORTS(CACHE(SUMMARY)) (see "Cache Summary reporting" on page 303 )

This generates a report with three sections:

- Cache Subsystem Summary
- Top-20 Device List by DASD I/O Rate
- Top-20 Device List by total I/O Rate

The Summary report provides an overview on all subsystems with the most relevant data. The two device lists contain the 20 devices which show the highest DASD I/O rates (devices that should be investigated for potential cache hit improvements) or that show the highest I/O rates (hopefully with the best cache hit rates).

**Subsystem-level reporting** - REPORTS(CACHE(SUBSYS)) (see "Subsystem-level reporting" on page 304 )

This generates a report with these sections:

- Cache Subsystem Status
- Cache Subsystem Overview
- Cache Subsystem Device Overview
- RAID Rank Activity

The subsystem-level report gives an overall view of the storage controller, that is the amount of cache storage and non-volatile storage installed, as well as the current status of the cache. In addition, the performance analyst finds the number of I/O requests sent to the control unit and their resolution in the cache (*hits*). Furthermore, a list of all volumes attached to the subsystem is part of the report, showing their specific utilization of the cache.

The suboptions SSID/EXSSID can be used to select or exclude specific control units.

**Device-level reporting** - REPORTS(CACHE(DEVICE)) (see "Device-level reporting" on page 311 )



This generates, in addition to the report previously described, a report with two sections:

- Cache Device Status
- Cache Device Activity

The device-level report provides detailed information for each single device attached to the selected control unit. The status section shows whether caching and DASD FAST WRITE are active, or whether the current device is part of a duplex pair. The report is intended to help analyze cache usage in detail on the basis of the information about the applications that access these volumes.

**Note:** When comparing I/O rates in the DASD Activity report and in the Cache Subsystem Activity report, you may see differences due to different ways how I/Os are counted:

- In the DASD Activity report, one I/O is counted for one SSCH or RSCH instruction. There can be record chaining, for example for paging I/O, which is not reflected in the SSCH count.
- In the Cache Subsystem Activity report, one I/O is counted for each cache request, and one I/O chain may cause several cache requests.

This different I/O counting can lead to higher or lower I/O rates in the Cache Subsystem Activity report than in the DASD Activity report.

## Cache Summary reporting

### Cache Subsystem Summary

The report offers you a top-down approach to analyze the storage subsystems in your configuration because you can see at a glance the most important data. Looking at this report, the storage subsystems causing problems can be easily identified and analyzed in a second Postprocessor run requesting more details.

C A C H E S U B S Y S T E M S U M M A R Y																	PAGE	1	
z/OS V2R2		SYSTEM ID SYS1		DATE 09/28/2016		INTERVAL 15.00.055													
		RPT VERSION V2R2 RMF		TIME 00.30.00															
SSID	CU-ID	TYPE	CACHE	NVS	I/O RATE	OFF RATE	--CACHE READ	HIT DFW	RATE- CFW	-----DASD STAGE	I/O RATE- DFWBP	ICL	BYP	OTHER	ASYNC RATE	TOTAL H/R	READ H/R	WRITE H/R	% READ
3000	1B30	2105-E20	3072	192	67.5	0.0	23.8	38.8	0.0	4.9	0.0	0.0	0.0	0.0	20.4	0.928	0.830	1.000	42.5
3001	1B87	2105-E20	3072	192	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	N/A	N/A	N/A	N/A
3002	1C1E	2105-E20	3072	192	27.2	0.0	23.3	0.0	0.0	3.9	0.0	0.0	0.0	0.0	0.0	0.857	0.857	N/A	100.0
3003	1C95	2105-E20	3072	192	7.6	0.0	6.8	0.0	0.0	0.8	0.0	0.0	0.0	0.0	0.0	0.891	0.891	N/A	100.0
3004	1D0C	2105-E20	3072	192	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	N/A	N/A	N/A	N/A

Figure 174. Cache Subsystem Summary Report - Subsystem Summary

### Top-20 device lists

In addition to the subsystem summary, the report consists of two top-20 lists of devices, sorted in descending order by DASD I/O rate and by total I/O rate. These two lists allow you to identify the volumes with the highest I/O rates to the lower interface of a subsystem as well as the volumes with the highest I/O rates in total. Solving a possible problem, one of the listed devices would probably be of most benefit to the overall subsystem.

# PP - CACHE

C A C H E S U B S Y S T E M S U M M A R Y														PAGE 2			
z/OS V2R2			SYSTEM ID OS04				DATE 06/05/2016			INTERVAL 14.59.996							
			RPT VERSION V2R2 RMF				TIME 09.30.00										
*** DEVICE LIST BY DASD I/O RATE ***																	
VOLUME SERIAL	DEV NUM	SSID	% I/O	I/O RATE	---CACHE READ	HIT DFW	RATE-- CFW	-----DASD STAGE	I/O DFWBP	RATE-- ICL	----- BYP	OTHER	ASYNC RATE	TOTAL H/R	READ H/R	WRITE H/R	% READ
PRD440	077E	00B1	25.7	53.1	21.3	16.7	0.0	13.9	0.0	1.3	0.0	0.0	1.6	0.714	0.604	1.000	67.9
PPDS14	0220	00CC	6.9	29.1	1.0	15.7	0.0	0.1	0.0	12.4	0.0	0.0	2.0	0.573	0.943	1.000	6.2
PRD437	0214	00CC	4.5	19.0	12.5	2.0	0.0	4.4	0.0	0.0	0.0	0.0	2.0	0.765	0.738	0.998	89.2
PPD026	0876	00F1	27.4	71.4	65.9	1.1	0.0	4.3	0.0	0.0	0.0	0.0	0.5	0.939	0.939	0.975	98.4
PRD554	06B9	00FE	59.1	32.4	28.9	0.0	0.0	3.5	0.0	0.0	0.0	0.0	0.0	0.893	0.893	1.000	99.9
PRD339	0231	00CC	2.4	10.3	7.2	0.0	0.0	3.1	0.0	0.0	0.0	0.0	0.0	0.698	0.698	N/A	0.0
...																	
*** DEVICE LIST BY TOTAL I/O RATE ***																	
VOLUME SERIAL	DEV NUM	SSID	% I/O	I/O RATE	---CACHE READ	HIT DFW	RATE-- CFW	-----DASD STAGE	I/O DFWBP	RATE-- ICL	----- BYP	OTHER	ASYNC RATE	TOTAL H/R	READ H/R	WRITE H/R	% READ
PPD026	0876	00F1	27.4	71.4	65.9	1.1	0.0	4.3	0.0	0.0	0.0	0.0	0.5	0.939	0.939	0.975	98.4
PRD440	077E	00B1	25.7	53.1	21.3	16.7	0.0	13.9	0.0	1.3	0.0	0.0	1.6	0.714	0.604	1.000	67.9
PRD327	0200	00CC	11.8	49.8	3.0	46.8	0.0	0.1	0.0	0.0	0.0	0.0	6.3	0.998	0.973	1.000	6.2
PRD343	0515	00E4	17.5	48.9	24.8	24.0	0.0	0.1	0.0	0.0	0.0	0.0	5.5	0.998	0.996	0.999	50.8
PBV321	022C	00CC	11.2	47.3	47.1	0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.1	1.000	1.000	1.000	99.7
PRD307	0507	00E4	15.3	42.8	4.1	38.2	0.0	0.5	0.0	0.0	0.0	0.0	13.0	0.988	0.891	1.000	10.8
...																	

Figure 175. Cache Subsystem Activity Report - Top-20 Device Lists

Table 126. Fields in the Cache Subsystem Summary Report

Field Heading	Meaning
Most of the fields in this report are identical to fields in the Cache Subsystem Activity report. Therefore, please refer to Table 127 on page 305 and Table 131 on page 311.	
<b>Note:</b> The reported storage capacities for cache and non-volatile storage (NVS) represents only the Cluster Processor Complex in the storage server, that controls the subsystem. Since a typical storage server has two clusters, you must double the reported capacities to get the actual sizes.	
CACHE	Amount of physical storage that is configured in this storage subsystem (in megabytes unless otherwise noted).
NVS	Amount of physical non-volatile storage (NVS) that is configured in this storage subsystem (in megabytes unless otherwise noted).
OFF RATE	Rate of I/O requests to non-cached devices during the reporting interval.
% I/O	Percentage of I/O requests to this volume, compared to the total number of I/O requests sent to the subsystem it is attached to. This is not the percentage of all I/O requests in the system.

## Subsystem-level reporting

### Cache subsystem status and overview

The first section shows the configuration of the selected storage subsystem. This includes the amount of cache configured (installed) and available for use, and how much storage, if any, is pinned because of a DASD failure. It also includes the amount of non-volatile storage needed for the DASD FAST WRITE function. Finally, the overall subsystem status is shown in terms of whether caching is active, or, for example, CACHE FAST WRITE is activated.

The second section provides details of the cache usage of the subsystem. Here, the analyst can see at a glance all the I/O requests to the subsystem, divided into the categories NORMAL, SEQUENTIAL, and CACHE FAST WRITE. READ and WRITE requests are shown separately as totals, rates, or ratios.

**Note:** All values shown as RATE are calculated on the basis of seconds that are shown in the CINT value in the report header.

C A C H E S U B S Y S T E M A C T I V I T Y														PAGE	1
z/OS V2R2		SYSTEM ID SYS1		DATE 09/28/2016		INTERVAL 15.00.000									
		RPT VERSION V2R2 RMF		TIME 00.30.00											
SUBSYSTEM	2105-01	CU-ID	1E91	SSID	3007	CDATE	09/28/2016	CTIME	00.29.55	CINT	15.00				
TYPE-MODEL	2105-E20	MANUF	IBM	PLANT	75	SERIAL	000000016374								
-----															
C A C H E S U B S Y S T E M S T A T U S															
-----															
SUBSYSTEM STORAGE		NON-VOLATILE STORAGE		STATUS											
CONFIGURED	3072.0M	CONFIGURED	192.0M	CACHING	-	ACTIVE									
AVAILABLE	2692.2M	PINNED	0.0	NON-VOLATILE STORAGE	-	ACTIVE									
PINNED	0.0			CACHE FAST WRITE	-	ACTIVE									
OFFLINE	0.0			IML DEVICE AVAILABLE	-	YES									
-----															
C A C H E S U B S Y S T E M O V E R V I E W															
-----															
TOTAL I/O	185787	CACHE I/O	185787	CACHE OFFLINE	0										
TOTAL H/R	0.893	CACHE H/R	0.893												
CACHE I/O	-----READ I/O REQUESTS-----					-----WRITE I/O REQUESTS-----					%				
REQUESTS	COUNT	RATE	HITS	RATE	H/R	COUNT	RATE	FAST	RATE	HITS	RATE	H/R	READ		
NORMAL	151012	167.0	132096	146.1	0.875	5339	5.9	5339	5.9	5339	5.9	1.000	96.6		
SEQUENTIAL	24588	27.2	23696	26.2	0.964	4848	5.4	4848	5.4	4848	5.4	1.000	83.5		
CFW DATA	0	0.0	0	0.0	N/A	0	0.0	0	0.0	0	0.0	N/A	N/A		
TOTAL	175600	194.2	155792	172.3	0.887	10187	11.3	10187	11.3	10187	11.3	1.000	94.5		
-----	---CACHE MISSES---					-----MISC-----			-----NON-CACHE I/O-----						
REQUESTS	READ	RATE	WRITE	RATE	TRACKS	RATE	COUNT	RATE	ICL	COUNT	RATE				
NORMAL	18916	20.9	0	0.0	25086	27.8	DFW BYPASS	0	0.0	ICL	0	0.0			
SEQUENTIAL	892	1.0	0	0.0	36075	39.9	CFW BYPASS	0	0.0	BYPASS	0	0.0			
CFW DATA	0	0.0	0	0.0			DFW INHIBIT	0	0.0	TOTAL	0	0.0			
TOTAL	19808	RATE	21.9												
-----	---CKD STATISTICS---		---RECORD CACHING---		---HOST ADAPTER ACTIVITY---				-----DISK ACTIVITY-----						
WRITE	0	READ MISSES	0	BYTES		BYTES	RESP	BYTES	BYTES						
WRITE HITS	0	WRITE PROM	472	/REQ		/SEC	TIME	BYTES	BYTES						
				982		1.1K	READ	16.912	65.9K	12.7M					
				15.9K		2.8M	WRITE	12.465	48.3K	4.1M					

Figure 176. Cache Subsystem Activity Report - Status and Overview

Table 127. Fields in the Cache Subsystem Activity Report - Header

Field Heading	Meaning
SUBSYSTEM	Storage subsystem type (as configured).
CU-ID	Physical control unit number of the caching subsystem. This is equal the lowest device number, or to the device that has been turned online first, respectively.
SSID	Subsystem identifier: a number assigned during installation of the subsystem that uniquely identifies the storage subsystem.
<p><b>Note:</b> Device reserve activity can cause a data gatherer interface to wait until a reserve has been released. This in turn can cause the cache interval to be much longer than a regular RMF interval.</p> <p>Therefore, CDATE, CTIME and CINT have been introduced to show the actual point in time to which the cache interval start is related, and the actual cache interval length. All rates shown in the report are based on CINT, not on INTERVAL.</p>	
CDATE	Date when the cache interval started.
CTIME	Time when the cache interval started.
CINT	Cache interval time.  In interval reports, the format is <i>mm.ss</i> , while in duration reports the format is <i>hh.mm.ss</i> .
TYPE-MODEL	Device type and model.
TYPE MODEL MANUF PLANT SERIAL	The hardware description of the disk system.
VOLSER	Volume serial number of the reported DASD device (only for device-level reporting).
NUM	Device number of the reported DASD device (only for device-level reporting).

## PP - CACHE

Table 128. Fields in the Cache Subsystem Activity Report - Subsystem Status

Field Heading	Meaning
<p><b>Note:</b> The reported storage capacities for cache and non-volatile storage (NVS) represents only the Cluster Processor Complex in the storage server, that controls the subsystem. Since a typical storage server has two clusters, you must double the reported capacities to get the actual sizes.</p>	
SUBSYSTEM STORAGE	Physical capacity of random access cache (in megabytes).
CONFIGURED	Amount of storage that is installed in this storage subsystem.
AVAILABLE	Amount of storage that is available for caching. This is the total cache size minus the amount used by the subsystem for the cache directory, minus the amount pinned and offline storage.
PINNED	Amount of storage that is unavailable because a DASD failure is preventing the subsystem from destaging the data. The data is pinned in cache.
OFFLINE	Amount of storage that is offline because of a host or subsystem error.
NON-VOLATILE STORAGE (NVS)	Physical capacity of random access storage with a backup battery power source (in megabytes).
CONFIGURED	Amount of NVS that is installed in this storage subsystem.
PINNED	Amount of NVS that is unavailable because a DASD failure is preventing the subsystem from destaging the data. The data is pinned in NVS.
STATUS	The caching status of the entire subsystem.
CACHING	<p>Overall caching status of the subsystem.</p> <p><b>ACTIVE</b> Caching is active (online and usable).</p> <p><b>DISABLED FOR MAINTENANCE</b> Cache has been disabled for maintenance.</p> <p><b>PENDING ACTIVE STATE</b> Caching is pending active, that is, cache is in the process of being brought online.</p> <p><b>INTERNAL ERROR TERMINATION</b> An internal error stopped cache (cache is offline)</p> <p><b>EXPLICIT HOST TERMINATION</b> Cache has been deactivated by request from host system or support facility.</p> <p><b>DEACTIVATION IN PROCESS</b> A command requesting deactivation of cache has been received, and destaging from cache to DASD is still in progress.</p> <p><b>DEACTIVATION FAILED</b> A command requesting deactivation of cache has been received, but destaging to DASD has failed.</p>

Table 128. Fields in the Cache Subsystem Activity Report - Subsystem Status (continued)

Field Heading	Meaning
NON-VOLATILE STORAGE	<p>Overall status of the non-volatile storage (NVS).</p> <p><b>ACTIVE</b> NVS is online and usable.</p> <p><b>PENDING DUE TO ERROR</b> A command requesting deactivation of NVS has been received but transfer from NVS to DASD has failed.</p> <p><b>DEACTIVATION IN PROCESS</b> A command requesting deactivation of NVS has been received, and destaging to DASD is still in progress.</p> <p><b>DISABLED FOR MAINTENANCE</b> NVS has been disabled for maintenance by the support facility.</p> <p><b>INTERNAL ERROR TERMINATION</b> An internal error caused termination of NVS.</p> <p><b>EXPLICIT HOST TERMINATION</b> NVS has been deactivated by request from host system or support facility.</p> <p><b>DASD FAST WRITE INHIBITED</b> DASD FAST WRITE is inhibited because the battery is defective.</p>
CACHE FAST WRITE	<p>Status of the CACHE FAST WRITE (CFW) option.</p> <p><b>ACTIVE</b> CFW is active.</p> <p><b>DEACTIVATED</b> CFW is deactivated. <b>Note:</b> CFW does not use NVS.</p>
IML DEVICE AVAILABLE	<p>Status of the IML device.</p> <p><b>NO</b> Device containing a diskette drive for loading the microcode is not available.</p> <p><b>YES</b> Device is operational.</p>

Table 129. Fields in the Cache Subsystem Activity Report - Subsystem Overview

Field Heading	Meaning
TOTAL I/O	Total number of I/O requests to cached devices in the storage subsystem.
TOTAL H/R	Ratio of I/Os that were processed within the cache (cache hits) based on the total number of I/Os.
CACHE I/O	<p>Total number of cacheable I/O requests to cached devices in the storage subsystem.</p> <p>This value excludes INHIBIT CACHE LOAD and CACHE BYPASS I/O requests.</p>
CACHE H/R	Ratio of I/Os that were processed within the cache (cache hits) based on the total number of cacheable I/O requests.
CACHE OFFLINE	Total number of I/O requests to non-cached devices in the storage subsystem.
CACHE I/O REQUEST - The channel command DEFINE EXTENT specifies the way the cache will be used. There are three categories (NORMAL, SEQUENTIAL, CFW DATA) and a TOTAL value:	
NORMAL	Cache will be managed by <i>least-recently-used (LRU)</i> algorithm for making cache space available.
SEQUENTIAL	Tracks following the track assigned in the current CCW chain are promoted. They will be transferred from DASD to cache in anticipation of a short-term requirement.
CFW DATA	WRITE and READ-AFTER-WRITE requests are processed in cache. The data might not be written to DASD. Because CFW does not use the NVS, the application is responsible for restoring the data after a cache or subsystem failure.
TOTAL	This is either the sum of I/O requests, the total I/O rate, or the average hit ratio for the three categories previously described.

## PP - CACHE

Table 129. Fields in the Cache Subsystem Activity Report - Subsystem Overview (continued)

Field Heading	Meaning
READ I/O REQUESTS	Cache I/O requests that searched or read data from DASD. This is the number of channel operations that had at least one SEARCH or READ command but no WRITE commands. This is counted for cached devices only.
COUNT	Total number of SEARCH/READ requests.
HITS	Number of SEARCH/READ requests that completed without accessing the DASD.
H/R	Hit Ratio - number of SEARCH/READ hits compared to the total number of SEARCH/READ requests.
WRITE I/O REQUESTS	Cache I/O requests that wrote data to DASD. This is the number of channel commands that had at least one WRITE command. It is counted for cached devices only.
COUNT	Total number of WRITE requests.
FAST	Total number of DASD/CACHE FAST WRITE requests.
HITS	Number of DASD/CACHE FAST WRITE requests that completed without accessing the DASD (fast write hit).
H/R	Number of DASD/CACHE FAST WRITE hits compared to the sum of all READ and WRITE requests (excluding ICL and BYPASS).
%READ	Percentage of READ requests compared to the sum of all READ and WRITE requests (excluding ICL and BYPASS).
CACHE MISSES	Cache misses are calculated as the difference between total I/O requests and the number of cache hits. They are shown for normal, sequential, and CFW requests.  A cache miss occurs if a record that is to be read from or written to a data set is not found in the cache. In either case, the track in which this record is located is searched on the DASD and transferred to cache. Subsequent requests referring to the same record can then be processed in cache only, provided that the data has not yet been moved out because cache space is needed for other data.  Other reasons for cache misses are DFW bypass and DFW inhibit. In the first case, a lack of NVS caused the data to be written to DASD immediately, while in the second case, usage of DFW is inhibited (for example, the device does not allow DFW).
READ	Number of SEARCH/READ requests that needed access to DASD because the data could not be found in the cache.
WRITE	Number of WRITE requests that needed access to DASD because the data could not be found in the cache.
TRACKS	Number of tracks transferred from DASD to cache.
TOTAL	Total number of I/O requests that needed access to DASD because the data could not be found in the cache.
MISC - Miscellaneous cache activities	
DFW BYPASS	Number of DASD FAST WRITE requests that would have resulted in a DFW hit; however, NVS was overutilized causing writes to be sent directly to DASD.  This value is also known as DFW RETRY.
CFW BYPASS	Number of operations that did not transfer a track from DASD into cache because no free segments were available. If no free segments are available in cache, there is no destaging in favor of I/O requests with the CACHE FAST WRITE attribute. The I/O goes directly to the DASD.
DFW INHIBIT	If DASD FAST WRITE is active, this is the number of WRITE requests that <i>inhibited</i> DASD FAST WRITE. If DASD FAST WRITE is inactive, this is the number of WRITE requests that directly accessed the DASD, even with DASD FAST WRITE turned on.
ASYNC	Number of tracks transferred from cache to DASD asynchronously to transfer from the channel to release space in the cache and the NVS.  For a duplex pair, this is the count of transfers from the cache to the secondary device of a duplex pair. Transfers from the cache to the primary device of a duplex pair are not counted.  A high number of ASYNC I/Os is an indicator for an over-committed cache or NVS.
NON-CACHE I/O - READ	requests that switched off cache processing.

Table 129. Fields in the Cache Subsystem Activity Report - Subsystem Overview (continued)

Field Heading	Meaning
ICL	Inhibit cache load. Number of I/O requests that inhibited load of data into cache although the data was not found in the cache. <b>Note:</b> If the data had been in the cache, it would have been counted as cache hit. Therefore, this is actually the number of ICL misses.
BYPASS	Number of I/O requests that explicitly bypassed the cache, irrespective of whether the data is in the cache or not.
TOTAL	Total number of I/O requests that bypassed the cache.
CKD STATISTICS - CKD (Count-Key-Data) is a format used to store data on DASD. The counts shown in this section are contained in the total write count.	
WRITES	Number of write I/O requests in CKD format.
WRITE HITS	Number of write I/O requests in CKD format that could be resolved in the cache.
RECORD CACHING - Record caching is done dynamically upon a decision made by DCME or the microcode. It may improve overall cache performance if caching of whole tracks would waste cache storage. The decision is based on the number of I/Os, the hit ratio, and the locality of reference of a certain entity of data.	
READ MISSES	Number of instances in which a record requested for READ was not found in the cache, and access to DASD was required.
WRITE PROM	Number of instances in which a record requested for WRITE was found in the cache, and access to DASD was not required.
HOST ADAPTER ACTIVITY <sup>1)</sup> - I/O activity of normal, sequential and CFW read and write requests.	
BYTES/REQ	The average number of transferred bytes per read and write request.
BYTES/SEC	The average number of transferred bytes per second for read and write requests.
DISK ACTIVITY <sup>1)</sup> - Transfer activity from hard disk to cache and vice versa.	
RESP TIME	Response time in milliseconds per read and write request.
BYTES/REQ	The average number of transferred bytes per read and write request.
BYTES/SEC	The average number of transferred bytes per second for read and write requests.
<sup>1)</sup> available for the IBM TotalStorage DS family	

**Duration reports:** An asterisk behind a status field or the field VOLUME SERIAL indicates that the contents of the field has changed during the duration interval.

### Cache subsystem device overview and RAID rank activity

The first section lists all the devices in the subsystem. Each line shows the most important statistics for the device it represents. The I/O rate is divided into two groups (cache hits and DASD I/O), showing the different types of I/O activity in each group.

A RAID rank is a set of physical volumes. Several logical volumes as well as parallel access volumes are associated with a single RAID rank. Such a subsystem consists of multiple RAID ranks. If several higher utilized logical volumes are mapped to the same RAID rank, DASD skew is likely to appear. Knowing which logical volumes are associated to a certain RAID rank allows the storage administrator to move logical volumes from one RAID rank to another and thus optimally balance the load on the RAID ranks.

The RAID Rank Activity section in this report only appears for each 2105 subsystem. For 2107 subsystems, RAID rank activity is measured in the ESS Statistics report. The RAID Rank Activity section provides information about each RAID rank belonging to the subsystem. It shows the physical characteristics of a RAID rank, the details for READ and WRITE requests, and lists the volumes with the highest activity.



# PP - CACHE

C A C H E S U B S Y S T E M A C T I V I T Y														PAGE 2			
z/OS V2R2				SYSTEM ID SYS1			DATE 09/28/2016			INTERVAL 15.00.000							
				RPT VERSION V2R2 RMF			TIME 00.30.00										
SUBSYSTEM	2105-01	CU-ID	1E91	SSID	3007	CDATE	09/28/2016	CTIME	00.29.55	CINT	15.00						
TYPE-MODEL	2105-E20	MANUF	IBM	PLANT	75	SERIAL	00000016374										
-----																	
C A C H E S U B S Y S T E M D E V I C E O V E R V I E W																	
-----																	
VOLUME SERIAL	DEV NUM	RRID	% I/O	I/O RATE	---CACHE READ	HIT DFW	RATE-- CFW	-----DASD STAGE	I/O DFWBP	-----RATE ICL	-----BYP	-----OTHER	ASYNC RATE	TOTAL H/R	READ H/R	WRITE H/R	% READ
*ALL			100.0	205.5	172.3	11.3	0.0	21.9	0.0	0.0	0.0	0.0	20.5	0.893	0.887	1.000	94.5
*CACHE-OFF			0.0	0.0													
*CACHE			100.0	205.5	172.3	11.3	0.0	21.9	0.0	0.0	0.0	0.0	20.5	0.893	0.887	1.000	94.5
NP1MHD	1E80	0700	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	N/A	N/A	N/A	N/A
NP1MHE	1E81	0700	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	N/A	N/A	N/A	N/A
NP1MHF	1E82	0700	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	N/A	N/A	N/A	N/A
NP1MHG	1E83	0700	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	N/A	N/A	N/A	N/A
NP1MHH	1E84	0700	8.1	16.5	14.7	0.0	0.0	1.8	0.0	0.0	0.0	0.0	0.0	0.891	0.891	N/A	100.0
NP1MHI	1E85	0700	17.3	35.6	27.9	0.0	0.0	7.6	0.0	0.0	0.0	0.0	0.0	0.785	0.785	N/A	100.0
NP1MHJ	1E86	0700	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	N/A	N/A	N/A	N/A
NP1MHK	1E87	0700	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	N/A	N/A	N/A	N/A
NP1MHL	1E88	0700	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	N/A	N/A	N/A	N/A
NP1MHM	1E89	0700	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	N/A	N/A	N/A	N/A
...																	
-----																	
R A I D R A N K A C T I V I T Y																	
-----																	
ID	RAID TYPE	DA	HDD	----- READ RATE	REQ AVG MB	----- MB/S	RTIME	----- WRITE RATE	REQ AVG MB	----- MB/S	RTIME	----- HIGHEST UTILIZED VOLUMES -----					
*ALL			7	68	0.054	3.7	16	10	0.105	1.0	32						
0700	RAID-5	19	7	68	0.054	3.7	16	10	0.105	1.0	32	NP1MJ2	NP1MJK	NP1MHI	NP1MI4	NP1MHN	NP1MHH

Figure 177. Cache Subsystem Activity Report - Device Overview

Table 130. Fields in the Cache Subsystem Activity Report - Device Overview

Field Heading	Meaning
VOLUME SERIAL	*ALL All volumes belonging to the reported storage subsystem. *CACHE-OFF All non-cached volumes. *CACHE All cached volumes.  For all other lines: Volume serial number.
DEV NUM	Device number.
RRID or XTNT POOL	RRID is displayed for 2105 subsystems and is the RAID rank identifier.  XTNT POOL is displayed for 2107 subsystems and is the extent pool identifier.
% I/O	Percentage of I/O requests to this volume or category, compared to the total number of I/O requests sent to the subsystem.
I/O RATE	Number of I/O requests per second during the reporting interval.
CACHE HIT RATE - I/O rate of all cache hits.	
READ	Rate of SEARCH/READ requests that completed without accessing the DASD.
DFW	Rate of DFW requests.
CFW	Rate of WRITE and READ-AFTER-WRITE requests that are processed in cache.
DASD I/O RATE - I/O rate of all requests that accessed DASD.	
STAGE	Rate of normal or sequential I/O requests that accessed DASD.
DFWBP	Rate of requests that caused DFW BYPASS.
ICL	Rate of Inhibit cache load requests.



Table 130. Fields in the Cache Subsystem Activity Report - Device Overview (continued)

Field Heading	Meaning
BYP	Rate of requests that explicitly bypassed the cache, irrespective of whether the data was in the cache or not.
OTHER	Rate of CFW BYPASS and DFW INHIBIT requests.
ASYNC RATE	Number of I/Os per second that caused asynchronous transfer from cache to DASD (destaging).
TOTAL H/R	Ratio of I/O requests that were processed within the cache (cache hits) to the total number of I/Os.
READ H/R	Number of READ request hits compared to all READ requests (excluding ICL and BYPASS).
WRITE H/R	Number of WRITE request hits compared to all WRITE requests (excluding ICL and BYPASS).
% READ	Percentage of READ requests compared to all READ and WRITE requests.

Table 131. Fields in the Cache Subsystem Activity Report - RAID Rank Activity

Field Heading	Meaning
ID	RAID rank ID. *ALL is shown for the summary of all RAID ranks.
RAID TYPE	RAID rank type.
DA	Device adapter ID.
HDD	Number of hard disk drives in the RAID rank.
<b>READ and WRITE Requests</b>	
RATE	Number of I/O requests per second.
AVG MB	Average number of megabytes transferred per I/O request.
MB/S	Bandwidth of the I/O requests.
RTIME	Average response time of an HDD I/O request (milliseconds). These HDD requests could be I/Os due to READ misses, sequential prestages, destages, and copy services.
<b>Volume Utilization</b>	
HIGHEST UTILIZED VOLUMES	Top six volumes with I/O activity sorted in descending order.

## Device-level reporting

### Cache device status

This section of the report presents the cache status of the device and the duplex pair.

### Cache device activity

The second section of the report contains the same type of information as the Cache Subsystem Overview section, but for only one device. The field contents are explained in Table 129 on page 307.

PP - CACHE

C A C H E D E V I C E A C T I V I T Y														PAGE	2
z/OS V2R2		SYSTEM ID SYS1		DATE 09/28/2016		INTERVAL 15.00.055									
		RPT VERSION V2R2 RMF		TIME 00.30.00											
SUBSYSTEM	2105-01	CU-ID	1E91	SSID	3007	CDATE	09/28/2016	CTIME	00.29.55	CINT	15.00				
TYPE-MODEL	2105-E20	MANUF	IBM	PLANT	75	SERIAL	000000016374								
VOLSER	NP1MHI	NUM	1E85	RRID	0700										
-----															
C A C H E D E V I C E S T A T U S															
-----															
C A C H E S T A T U S				D U P L E X P A I R S T A T U S											
CACHING	- ACTIVE			DUPLEX PAIR	- NOT ESTABLISHED										
DASD FAST WRITE	- ACTIVE			STATUS	- N/A										
PINNED DATA	- NONE			DUAL COPY VOLUME	- N/A										
-----															
C A C H E D E V I C E A C T I V I T Y															
-----															
TOTAL I/O	32141	CACHE I/O	32141	CACHE OFFLINE	N/A										
TOTAL H/R	0.785	CACHE H/R	0.785												
-----															
CACHE I/O	-----READ I/O REQUESTS-----				-----WRITE I/O REQUESTS-----								%		
REQUESTS	COUNT	RATE	HITS	RATE	H/R	COUNT	RATE	FAST	RATE	HITS	RATE	H/R	READ		
NORMAL	26728	29.6	19998	22.1	0.748	0	0.0	0	0.0	0	0.0	N/A	100.0		
SEQUENTIAL	5413	6.0	5241	5.8	0.968	0	0.0	0	0.0	0	0.0	N/A	100.0		
CFW DATA	0	0.0	0	0.0	N/A	0	0.0	0	0.0	0	0.0	N/A	N/A		
TOTAL	32141	35.6	25239	27.9	0.785	0	0.0	0	0.0	0	0.0	N/A	100.0		
-----															
-----CACHE MISSES-----				-----MISC-----				-----NON-CACHE I/O-----							
REQUESTS	READ	RATE	WRITE	RATE	TRACKS	RATE	COUNT	RATE					COUNT	RATE	
NORMAL	6730	7.4	0	0.0	6899	7.6	DFW BYPASS	0	0.0	ICL				0	0.0
SEQUENTIAL	172	0.2	0	0.0	6026	6.7	CFW BYPASS	0	0.0	BYPASS				0	0.0
CFW DATA	0	0.0	0	0.0			DFW INHIBIT	0	0.0	TOTAL				0	0.0
							ASYNC (TRKS)	0	0.0						
TOTAL	6902	RATE	7.6												
-----															
---CKD STATISTICS---		---RECORD CACHING---			---HOST ADAPTER ACTIVITY---				-----DISK ACTIVITY-----						
WRITE	0	READ MISSES	0	BYTES BYTES				RESP BYTES BYTES							
WRITE HITS	0	WRITE PROM	28	/REQ /SEC				TIME /REQ /SEC							
				READ 736 146				READ 14.132 63.1K 743.9K							
				WRITE 15.9K 471.6K				WRITE 13.812 50.9K 648.1K							

Figure 178. Cache Subsystem Activity Report - Cache Device Activity (device-level reporting)

**Note:** In the header of the Cache Device Activity Report, RRID shown for 2105 subsystems is the RAID rank identifier. EXTENT POOL shown for 2107 subsystems is the extent pool identifier.

Table 132. Fields in the Cache Subsystem Activity Report - Cache Device Status

Field Heading	Meaning
CACHE STATUS	Status of the cache from the perspective of the device being reported.
CACHING	<p>Caching status of the device.</p> <p><b>ACTIVE</b> Caching is active; requests to the reported device can be processed without DASD access.</p> <p><b>DEACTIVATE PENDING</b> Cache has been deactivated on request from host system or support facility, but transfer of modified data to DASD has failed.</p> <p><b>DEACTIVATED</b> Caching has been deactivated for the reported device.</p>

Table 132. Fields in the Cache Subsystem Activity Report - Cache Device Status (continued)

Field Heading	Meaning
DASD FAST WRITE	<p>Status of the DASD FAST WRITE option.</p> <p><b>ACTIVE</b> DASD FAST WRITE requests can be processed for this device.</p> <p><b>DEACTIVATION PENDING</b> DASD FAST WRITE has been terminated on request by host system or support facility, but transfer of modified data to DASD is in progress or has failed.</p> <p><b>DEACTIVATED</b> DASD FAST WRITE requests are ignored for this device.</p>
PINNED DATA	<p>A device has failed, and data that has not yet been written to DASD is pinned in cache or NVS for later recovery.</p> <p><b>NONE</b> No data is pinned for the reported device.</p> <p><b>EXISTS, DFW NOT SUSPENDED</b> Pinned data exists, but DASD FAST WRITE has not been suspended.</p> <p><b>EXISTS, DFW SUSPENDED</b> Pinned data exists, and DASD FAST WRITE has been suspended.</p>
<p>DUPLEX PAIR STATUS - A duplex pair of devices can be established, to ensure a maximum of data security. During normal processing, all I/O requests are made on behalf of the primary device. Data is written to the primary device, either directly or via cache. In either case, the control unit ensures that the data is copied to the secondary device. This happens asynchronously to the transfer from the channel.</p>	
DUPLEX PAIR	<p>Overall dual copy status.</p> <p><b>NOT ESTABLISHED</b> The device is configured as a simplex device.</p> <p><b>ACTIVE</b> The device reported is either the primary or the secondary device of a duplex pair.</p> <p><b>PENDING</b> The control unit copies data to establish the duplex pair to which the reported device belongs.</p> <p><b>SUSPENDED</b> An error in one device of the duplex pair caused the duplex pair to be suspended.</p>
STATUS	<p>The status of this device as part of a duplex pair.</p> <p><b>PRIMARY</b> The device is the primary device of the duplex pair. All channel operations are associated with this device.</p> <p><b>SECONDARY</b> The device is the secondary device of the duplex pair. No regular I/O is possible to this device. I/O operations to the primary device are duplexed to this device by the control unit.</p> <p><b>N/A</b> The device has not been established as part of a duplex pair.</p>
DUAL COPY VOLUME	<p>Identification of the other device of a duplex pair.</p> <p><b>nnnn</b> Device address of the other device of the duplex pair, if a duplex pair is established.</p> <p><b>N/A</b> The device has not been established as part of a duplex pair.</p>

The report contains less information for devices which had the cache offline at the end of the interval.

# PP - CACHE

C A C H E   S U B S Y S T E M   A C T I V I T Y											
z/OS V2R2		SYSTEM ID		SYS1	DATE	09/28/2016	INTERVAL	12.00.000			
		RPT VERSION		V2R2	RMF	TIME	07.12.00				
SUBSYSTEM	3990-06	CU-ID	02A0	SSID	0016	CDATE	09/28/2016	CTIME	07.11.55	CINT	12.00
TYPE-MODEL	9396-001	MANUF	IBM	PLANT	92	SERIAL	000000012345				
VOLSER	RZ3381	NUM	02B6	RRID	N/A						
----- C A C H E   D E V I C E   S T A T U S -----											
CACHE STATUS					DUPLEX PAIR STATUS						
CACHING	- DEACTIVATED				DUPLEX PAIR NOT ESTABLISHED						
DASD FAST WRITE	- DEACTIVATED										
PINNED DATA	- NONE										
----- C A C H E   D E V I C E   A C T I V I T Y -----											
TOTAL I/O	---	CACHE I/O	---	CACHE OFFLINE	8772						
TOTAL H/R	---	CACHE H/R	---								

Figure 179. Cache Subsystem Activity Report - Cache Offline

## Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following table shows the overview condition names for the Overview report.

Table 133. Overview names in the Cache Subsystem Activity Report

Field Heading or Meaning	Subsystem Report	Device Report
Subsystem Status / Device Status.		
CACHING	CASSC	CADSC
NON-VOLATILE STORAGE	CASSNVS	
Subsystem Overview / Device Activity		
TOTAL I/O	CASTOT	CADTOT
CACHE I/O	CASCTOT	CADCTOT
CACHE OFFLINE	CASCOFF	
TOTAL H/R	CASHRT	CADHRT
CACHE H/R	CASHR	CADHR
READ I/O REQUESTS RATE NORMAL	CASRN	CADRN
READ I/O REQUESTS RATE SEQUENTIAL	CASRS	CADRS
READ I/O REQUESTS RATE CFW DATA	CASRC	CADRC
READ I/O REQUESTS RATE TOTAL	CASRT	CADRT
READ I/O REQUESTS HITS RATE NORMAL	CASRHN	CADRHN
READ I/O REQUESTS HITS RATE SEQUENTIAL	CASRHS	CADRHS
READ I/O REQUESTS HITS RATE CFW DATA	CASRHC	CADRHC
READ I/O REQUESTS HITS RATE TOTAL	CASRHT	CADRHT
READ I/O REQUESTS H/R NORMAL	CASRHRN	CADRHRN
READ I/O REQUESTS H/R SEQUENTIAL	CASRHRS	CADRHRS
READ I/O REQUESTS H/R CFW DATA	CASRHRC	CADRHRC

Table 133. Overview names in the Cache Subsystem Activity Report (continued)

Field Heading or Meaning	Subsystem Report	Device Report
READ I/O REQUESTS H/R TOTAL	CASRHRT	CADRHRT
WRITE I/O REQUESTS RATE NORMAL	CASWN	CADWN
WRITE I/O REQUESTS RATE SEQUENTIAL	CASWS	CADWS
WRITE I/O REQUESTS RATE CFW DATA	CASWC	CADWC
WRITE I/O REQUESTS RATE TOTAL	CASWT	CADWT
WRITE I/O REQUESTS FAST WRITE RATE NORMAL	CASWFN	CADWFN
WRITE I/O REQUESTS FAST WRITE RATE SEQUENTIAL	CASWFS	CADWFS
WRITE I/O REQUESTS FAST WRITE RATE CFW DATA	CASWFC	CADWFC
WRITE I/O REQUESTS FAST WRITE RATE TOTAL	CASWFT	CADWFT
WRITE I/O REQUESTS HITS RATE NORMAL	CASWHN	CADWHN
WRITE I/O REQUESTS HITS RATE SEQUENTIAL	CASWHS	CADWHS
WRITE I/O REQUESTS HITS RATE CFW DATA	CASWHC	CADWHC
WRITE I/O REQUESTS HITS RATE TOTAL	CASWHT	CADWHT
WRITE I/O REQUESTS H/R NORMAL	CASWHRN	CADWHRN
WRITE I/O REQUESTS H/R SEQUENTIAL	CASWHRS	CADWHRS
WRITE I/O REQUESTS H/R CFW DATA	CASWHRC	CADWHRC
WRITE I/O REQUESTS H/R TOTAL	CASWHRT	CADWHRT
% READ NORMAL	CASRWN	CADRWN
% READ SEQUENTIAL	CASRWS	CADRWS
% READ CFW DATA	CASRWC	CADRWC
% READ TOTAL	CASRWT	CADRWT
CACHE MISSES READ RATE NORMAL	CASMRN	CADMRN
CACHE MISSES READ RATE SEQUENTIAL	CASMRS	CADMRS
CACHE MISSES READ RATE CFW DATA	CASMRC	CADMRC
CACHE MISSES WRITE RATE NORMAL	CASMWN	CADMWN
CACHE MISSES WRITE RATE SEQUENTIAL	CASMWS	CADMWS
CACHE MISSES WRITE RATE CFW DATA	CASMWC	CADMWC
CACHE MISSES TRACKS RATE NORMAL	CASMTN	CADMTN
CACHE MISSES TRACKS RATE SEQUENTIAL	CASMTS	CADMTS
CACHE MISSES RATE TOTAL	CASMT	CADMT
MISC (Miscellaneous) DFW BYPASS RATE	CASDFWB	CADDFWB
MISC (Miscellaneous) CFW BYPASS RATE	CASCFWB	CADCFWB
MISC (Miscellaneous) DFW INHIBIT RATE	CASDFWI	CADDFWI
MISC (Miscellaneous) ASYNC(TRKS) RATE	CASASYNC	CADASYNC
NON CACHE I/O ICL RATE	CASNCICL	CADNCICL
NON CACHE I/O BYPASS RATE	CASNCB	CADNCB
NON CACHE I/O TOTAL RATE	CASNCT	CADNCT
HOST ADAPTER ACTIVITY BYTES/REQ READ	CASBRR	CADBRR

Table 133. Overview names in the Cache Subsystem Activity Report (continued)

Field Heading or Meaning	Subsystem Report	Device Report
HOST ADAPTER ACTIVITY BYTES/SEC READ	CASBRS	CADBRS
HOST ADAPTER ACTIVITY BYTES/REQ WRITE	CASBWR	CADBWR
HOST ADAPTER ACTIVITY BYTES/SEC WRITE	CASBWS	CADBWS
DISK ACTIVITY RESP TIME READ	CASDRRT	CADDRRT
DISK ACTIVITY BYTES/REQ READ	CASDRBR	CADDRBR
DISK ACTIVITY BYTES/SEC READ	CASDRBS	CADDRBS
DISK ACTIVITY RESP TIME WRITE	CASDWRT	CADDWRT
DISK ACTIVITY BYTES/REQ WRITE	CASDWBR	CADDWBR
DISK ACTIVITY BYTES/SEC WRITE	CASDWBS	CADDWBS
Subsystem Device Overview		
I/O RATE (volser)	CADT	
I/O RATE (*ALL)	CASAT	
I/O RATE (*CACHE)	CASCT	
I/O RATE (*CACHE-OFF)	CASOT	
DASD I/O RATE STAGE (volser)	CADSTG	
DASD I/O RATE STAGE (*ALL)	CASASTG	
DASD I/O RATE STAGE (*CACHE)	CASCSTG	
% I/O (*CACHE-OFF)	CASCOIO	
RAID RANK READ REQ RATE	CARRRT	
RAID RANK READ REQ AVG MB	CARRMB	
RAID RANK READ REQ MB/S	CARRMBS	
RAID RANK READ REQ RTIME	CARRRTIM	
RAID RANK WRITE REQ RATE	CARWRT	
RAID RANK WRITE REQ AVG MB	CARWMB	
RAID RANK WRITE REQ MB/S	CARWMBS	
RAID RANK WRITE REQ RTIME	CARWRTIM	

## CF - Coupling Facility Activity report

A Coupling Facility Activity report is produced for each coupling facility attached to the sysplex. It provides the following information:

- Coupling Facility usage summary
- Coupling Facility structure activity
- Subchannel activity
- CF to CF activity

### How to request this report

Monitor III gathers data for this report automatically. If you want to suppress gathering, you have to disable writing SMF record type 74.4.

To produce this report, specify  
SYSRPTS(CF)

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

**Example URL for the DDS API:**

<http://ddshost:8803/gpm/rmfpp.xml?reports=CF>

## Contents of the report

A Coupling Facility Activity report is produced for each coupling facility attached to the sysplex. Figure 180 on page 318 gives an example of the overall structure of the Coupling Facility Activity report. It shows the sequencing of the report sections:

- Coupling Facility Usage Summary
- Coupling Facility Structure Activity
- Subchannel Activity
- CF to CF Activity

For a complete Coupling Facility Activity report, it is recommended to combine data from all of the systems in the sysplex. If data from one or more systems is missing, the Structure and Subchannel Activity sections of the report are incomplete. In addition, the PRIM (primary) and SEC (secondary) indicators of synchronously duplexed structures might be missing in the Usage Summary section because this information is gathered only on one member of the sysplex (sysplex master gathering).

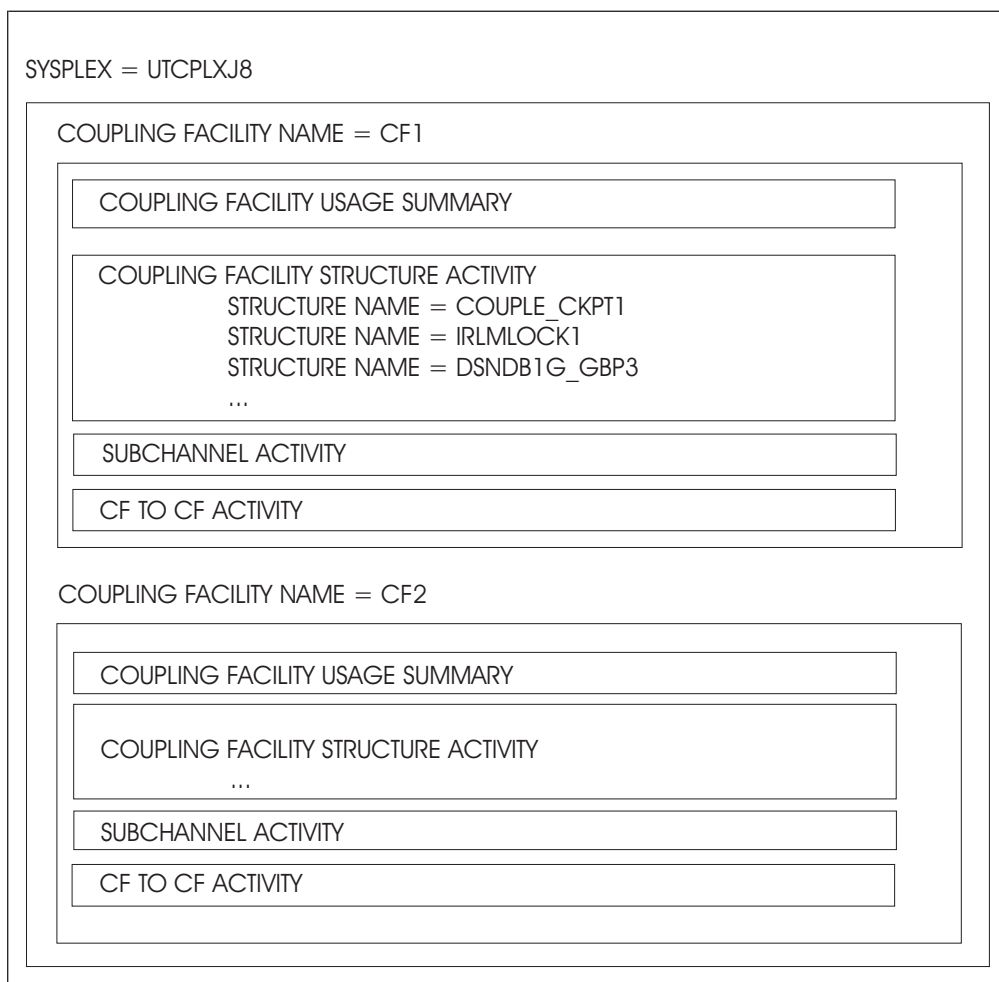


Figure 180. Structure of the Coupling Facility Activity report

### Coupling Facility Usage Summary section

This section of the Coupling Facility Activity report gives a snapshot of coupling facility storage and specific structure data at the end of the reporting interval. It lists all the structures occupying space either in the coupling facility real storage or in storage class memory (SCM), including those with no currently active connections. For asynchronously duplexed secondary structures, it provides a summary on consumed processing times and performance counts related to asynchronous duplexing. Data shown for asynchronous duplex sync up requests are accumulated over all reported systems that have a connector to the structure.

All structure summary data is grouped in the report by structure type (LIST, LOCK, CACHE and UNKN in that order). Within the structure type, the report lines are in alphanumeric order by structure name. This ordering puts the structure data in the same relative position across interval reports.



COUPLING FACILITY ACTIVITY												PAGE 1
z/OS V2R2		SYSPLEX UTCPLXJ8 RPT VERSION V2R2 RMF		DATE 09/28/2016 TIME 12.00.00		INTERVAL 030.00.000 CYCLE 01.000 SECONDS						
-----												
COUPLING FACILITY NAME = X7CFP87												
TOTAL SAMPLES(AVG) = 1800 (MAX) = 1800 (MIN) = 1799												
-----												
COUPLING FACILITY USAGE SUMMARY												
-----												
GENERAL STRUCTURE SUMMARY												
-----												
TYPE	STRUCTURE NAME	STATUS CHG	ALLOC SIZE	% OF CF STOR	# REQ	% OF ALL REQ	% OF CF UTIL	AVG REQ/ SEC	LST/DIR ENTRIES TOT/CUR	DATA ELEMENTS TOT/CUR	LOCK ENTRIES TOT/CUR	DIR REC XI'S
LIST	DBSVPLX7_SCA	ACTIVE PRIM	5M	0.0	10754	1.5	1.8	5.97	3484 182	6743 622	N/A N/A	N/A N/A
	DFHXQLS_G2POOL1	ACTIVE	63M	0.4	0	0.0	0.0	0.00	96K 1822	96K 1859	N/A N/A	N/A N/A
LIST	ISTGENERIC	ACTIVE PRIM	11M	0.0	20850	0.0	0.0	11.58	25K 47	499 2	4 0	N/A N/A
	THRLSTSCMKP1_1	ACTIVE	4G	4.1	20114K	4.4	16.5	11174	1428K 1262K	8569K 7060K	1024 0	N/A N/A
LOCK	DBTPLX5_ALOCK	ACTIVE SEC A	44M	0.3	68212	9.4	9.2	37.90	11K 0	0 0	17M 0	N/A N/A
...												
CACHE	IRRXCF00_B002	ACTIVE	2M	0.0	0	0.0	0.0	0.00	184 4	181 1	N/A N/A	0 0
-----												
SCM STRUCTURE SUMMARY												
-----												
TYPE	STRUCTURE NAME	SCM SPACE ALG %USED	SPACE MAX/ %USED	AUGMENTED EST. MAX/ %USED	LST ENTRY EST. MAX/ CUR	LST ELEM EST. MAX/ CUR	--- SCM READ CNT/BYTE X'FERRED	--- SCM WRITE CNT/BYTE X'FERRED	--- SCM READ AVG ST/ STD_DEV	--- SCM WRITE AVG ST/ STD_DEV	SCM AUX ENABLED CMD/%ALL	DELAYED FAULTS CNT/%ALL
LIST	THRLSTSCMKP1_1	KP1	16384M 0.0	754M 0.3	9585K 0	57508K 0	4790 5023M	3944 4136M	1353.8 832.0	1762.3 736.6	0 0.0	58749 0.0
-----												
ASYNCHRONOUS CF DUPLEXING SUMMARY												
-----												
TYPE	STRUCTURE NAME	TOTAL	ASYNC DUPLEX --TRANSMIT TIME-- AVG STD_DEV	CF OPERATIONS --SERVICE TIME-- AVG STD_DEV	---	ASYNC DUPLEX TOTAL #SUSPEND	SYNC UP REQUESTS --SUSPEND TIME-- AVG STD_DEV	---	---	---	---	---
LOCK	DBTPLX5_ALOCK	19432K	1.1	22.6	1.2	23.8	43197	1	744.0	0.0		
	DBTSTPLX_LCK	6700K	1.1	22.2	1.1	23.8	7420	0	0.0	0.0		
-----												
STORAGE SUMMARY												
-----												
			ALLOC SIZE	% OF CF STORAGE		----- DUMP SPACE -----	% IN USE	MAX % REQUESTED				
	TOTAL CF STORAGE USED BY STRUCTURES		6202M	6.2								
	TOTAL CF DUMP STORAGE		1024M	1.0	0.0	0.0						
	TOTAL CF AUGMENTED SPACE		2M	0.0								
	TOTAL CF STORAGE AVAILABLE		93138M	92.8								
	TOTAL CF STORAGE SIZE		100366M									
			ALLOC SIZE	% ALLOCATED								
	TOTAL CONTROL STORAGE DEFINED		100366M	5.7								
	TOTAL DATA STORAGE DEFINED		0K	0.0								
			ASSIGNED	% IN USE	SUM MAX SCM							
	TOTAL CF STORAGE CLASS MEMORY		16384M	0.0	16384M							
-----												
PROCESSOR SUMMARY												
-----												
COUPLING FACILITY		2827	MODEL H66	CFLEVEL 19	DYNDISP OFF							
AVERAGE CF UTILIZATION (% BUSY)		53.9	LOGICAL PROCESSORS:		DEFINED 2	EFFECTIVE 2.0						
					SHARED 0	AVG WEIGHT 0.0						

Figure 181. Coupling Facility Activity Report - Usage Summary

The report summarizes request activity for active structures. This activity is described in more detail in the Coupling Facility Structure Activity section of the report. It is included in this report to give the customer a quick view of the relative amount of activity among the structures in a coupling facility.

The following table explains the field headings in the Coupling Facility Usage Summary section.

Table 134. Fields in the Coupling Facility Activity Report - Usage Summary

Field Heading	Meaning
TOTAL SAMPLES (AVG), (MAX), (MIN)	Average (AVG), the maximum (MAX) and minimum (MIN) sample count for all systems connected to this coupling facility.
GENERAL STRUCTURE SUMMARY	
TYPE	Indicates whether the structure is a list, lock, or cache structure. The structures being reported are grouped by structure type.  UNKN indicates a structure for which there was no activity during the interval but that is still allocated in the coupling facility. There are no structure activity details for UNKN structures.
STRUCTURE NAME	The name given to the structure by the coupling facility policy specification in the Function Couple Data Set. It is up to 16 characters and is unique within a sysplex.
STATUS	Indicates status of the structure at the end of the interval:  <b>ACTIVE</b> At least one system is connected to the structure. If a structure became active during this interval, the report gives the partial interval activity data. In the unlikely event a structure becomes active several times during an interval, only the last activation is reported.  <b>ACTIVE PRIM</b> The structure is the rebuilt-old (primary) structure in a duplexing rebuild process. An appended "A" indicates that the structure is asynchronously duplexed.  <b>ACTIVE SEC</b> The structure is the rebuilt-new (secondary) structure in a duplexing rebuild process. An appended "A" indicates that the structure is asynchronously duplexed.  <b>INACTV</b> No system is connected to the structure but it still occupies storage in the coupling facility. The structure will not show any request activity because RMF was unable to gather end-of-interval data for calculating delta values. A structure is inactive while it is undergoing recovery operations or being moved to another coupling facility, or it was specified by the owning subsystem as a persistent structure. There are no structure activity details for an inactive structure.  <b>UNALLOC</b> No system is connected to the structure and it no longer occupies storage in the coupling facility. The structure was active earlier in the interval but no activity data is shown because RMF was unable to collect end-of-interval data for calculating delta values. There are no structure activity details reported for an unallocated structure.  The PRIM and SEC indicators of synchronously duplexed structures might not appear if data from one or more systems in the sysplex is missing.
CHG	X indicates that the status of this structure changed during the reporting interval.
ALLOC SIZE	The number of bytes set aside in the coupling facility for this structure by the coupling facility policy in the Function Couple Data Set. Storage is allocated in increments of 4K bytes. This storage consists of both control and data storage.
% OF CF STOR	The percentage of the total coupling facility storage allocated to this structure.

Table 134. Fields in the Coupling Facility Activity Report - Usage Summary (continued)

Field Heading	Meaning
# REQ	The number of requests processed by the coupling facility against this structure. This is the same number as appears in the TOTAL line of the Coupling Facility Structure Activity report.
% OF ALL REQ	The percentage of all requests attributable to this structure. Use this field for a quick idea of where the activity occurred during the interval.
% OF CF UTIL	The percentage of CF processor time used by the structure. The structure execution time is related to the total CF-wide processor busy time. The sum of the values in this column is less than 100%, because not all CF processor time is attributable to structures.  N/A is shown in this field if the CF level is lower than 15.
AVG REQ/SEC	The average number of requests per second for this structure.
LST/DIR ENTRIES	<b>TOT</b> Maximum number of list or directory entries that can reside in coupling facility real storage for the structure.  <b>CUR</b> Number of structure list or directory entries which are currently in use and reside in coupling facility real storage.  N/A in this and the following fields indicates that the information is not applicable.
DATA ELEMENTS	<b>TOT</b> Maximum number of list elements that can reside in coupling facility real storage.  <b>CUR</b> Number of structure list elements which are currently in use and reside in coupling facility real storage.
LOCK ENTRIES	<b>TOT</b> The total number of lock table entries.  <b>CUR</b> The non-zero lock table count found.
DIR REC/ DIR REC XI'S	Number of Cache directory reclaims.  Directory reclaims occur when the total number of used unique entities exceeds the total number of directories. Whenever this shortage of directory entries occurs, the coupling facility will reclaim in-use directory entries associated with unchanged data. All users of that data must be notified that their copy of the data is invalid. As a consequence, it may happen that this data must be re-read from DASD and registered to the coupling facility again.  Directory reclaim activity can be avoided by increasing the directory entries for a particular structure.  The second value is the number of reclaims that caused an XI (see XI field in the Structure Activity section). A high value is an indicator for a performance problem in this structure.
SCM STRUCTURE SUMMARY <sup>(See note 2)</sup>	
ALG	Type of algorithm that is used by the coupling facility to control the movement of structure objects between coupling facility real storage and storage class memory:  <b>KP1</b> KeyPriority1 <b>UNK</b> Unknown
SCM SPACE	<b>MAX</b> Maximum amount of storage class memory that this structure can use (in bytes).  <b>%USED</b> Percentage of maximum amount of storage class memory that is in use by this structure.
AUGMENTED	<b>EST.MAX</b> Estimated maximum amount of CF space that may be assigned as augmented space for this structure (in bytes).  <b>%USED</b> Percentage of maximum augmented space that is in use by this structure.

**PP - CF**

Table 134. Fields in the Coupling Facility Activity Report - Usage Summary (continued)

Field Heading	Meaning
LST ENTRY	<p><b>EST.MAX</b> Estimated maximum number of list entries that may reside in storage class memory for this structure.</p> <p><b>CUR</b> Number of existing structure list entries that reside in storage class memory.</p>
LST ELEM	<p><b>EST.MAX</b> Estimated maximum number of list elements that may reside in storage class memory for this structure.</p> <p><b>CUR</b> Number of existing structure list elements that reside in storage class memory.</p>
SCM READ	<p><b>CNT</b> The number of read operations against storage class memory that were either initiated</p> <ul style="list-style-type: none"> <li>• by a reference to list structure objects residing in storage class memory, or</li> <li>• as a prefetch operation in order to retrieve list structure objects in storage class memory that are expected to be referenced.</li> </ul> <p><b>BYTE X'FERRED</b> SCM read bytes transferred. This is the number of bytes transferred from storage class memory to CF.</p> <p><b>AVG ST</b> Average service time per SCM read operation to storage class memory in microseconds.</p> <p><b>STD_DEV</b> Standard deviation of the service time for SCM read operations to storage class memory in microseconds.</p>
SCM WRITE	<p><b>CNT</b> The number of list write operations performed to storage class memory.</p> <p><b>BYTE X'FERRED</b> SCM write bytes transferred. This is the number of bytes transferred from CF storage to storage class memory.</p> <p><b>AVG ST</b> Average service time per SCM write operation to storage class memory in microseconds.</p> <p><b>STD_DEV</b> Standard deviation of the service time for SCM write operations to storage class memory in microseconds.</p>
SCM AUX ENABLED	<p><b>CMD</b> SCM auxiliary enabled command count. This is the number of commands that required the use of CF auxiliary frames.</p> <p><b>%ALL</b> Percentage of the SCM auxiliary enabled command count in relation to all requests for this structure.</p>
DELAYED FAULTS	<p><b>CNT</b> Number of commands and for multiple list-entry commands, the number of list item references that were delayed due to a fault condition resulting in a required access to storage class memory.</p> <p><b>%ALL</b> Percentage of delayed faults in relation to all requests for the structure.</p>
<p>1 ASYNCHRONOUS CF DUPLEXING SUMMARY<sup>(See note 3)</sup></p>	

Table 134. Fields in the Coupling Facility Activity Report - Usage Summary (continued)

Field Heading	Meaning
ASYNC DUPLEX CF OPERATIONS	<p><b>TOTAL</b> Number of asynchronous duplex operations transmitted from the primary to the secondary structure that completed in the secondary structure.</p> <p><b>TRANSMIT TIME AVG</b> Average asynchronous duplex operation transmission time for operations sent from the primary to the secondary structure, in microseconds.</p> <p><b>TRANSMIT TIME STD_DEV</b> Standard deviation of the average asynchronous duplex operation transmission time.</p> <p><b>SERVICE TIME AVG</b> Average service time to transfer the asynchronous duplex operation to the secondary structure and complete the operation in the secondary structure, in microseconds.</p> <p><b>SERVICE TIME STD_DEV</b> Standard deviation of the average service time to transfer and complete the asynchronous duplex operation in the secondary.</p>
ASYNC DUPLEX SYNC_UP REQUESTS	<p><b>TOTAL</b> Total number of user requests that requested sync up with the primary structure. User requests are summed up over all systems connected to the structure.</p> <p><b>#SUSPEND</b> Number of user requests that were suspended waiting for asynchronous duplex operations to complete in the secondary structure. Suspend counts are summed up over all systems connected to the structure.</p> <p><b>SUSPEND TIME AVG</b> Average accumulated suspend time for suspended requests waiting for asynchronous duplex operations to complete in the secondary structure, in microseconds.</p> <p><b>SUSPEND TIME STD_DEV</b> Standard deviation of the average accumulated suspend time.</p>
STORAGE SUMMARY	
TOTAL CF STORAGE USED BY STRUCTURES	<p>The total amount of coupling facility storage that is used by structures and the percentage of the total coupling facility storage allocated to these structures.</p> <p>These totals do not necessarily represent 100% of the facility activity for the interval. There is some amount of storage and request activity overhead that is not attributable to individual structures. For example, the total for # REQ will usually be less than the sum of the TOTAL REQ from the Subchannel Activity Report because the subchannel numbers include facility management command counts whereas the structure numbers do not.</p>
TOTAL CF DUMP STORAGE	Amount and percentage of coupling facility space allocated as dump space.
DUMP SPACE	<p><b>% IN USE</b> The percentage of dump space in use at end of the interval. This amount is a sampled value so it is intended to show trends, not instantaneous peaks.</p> <p><b>MAX % REQUESTED</b> The maximum percentage of dump space requested since the coupling facility dump storage was allocated.</p> <p>This high water mark is maintained by the coupling facility hardware so is reset only when dump space is reinitialized. If this percentage is over 100, it means at least one dump has been lost or truncated since the most recent allocation of dump space. If the percentage is close to or over 100, you should increase the dump space allocation by modifying the coupling facility policy for dump space and activating the modified policy.</p>

**PP - CF**

Table 134. Fields in the Coupling Facility Activity Report - Usage Summary (continued)

Field Heading	Meaning
TOTAL CF AUGMENTED SPACE	<p><b>ALLOC SIZE</b> Total amount of CF storage used by all structures as augmented space (in bytes).</p> <p><b>% OF CF STORAGE</b> Percentage of CF storage used by all structures as augmented space.</p>
TOTAL CF STORAGE AVAILABLE	The amount and percentage of coupling facility space that is not allocated to any structure, not allocated as dump space, and not allocated as augmented space.
TOTAL CF STORAGE SIZE	The total amount of storage in the coupling facility, including both allocated and available space. This value does not include the storage required by the coupling facility code itself, so that it differs from the storage assigned to the coupling facility on the HMC.
TOTAL CONTROL STORAGE DEFINED, TOTAL DATA STORAGE DEFINED, % ALLOCATED	<p>The amount of coupling facility storage that is allowed to be occupied by control information (CONTROL STORAGE) or data (DATA STORAGE).</p> <p>For each structure, plus the dump area, a certain amount of control and data storage is allocated. The coupling facility defines an area called control storage; structure control information is restricted to that area. The remaining storage is called data storage and is used for structure data. If the data storage area becomes full, structure data can then be allocated from the control storage area. If TOTAL DATA STORAGE DEFINED is zero, it means control information can reside anywhere on the coupling facility and there are no allocation restrictions.</p> <p>If the % ALLOCATED field for control storage shows a percentage approaching 100, it means the control storage is close to being completely allocated even though the CF SPACE AVAILABLE field may still show an amount of total free space. Possible customer actions include:</p> <ul style="list-style-type: none"> <li>• Changing structure preference lists in the coupling facility policy specification to direct some structures away from this facility.</li> <li>• Adding another coupling facility to the sysplex.</li> </ul>
TOTAL CF STORAGE CLASS MEMORY	<p><b>ASSIGNED</b> Total CF storage class memory. This is the amount of storage class memory that may be concurrently used as structure extensions. Storage is assigned in increments of 4K bytes.</p> <p><b>% IN USE</b> Percentage of storage class memory that is in use by all structures of the coupling facility.</p> <p><b>SUM MAX SCM</b> Sum of the storage class memory maxima defined for all structures of the coupling facility.</p>
PROCESSOR SUMMARY	
COUPLING FACILITY	Coupling facility processor type.
MODEL	Coupling facility processor model.
CFLEVEL	Coupling facility architected function level.
DYNDISP <sup>(See note 1)</sup>	The dynamic CF dispatching status (ON, OFF, or THIN). THIN indicates that coupling thin interrupts are enabled for the coupling facility (only for CFLEVEL 19 or higher).

Table 134. Fields in the Coupling Facility Activity Report - Usage Summary (continued)

Field Heading	Meaning
AVERAGE CF UTILIZATION (% BUSY)	<p>Average value of CPU utilizations within the coupling facility.</p> <p>The utilization of the individual CPs in the coupling facility is recorded in the SMF 74, Subtype 4, Processor Data Section.</p> <p>In case of a stand-alone coupling facility, the utilization of the individual CPs should be approximately the same. In a PR/SM environment where this CP is shared with other partitions the utilization is the logical utilization of the CP (that is, only the utilization by the coupling facility). The CPU Activity report can be used to determine the total utilization of the CP.</p> <p>If the average utilization is high, you can take the following actions:</p> <ol style="list-style-type: none"> <li>1. In a PR/SM environment, you can dedicate the CP to the integrated coupling facility or assign additional CPs to the partition.</li> <li>2. Move structures to a coupling facility with lower utilization.</li> <li>3. Consider additional or larger coupling facilities.</li> </ol>
LOGICAL PROCESSORS DEFINED <sup>(See note 1)</sup>	Number of logical processors defined for the coupling facility.
LOGICAL PROCESSORS EFFECTIVE	<p>Number of effective available logical processors in a shared environment. This value is only useful in CFCC environment. CFCC measures the time of real command execution as well as the time waiting for work. The reported value shows the ratio of the LPAR dispatch time (CFCC execute and wait time) to the RMF interval length.</p> <p>For example, if a CFCC CEC contains 6 LPs, and the measured CF LPAR has two logical processors and is limited at 5 % the number of effective LPs is 0.3</p> <p>Please, refer to the CPU Activity report in case of an ICMF LPAR.</p>
LOGICAL PROCESSORS SHARED	The number of shared processors defined for the coupling facility.
LOGICAL PROCESSORS AVG WEIGHT <sup>(See note 1)</sup>	The average weight of shared processors, which is the sum of shared processor weights related to the number of shared processors.

**Notes:**

1. For CFLEVEL lower than 15, this field is not displayed.
2. SCM statistics are included in the SCM Structure Summary only for those structures that can make use of the SCM storage extension and have set a non-zero maximum SCM size. If none of the structures is configured to exploit SCM, the SCM Structure Summary displays message: "NO STORAGE CLASS MEMORY DATA AVAILABLE".
3. If no structure is configured to exploit asynchronous duplexing, the Asynchronous CF Duplexing Summary includes an informational message "NO ASYNCHRONOUS CF DUPLEXING DATA AVAILABLE".

**Coupling Facility Structure Activity section**

This section of the Coupling Facility Activity report has detail for each active structure in the coupling facility, including activity data for each system connected to the structure during the reporting interval.

COUPLING FACILITY ACTIVITY												PAGE 3	
z/OS V2R2		SYSPLEX UTCPLXJ8			DATE 09/28/2016			INTERVAL 030.00.000					
		RPT VERSION V2R2 RMF			TIME 12.00.00			CYCLE 01.000 SECONDS					
-----													
COUPLING FACILITY NAME = X7CFP87													
-----													
COUPLING FACILITY STRUCTURE ACTIVITY													
-----													
STRUCTURE NAME = DBSVPLX7_SCA TYPE = LIST STATUS = ACTIVE SECONDARY													
# REQ ----- REQUESTS ----- DELAYED REQUESTS -----													
SYSTEM	TOTAL	#	% OF	-SERV	TIME(MIC)-	REASON	#	% OF	---	AVG	TIME(MIC)	----	
NAME	AVG/SEC	REQ	ALL	AVG	STD_DEV	REQ	REQ	REQ	/DEL	STD_DEV	/ALL		
R7D	7887	SYNC	32	0.4	74.9	12.5	NO SCH	0	0.0	0.0	0.0	0.0	0.0
	4.38	ASYN	7855	100	104.8	41.3	PR WT	0	0.0	0.0	0.0	0.0	0.0
		CHNGD	0	0.0	INCLUDED	IN ASYN	PR CMP	2960	37.5	7.6	14.4	2.8	
		SUPPR	0	0.0			DUMP	0	0.0	0.0	0.0	0.0	
...													
TOTAL	7887	SYNC	32	0.4	74.9	12.5	NO SCH	0	0.0	0.0	0.0	0.0	0.0
	4.38	ASYN	7855	100	104.8	41.3	PR WT	0	0.0	0.0	0.0	0.0	0.0
		CHNGD	0	0.0			PR CMP	2960	37.5	7.6	14.4	2.8	
		SUPPR	0	0.0			DUMP	0	0.0	0.0	0.0	0.0	
-----													
STRUCTURE NAME = DBSVPLX7_LOCK1 TYPE = LOCK STATUS = ACTIVE PRIMARY													
# REQ ----- REQUESTS ----- DELAYED REQUESTS -----													
SYSTEM	TOTAL	#	% OF	-SERV	TIME(MIC)-	REASON	#	% OF	---	AVG	TIME(MIC)	----	EXTERNAL REQUEST
NAME	AVG/SEC	REQ	ALL	AVG	STD_DEV	REQ	REQ	REQ	/DEL	STD_DEV	/ALL		CONTENTIONS
R7D	68208	SYNC	272	0.4	71.7	18.2	NO SCH	0	0.0	0.0	0.0	0.0	REQ TOTAL 82K
	37.89	ASYN	68K	100	98.4	50.6	PR WT	68K	100	1.1	0.4	1.1	REQ DEFERRED 8605
		CHNGD	0	0.0	INCLUDED	IN ASYN	PR CMP	21K	31.1	5.9	11.8	1.8	-CONT 0
		SUPPR	0	0.0									-FALSE CONT 0
...													
TOTAL	68208	SYNC	272	0.4	71.7	18.2	NO SCH	0	0.0	0.0	0.0	0.0	REQ TOTAL 82K
	37.89	ASYN	68K	100	98.4	50.6	PR WT	68K	100	1.1	0.4	1.1	REQ DEFERRED 8605
		CHNGD	0	0.0			PR CMP	21K	31.1	5.9	11.8	1.8	-CONT 0
		SUPPR	0	0.0									-FALSE CONT 0
-----													
STRUCTURE NAME = IRRXCF00_B003 TYPE = CACHE STATUS = ACTIVE													
# REQ ----- REQUESTS ----- DELAYED REQUESTS -----													
SYSTEM	TOTAL	#	% OF	-SERV	TIME(MIC)-	REASON	#	% OF	---	AVG	TIME(MIC)	----	
NAME	AVG/SEC	REQ	ALL	AVG	STD_DEV	REQ	REQ	REQ	/DEL	STD_DEV	/ALL		
R7D	0	SYNC	0	0.0	0.0	0.0	NO SCH	0	0.0	0.0	0.0	0.0	0.0
	0.00	ASYN	0	0.0	0.0	0.0	PR WT	0	0.0	0.0	0.0	0.0	0.0
		CHNGD	0	0.0	INCLUDED	IN ASYN	PR CMP	0	0.0	0.0	0.0	0.0	0.0
		SUPPR	0	0.0			DUMP	0	0.0	0.0	0.0	0.0	0.0
...													
TOTAL	0	SYNC	0	0.0	0.0	0.0	NO SCH	0	0.0	0.0	0.0	0.0	-- DATA ACCESS ---
	0.00	ASYN	0	0.0	0.0	0.0	PR WT	0	0.0	0.0	0.0	0.0	READS 0
		CHNGD	0	0.0			PR CMP	0	0.0	0.0	0.0	0.0	WRITES 0
		SUPPR	0	0.0			DUMP	0	0.0	0.0	0.0	0.0	CASTOUTS 0
													XI'S 0

Figure 182. Coupling Facility Activity Report - Structure Activity

The following table explains the field headings in the Structure Activity section.

Table 135. Fields in the Coupling Facility Activity Report - Structure Activity

Field Heading	Meaning
STRUCTURE NAME	The name given to the structure by the coupling facility policy specification in the Function Couple Data Set. It is up to 16 characters and is unique within a sysplex.
TYPE	Indicates whether the structure is a list, lock, or cache structure. If it is a lock structure, then the contention counts are included in the report.
STATUS	Indicates status of the structure at the end of the interval. For the description of possible values refer to Table 134 on page 320.



Table 135. Fields in the Coupling Facility Activity Report - Structure Activity (continued)

Field Heading	Meaning
SYSTEM NAME	The system name for the system connected to the structure (from IEASYSxx Parmlib member, SYSNAME parameter)  The name is preceded by an '*' if the data for this system is incomplete for this interval, for example because the gatherer has been stopped.
# REQ TOTAL # REQ AVG/SEC	The sum of all requests (internal and external) that utilize the subchannel. Specifically: <ul style="list-style-type: none"> <li>External requests to send/receive data on behalf of a structure. The sum of synchronous and asynchronous requests completed against any structure within this coupling facility per second. This includes requests that changed from synchronous to asynchronous.</li> <li>Internal requests that utilize the subchannels (but are not aggregated by the structure).</li> </ul> <b>TOTAL</b> Total number of requests <b>AVG/SEC</b> Average number of requests per second for this structure  This field offers a quick way of determining which systems are generating the most activity for a given structure, and indicates where to focus tuning or load balancing efforts.
REQUESTS	The requests are shown in four categories described hereafter: SYNC, ASYNC, CHNGD, and SUPPR.
SYNC	Total number of hardware operations that started and completed synchronously to the coupling facility on behalf of connectors to the structure.
ASYNC	Total number of hardware operations that started and completed asynchronously to the coupling facility on behalf of connectors to the structure.  The service time is the time for all ASYNC requests (ASYNC and CHNGD).
CHNGD	Total number of hardware operations that changed from synchronous to asynchronous because the operation could not be serviced as synchronous operation. This field reports only those operations which were changed due to a subchannel busy condition and can be used as an indicator of a shortage of subchannel resources.  Conversions caused by heuristic sync/async algorithms used to optimize the coupling efficiency of workloads using the CF are not included.
SUPPR	Number of requests whose execution was suppressed by the coupling facility in order to avoid a potential serialization deadlock condition across a duplexed pair of structures. This field does not apply to asynchronously duplexed structures.
# REQ % OF ALL (valid for SYNC, ASYNC, CHNGD, and SUPPR)	The number of requests for this structure, and the percentage this represents of all requests for this structure from any system.
SERVICE TIME - AVG	The average time in microseconds required to satisfy a coupling facility request for this structure.
SERVICE TIME - STD_DEV	The standard deviation of service time for this structure.  Even though the average time looks acceptable, the standard deviation could be high, indicating that there is a wide fluctuation in service times for requests. In this case, analyze the coupling facility configuration for possible path or coupling facility bottlenecks in the <i>Subchannel Activity</i> section.
DELAYED REQUESTS	These columns list possible contention reasons for requests sent to the coupling facility.

Table 135. Fields in the Coupling Facility Activity Report - Structure Activity (continued)

Field Heading	Meaning
REASON	The reason for a delayed request can be either a subchannel contention (NO SCH) or a dump serialization (DUMP).  For synchronous duplexed requests, also peer subchannel wait time (PR WT) and waiting-for-peer-completion time (PR CMP) is reported. A duplexed request requires two subchannels. PR WT is the time (in microseconds) between the moment when the request was sent to the other duplexed structure instance and when it is sent to this one. PR CMP is the time (in microseconds) between the moment when this structure responded to z/OS and when the other structure instance responded. Both subchannels are busy until the responses from both structure instances are processed by z/OS.
# REQ	The total number and the percentage of requests delayed in the interval.
% of REQ	
AVG TIME - /DEL	The average delay time in microseconds over all delayed requests.
AVG TIME - STD_DEV	The standard deviation to the average delay time.
AVG TIME - /ALL	The average delay time in microseconds over all requests, whether delayed or not.
EXTERNAL REQUEST CONTENTIONS	These values are available for all serialized list structures.
REQ TOTAL	The number of requests against this structure.
REQ DEFERRED	The number of requests running into a lock contention
EXTERNAL REQUEST CONTENTIONS	These values are available for all lock structures.
REQ	Total requests issued for the lock structure
REQ DEFERRED	Subset of the above field indicating the number of requests that were unable to complete within the request issuer's thread. That is, any request that needed additional processing to complete.
-CONT	A subset of the REQ DELAYED field. It presents the number of requests delayed due to contention on a lock. <b>Example:</b>  A lock is held by an EXCLUSIVE request, and another request is made for the same lock with EXCLUSIVE or SHARE specified. If this number is high it could indicate an impact to the end user of the application or subsystem owning the lock structure. Refer to that application's traces or reports for more detail on what locks caused the heavy contention.
-FALSE CONT	A subset of the CONT field showing the number of requests that experience "hash contention". This occurs because a hashing algorithm is used to map a lock request to a lock table entry. When more than one lock request maps to the same entry, there is the potential for contention delay. You may need to increase the size of the lock table. <b>Note:</b> It is possible for an application to have unusual lock reference patterns that cause storage contention regardless of the size of the lock structure.
TOTAL	This row of data gives totals (or overall averages and percentages) for all the systems connected to the structure,
DATA ACCESS	This information is shown for cache structures.
READS	The number of occurrences the coupling facility returned data on a read request by any connector (read hit).  Directory only caches will always have a zero value reported since there are no data to be returned.
WRITES	The number of occurrences data has been written to the cache structure.  Directory only caches will always have a zero value reported since there are no data writes possible.

Table 135. Fields in the Coupling Facility Activity Report - Structure Activity (continued)

Field Heading	Meaning
CASTOUTS	<p>The number of times CASTOUT processing occurs.</p> <p>This is the process of writing changed cache data to permanent storage.</p> <p>This counter is of interest for store-in cache structures (for example, DB2 global buffer pool structures) in determining the volume of changed data being removed from the structure.</p>
XI'S	<p>The number of times a data item residing in a local buffer pool was marked invalid by the coupling facility.</p> <p>XI's count values are seen for directory, store-in and store-thru caches. This count reflects the amount of data sharing among the users of the cache and the amount of write or update activity against the data bases.</p>

### Subchannel Activity section

This section contains a summary line for each system attached to the coupling facility. MVS treats the set of available subchannels for a coupling facility as a pool of resources for any request to that facility. Therefore, the subchannel activity data is not reported by individual subchannel. MVS handles the load balancing across the subchannels automatically.

```

          COUPLING FACILITY ACTIVITY

z/OS V2R2          SYSPLEX UTCPLXJ8          DATE 09/28/2016          INTERVAL 030.00.000
                   RPT VERSION V2R2 RMF          TIME 13.00.00          CYCLE 01.000 SECONDS

-----
COUPLING FACILITY NAME = CX7CFP87
-----
                                SUBCHANNEL ACTIVITY
-----
# REQ          # REQUESTS          # % OF          DELAYED REQUESTS
SYSTEM TOTAL  -- CF LINKS -- PTH  # -SERVICE TIME(MIC)-  # % OF          # % OF          AVG TIME(MIC)
NAME  AVG/SEC TYPE GEN USE BUSY  REQ  AVG  STD_DEV  REQ  REQ  /DEL  STD_DEV  /ALL
-----
R7D   3599K CIB  4  4  0  SYNC  3301K  23.5  9.7  LIST/CACHE  0  0.0  0.0  0.0  0.0
      1999.5 SUBCH 28 28  ASYNC  236454  84.1  139.7  LOCK  0  0.0  0.0  0.0  0.0
      CHANGED 0 INCLUDED IN ASYNC TOTAL  0  0.0
      UNSUCC 0  0.0  0.0
R70   436212K ICP  4  4  3482 SYNC  435637K  5.2  3.6  LIST/CACHE  18K  0.0  285.0  215.6  0.0
      242340 SUBCH 28 28  ASYNC  141411  70.9  152.5  LOCK  0  0.0  0.0  0.0  0.0
      CHANGED 17622 INCLUDED IN ASYNC TOTAL  18K  0.0
      UNSUCC 0  0.0  0.0

...

                                CHANNEL PATH DETAILS
-----
SYSTEM NAME  ID  TYPE  OPERATION MODE  DEGRADED  DISTANCE  PCHID  AID  PORT  IOP  IDS
-----
R7D          C4  CIB  1X IFB HCA3-0 LR  N  <1  704  000D  01  06
           C5  CIB  1X IFB HCA3-0 LR  N  <1  705  000D  01  06
           C6  CIB  1X IFB HCA2-0 LR  N  <1  706  000C  02  05
           C7  CIB  1X IFB HCA2-0 LR  N  <1  707  000C  02  05
    
```

Figure 183. Coupling Facility Activity Report - Subchannel Activity

Table 136. Fields in the Coupling Facility Activity Report - Subchannel Activity

Field Heading	Meaning
SYSTEM NAME	<p>The name of the system attached to the coupling facility (from IEASYSxx Parmlib member, SYSNAME parameter).</p> <p>The name is preceded by an '*' if the data for this system is incomplete for this interval, for example because the gatherer has been stopped.</p>

Table 136. Fields in the Coupling Facility Activity Report - Subchannel Activity (continued)

Field Heading	Meaning
# REQ TOTAL # REQ AVG/SEC	<p><b>TOTAL</b> Total number of requests to this facility. This number will usually be greater than the sum of the individual structure values from the previous report section because it includes global coupling facility commands that are not attributable to any structure.</p> <p><b>AVG/SEC</b> Average number of requests per second for this facility.</p> <p>This field can be used as a quick way of determining which systems are generating the most activity for a given facility which in turn indicates where to focus tuning or load balancing efforts.</p>
CF LINKS	<p><b>TYPE</b> Channel path type.</p> <p><b>GEN</b> Number of subchannels that are defined.</p> <p><b>USE</b> Number of subchannels MVS is currently using for coupling facility requests.</p>
PTH BUSY	<p>Path busy - the number of times a coupling facility request was rejected because all paths to the coupling facility were busy.</p> <p>A high count combined with elongated service times for requests indicates a capacity constraint in the coupling facility. If coupling facility channels are being shared among PR/SM partitions, the contention could be coming from a remote partition.</p> <p>Identifying path contention: There can be path contention even when this count is low. In fact, in a non-PR/SM environment where the subchannels are properly configured, the total number of delayed requests, and not PTH BUSY, is the indicator for path contention. If this value is high, it means MVS is delaying the coupling facility requests and in effect gating the workload before it reaches the physical paths. Before concluding you have a capacity problem, however, be sure to check that the correct number of subchannels are defined in the I/O gen.</p> <p>PR/SM environment only: If coupling facility channels are being shared among PR/SM partitions, PTH BUSY behaves differently. You potentially have many MVS subchannels mapped to only a few coupling facility command buffers. You could have a case where the subchannels were properly configured (or even under-configured), subchannel busy is low, but path busy is high. This means the contention is due to activity from a remote partition.</p>
REQUESTS - The requests are shown in four categories.	
# REQ SYNC	Number of requests from this system to the coupling facility that started as synchronous requests which are completed (synchronously or asynchronously).
# REQ ASYNC	Number of completed requests which have been started as asynchronous requests.
# REQ CHANGED	Number of requests changed from synchronous to asynchronous because the requests could not be serviced as synchronous request.
# REQ UNSUCC	Number of requests which could not be completed due to hardware problems. This number should normally be zero. If it is non-zero, there is a hardware problem that needs to be investigated. The reason it is reported here is to judge to what impact extent hardware problem(s) impact coupling facility performance.
SERVICE TIME - AVG SERVICE TIME - STD_DEV	The average service time in microseconds and the standard deviation of the service time spent for requests to the coupling facility. The average service time in conjunction with its standard deviation can be used to determine potential impacts to the end user. Even though the average service time is low the standard deviation can be high indicating a wide fluctuation. This category is for the request types SYNC, ASYNC, and UNSUCC, the fields are not applicable for column CHANGED.
DELAYED REQUESTS - These columns lists possible contention reasons for requests sent to the coupling facility.	
# REQ LIST/CACHE	Number of delayed requests across all LIST and CACHE structures.
# REQ LOCK	Number of delayed requests across all LOCK structures.
# REQ TOTAL	Number of delayed requests across all structures.
% OF REQ	The percentage of requests delayed, related to the number of List/Cache requests, Lock requests and total requests.
AVG TIME - /DEL	The average delay time in microseconds over all delayed requests.

Table 136. Fields in the Coupling Facility Activity Report - Subchannel Activity (continued)

Field Heading	Meaning
AVG TIME - STD_DEV	The standard deviation to the average delay time.
AVG TIME - /ALL	The average delay time in microseconds over all requests, whether delayed or not.

Table 137. Fields in the Coupling Facility Activity Report - Subchannel Activity - Channel Path Details

Field Heading	Meaning
<b>Note:</b> If the hardware cannot provide values for a measurement, the field remains blank.	
SYSTEM NAME	The name of the system attached to the coupling facility (from IEASYSxx Parmlib member, SYSNAME parameter).
ID	The hexadecimal identifier of a channel path (CHPID) that is connected to the coupling facility.
TYPE	Channel path type.
OPERATION MODE	Channel path operation mode. It describes the data rate, bandwidth, protocol, and adapter type of the channel path.  A data rate of, for example, 1GBIT denotes a rate of 1.0625 gigabit per second.  A bandwidth of, for example, 12X denotes a twelve-fold bandwidth.  Protocols: <ul style="list-style-type: none"> <li>• IFB – InFiniBand</li> <li>• IFB3 – InFiniBand 3</li> <li>• GEN3 – PCIe third generation protocol</li> </ul> Adapter types: <ul style="list-style-type: none"> <li>• HCA2-O – Host Channel Adapter2-optical</li> <li>• HCA2-O LR – Host Channel Adapter2-optical long reach</li> <li>• HCA3-O – Host Channel Adapter3-optical</li> <li>• HCA3-O LR – Host Channel Adapter3-optical long reach</li> <li>• PCIE-O SR – Peripheral Component Interconnect Express short reach</li> </ul> Unknown operation mode: <ul style="list-style-type: none"> <li>• UNKNOWN</li> </ul>
DEGRADED	Character Y in this column indicates that the channel path is operating at reduced capacity (degraded) or not operating at all.
DISTANCE	Estimated distance in kilometers. The value is calculated as follows: Average round-trip path time in microseconds ----- 10 microseconds / kilometer  A value of zero means that the time was not measured.
PCHID	Physical channel identifier.
AID	The hexadecimal coupling adapter identifier associated with the channel path.
PORT	The hexadecimal port associated with the channel path.
IOP IDS	The hexadecimal identifiers of I/O processors (System Assist Processors) to which the channel path is accessible.

### CF to CF Activity section

COUPLING FACILITY ACTIVITY																
z/OS V2R2		SYSPLEX UTCPLXJ8			DATE 04/11/2016			INTERVAL 030.00.000							PAGE 6	
		RPT VERSION V2R2 RMF			TIME 13.39.00			CYCLE 1.000 SECONDS								
-----																
COUPLING FACILITY NAME = X7CFP87																
-----																
CF TO CF ACTIVITY																
-----																
PEER CF	-RECEIVER- TYPE USE	--SENDER-- TYPE USE				REQUESTS				DELAYED REQUESTS						
						# REQ	AVG/ SEC	-SERVICE TIME(MIC)- AVG	STD_DEV	# REQ	% OF REQ	----- /DEL	AVG TIME(MIC) STD_DEV	----- /ALL		
X7CFH89	CS5	3	CS5	3	SYNC	0	0.0	0.0	0.0	0	0.0	0.0	0.0	0.0		
X7CFP87	ICP	8	ICP	8	SYNC	0	0.0	0.0	0.0	0	0.0	0.0	0.0	0.0		
X7CFP89	CIB	2	CIB	2	SYNC	0	0.0	0.0	0.0	0	0.0	0.0	0.0	0.0		
-----																
CHANNEL PATH DETAILS																
-----																
PEER CF	ID	TYPE	R/S	OPERATION MODE			DEGRADED	DISTANCE								
X7CFH89	24	CS5	R	8X	GEN3	PCIE-0	SR	N	<1							
	24	CS5	S	8X	GEN3	PCIE-0	SR	N	<1							
	25	CS5	R	8X	GEN3	PCIE-0	SR	N	<1							
	25	CS5	S	8X	GEN3	PCIE-0	SR	N	<1							
	26	CS5	R	8X	GEN3	PCIE-0	SR	N	<1							
	26	CS5	S	8X	GEN3	PCIE-0	SR	N	<1							
X7CFP89	B9	CIB	R	1X	IFB	HCA3-0	LR	N	<1							
	B9	CIB	S	1X	IFB	HCA3-0	LR	N	<1							
	BA	CIB	R	1X	IFB	HCA3-0	LR	N	<1							
	BA	CIB	S	1X	IFB	HCA3-0	LR	N	<1							

Figure 184. Coupling Facility Activity Report - CF to CF Activity

Table 138. Fields in the CF to CF Activity Section

Field Heading	Meaning
PEER CF	Name of the remote coupling facility.
RECEIVER SENDER	<p><b>TYPE</b> CF link type of receiver/sender channel paths.</p> <p><b>USE</b> The number of receiver/sender paths of named type used for coupling facility communication.</p> <p>Detail data on sender channel paths could be blank if you are using data from an old RMF gatherer or have preallocated data from a previous release of RMF.</p>
REQUESTS	The requests are synchronous (SYNC).
# REQ	<p>The sum of the following signals that have been sent from the subject CF to the remote CF:</p> <ul style="list-style-type: none"> <li>• Number of halt execution signals.</li> <li>• Number of ready to complete signals.</li> <li>• Number of ready to execute signals.</li> <li>• Number of request suppression signals.</li> <li>• Number of request for suppression accepted signals.</li> </ul>
AVG/SEC	Average number of signals/messages per second.
SERVICE TIME - AVG	The average service time in microseconds for all kind of signals that have been sent from the subject CF to the remote CF, including redrives, excluding any delay time.
SERVICE TIME - STD_DEV	The standard deviation of the average service time.
DELAYED REQUESTS	The delayed requests are synchronous (SYNC).
# REQ	The number of signals of all types which have experienced a delay in being sent from the subject CF to this remote CF.
% OF REQ	The percentage of requests delayed.
AVG TIME - /DEL	The average delay time in microseconds over all delayed requests.
AVG TIME - STD_DEV	The standard deviation to the average delay time.
AVG TIME - /ALL	The average delay time in microseconds over all requests, whether delayed or not.

**Note:** If the hardware cannot provide values for a measurement, the field remains blank.

Table 139. Fields in the Coupling Facility Activity Report - CF to CF Activity - Channel Path Details

Field Heading	Meaning
PEER CF	Name of the remote coupling facility.
ID	The hexadecimal identifier of a channel path (CHPID) that is connecting both coupling facilities with each other.
TYPE	Channel path type.
R/S	<b>R</b> Receiver channel path. <b>S</b> Sender channel path.
OPERATION MODE	Channel path operation mode. It describes the data rate, bandwidth, protocol, and adapter type of the channel path.  For more information about displayed values, refer to Table 137 on page 331.
DEGRADED	Character <b>Y</b> in this column indicates that the channel path is operating at reduced capacity (degraded) or not operating at all.
DISTANCE	Estimated distance in kilometers.  For more information, refer to Table 137 on page 331.

## Spreadsheet and Overview reference

You can make this report available through Overview records in a spreadsheet, using the Spreadsheet Reporter. The following table shows all criteria and the corresponding Overview names for creating Overview records. For details, see the *z/OS RMF User's Guide*.

Table 140. Overview names in the Coupling Facility Activity Report

Field Heading or Meaning	Overview Name
Average service time of SYNC operations	SYNCST
SYNC operation rate	SYNCRT
Average service time of ASYNC operations	ASYNCS
Ended ASYNC operation rate	ASYNCR
Percentage of changed operations	CHNGDP
Changed operation rate	CHNGDRT
Path busy rate	PBSY
Percent delayed requests	DREQP
CF processor utilization	CFUTIL
Directory reclaims	DIRRCLM
List/directory entries: current to total ratio	LDECTR
Data elements: current to total ratio	DECTR
Lock entries: current to total ratio	LECTR
Cache read request rate	CREADRT
Cache write request rate	CWRITERT
Cache castout rate	CCOUTRT
Cache cross invalidation rate	CXIRT
Total requests to lock structure or serialized list structure	LCKREQ
Contention on lock structure	LCKCONT

Table 140. Overview names in the Coupling Facility Activity Report (continued)

Field Heading or Meaning	Overview Name
False contention on lock structure	LCKFCNT
Percentage of CF utilization	STUTILP
Percentage of subchannel busy	SUBCHBP
Percentage of storage class memory in use	SCMIUP
Percentage of augmented space in use	AUGMIUP
SCM list entry current to total ratio	SCMLCTR
SCM list element current to total ratio	SCMLECTR
Average service time per SCM read operation	SCMRST
Average service time per SCM write operation	SCMWST
SCM auxiliary enabled commands to total request ratio	SCMAUXR
SCM delayed faults to total request ratio	SCMDFR

---

## CHAN - Channel Path Activity report

The Channel Path Activity report provides information about channel path use.

The report identifies each channel path by identifier and channel path type, and reports both the total channel utilization by the central processing complex (CPC) and the channel utilization of the individual system image (partition).

Data for total utilization and partition utilization is gathered independently. Because the internal interval used to gather this data is a few seconds, the total utilization and the sum of the partition's utilization sharing that channel might differ if a short RMF interval is specified. If the interval is too small or the appropriate data cannot be gathered, dashes (---) are displayed instead of data. Please refer to the information APAR II05151 for a list of channel types for which channel utilization data is not gathered.

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

**NOW ONLINE**

Brought online during the interval and still online at the end of the interval

**NOW OFFLINE**

Taken offline during the interval and still offline at the end of the interval

**OFFLINE**

Offline for the entire interval

**DELETED**

Deleted during the interval

**MODIFIED**

Modified during the interval

**INSTALLED**

Installed during the interval



For all channels that are managed by dynamic channel path management (DCM), additional information is available. DCM allows an installation to identify channels that they wish to be managed dynamically. These channels are not assigned permanently to a specific control unit, but belong to a pool of channels. Based on workload requirements in the system, these channels will be assigned dynamically by DCM. On top of the report, there is a consolidated data section for managed channels displaying the total number of channel paths for each type and the average activity data. The character **M** as suffix of the acronym for the channel path type is an indicator that the channel is managed by DCM.

## Duration report

Any channel that moved online or offline during the duration interval is indicated by an asterisk following the channel identifier.

In this report, the mode of the central processing complex (CPC) can be the following:

### BASIC

The report shows all channels configured in the system. Only data for total utilization is reported. The partition utilization column is blank.

### NOW BASIC

The report shows the last active mode. If you combine the SMF records from before and after a power-on-reset (POR) and changed the mode, two modes appear in the SMF records. By combining the intervals of the SMF records into one duration report, RMF displays the last active mode in the mode field. The partition utilization column is blank.

**LPAR** The report shows the individual PR/SM logical partition's utilization and the total utilization of the shared ESCON channels, and the partition's and total utilization of the unshared channels.

### NOW LPAR

The report shows the last active mode. If you combine the SMF records from before and after a POR and changed the mode, two modes appear in the SMF records. By combining the intervals of the SMF records into one duration report, RMF displays the last active mode in the mode field. The partition utilization column is blank.

You can use channel path activity information together with I/O device activity and I/O queuing activity information to identify performance bottlenecks associated with channel paths. To find out which logical control unit is using the channel, look in the I/O Queuing Activity report. From there you can go to check device response times. For example, if a channel path to a device shows excessive use, you could define additional paths to the device or introduce a different job mix to produce better performance.

## How to request this report

Monitor I gathers data for this report automatically. If you want to suppress gathering, you need to specify NOCHAN.

To produce this report, specify  
REPORTS(CHAN)

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

**Example URL for the DDS API:**

http://ddshost:8803/gpm/rmfpp.xml?reports=CHAN

## Contents of the report

### Notes:

1. On a machine running in LPAR mode, but with only one LPAR defined, the *PART* columns for the *READ*, *WRITE* and *UTILIZATION* fields display a zero value for channels of type FC (FICON).
2. When Channel Path Measurement Facility (CPMF) is not available, for example, on z/OS systems running as z/VM guests, RMF uses sampled data from SRM so that the reported channel utilization is only an approximate value. With increasing channel speed, the channel utilization value becomes more and more inaccurate. Therefore, in such cases, RMF does not provide accurate values of FICON channel utilization.

Beginning with z990 processors, the channel data from SRM is no longer available. As a result, the channel utilization data on a z/OS system running as z/VM guest, is reported as -----

CHANNEL PATH ACTIVITY															PAGE	1	
z/OS V2R2		SYSTEM ID CB88		DATE 09/28/2016		INTERVAL 14.59.999											
		RPT VERSION V2R2 RMF		TIME 08.00.00		CYCLE 1.000 SECONDS											
IODF = 8E		CR-DATE: 09/14/2016		CR-TIME: 16.47.03		ACT: POR		MODE: LPAR		CPMF: EXTENDED MODE		CSSID: 2					
DETAILS FOR ALL CHANNELS																	
CHANNEL PATH		UTILIZATION(%)		READ(MB/SEC)		WRITE(MB/SEC)		FICON OPERATIONS			ZHPF OPERATIONS						
ID	TYPE	G	SPEED	SHR	PART	TOTAL	BUS	PART	TOTAL	PART	TOTAL	RATE	ACTIVE	DEFER	RATE	ACTIVE	DEFER
00	OSD		10G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
02	OSE		1G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
04	OSD		10G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
05	OSD		10G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
...																	
74	CIB			Y	-----	-----											
75	CIB			Y	-----	-----											
92	ICP			Y	-----	-----											
93	ICP			Y	-----	-----											
C0	FC_S	13	8G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.0	0.0	0.0	0.0	0.0	0.0
CHANNEL PATH		UTILIZATION(%)		READ(MB/SEC)		WRITE(MB/SEC)		PHYSICAL NETWORK IDS									
ID	TYPE	G	SPEED	SHR	PART	TOTAL	BUS	PART	TOTAL	PART	TOTAL	PORT 1	PORT 2				
07	OSD		1G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	NETWORK1	NETWORK1				
08	OSD		1G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	NETWORK1					
0C	OSD		1G	Y	0.00	0.00	0.09	0.00	0.00	0.00	0.00	NETWORK7C5	NETWORK7C5				
13	OSD		1G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	NETWORK7C2	NETWORK7C2				
15	OSD		10G	Y	0.00	0.00	0.16	0.00	0.00	0.00	0.00	NETWORK1					
...																	
CHANNEL PATH		WRITE(B/SEC)		MESSAGE RATE		MESSAGE SIZE		SEND FAIL		RECEIVE FAIL		PHYSICAL NETWORK ID					
ID	TYPE	G	SHR	PART	TOTAL	PART	TOTAL	PART	TOTAL	PART	TOTAL	PORT 1	PORT 2				
E5	IQD			0	0	0	0	-----	-----	0	0	0	0				
F1	IQD			0	0	0	0	-----	-----	0	0	0	0				
F2	IQD			0	0	0	0	-----	-----	0	0	0	0				

Figure 185. Channel Path Activity report

Table 141. Fields in the Channel Path Activity report

Field Heading	Meaning
IODF = xx	The IODF number where xx is the suffix of the IODF data set name.
CR-DATE: mm/dd/yyyy	The creation date of the IODF.
CR-TIME: hh.mm.ss	The creation time of the IODF.
ACT: text	The configuration state where text indicates how the IODF was activated.

Table 141. Fields in the Channel Path Activity report (continued)

Field Heading	Meaning
MODE	The mode of the central processing complex (CPC): <b>BASIC</b> The report shows all channels configured in the system. <b>LPAR</b> The report shows both the total utilization and the individual partition's utilization of all channels configured to the logical partition.
CPMF	The availability of the Channel Path Measurement Facility (CPMF). CPMF allows RMF to report channel utilization information for individual partitions. The value can be: <b>COMPATIBILITY MODE</b> CPMF is running in compatibility mode. <b>EXTENDED MODE</b> CPMF is running in extended mode. <b>NOT AVAILABLE</b> CPMF is not available on the system.  The indication (CHANGED) will be shown if the CPMF mode has changed during the reporting interval. In that case, only TOTAL values will be reported.  For more information about CPMF, see the data area IRACPMB in MVS Data Areas Volume 2 available from the z/OS Internet Library.
CSSID	This field is shown only for z990 processors or follow-on processors and denotes the ID of the monitored logical channel subsystem.
CHANNEL GROUP G NO	For each channel type which is managed by DCM, a summary line is shown with the average values for all channels in this group.  G indicates the generation and is used to differentiate between channels of the same channel type, when one has significant differences from the other. Newer generations with significant differences (for example, the channel throughput) are indicated by a number (1, 2, ...). For example, for a FICON channel, a number 1 indicates that the channel has an auto-negotiated throughput of 1 Gbit/sec, or a number 2 indicates a throughput of 2 Gbit/sec.  The number of channels of the group is given in column NO.
CHANNEL PATH ID	The hexadecimal channel path identifier (CHPID).
CHANNEL PATH TYPE	Type of channel path.  You may issue the console command D M=CHP(xx) to see an explanation of the channel path type.  If RMF encounters an error while processing the TYPE data, this field is blank. RMF continues to measure channel path activity. Check the operator console for messages.
CHANNEL PATH G	This column indicates the generation and is used to differentiate between channels of the same channel type, when one has significant differences from the other. Newer generations with significant differences are indicated by a number (1, 2, ...). For example, for z/OS, a number 2 indicates that a FICON channel has auto negotiated to a link speed of 2 GB/sec.
CHANNEL PATH SPEED	The channel path speed in bits per second at the end of the interval.
CHANNEL PATH SHR	The indication of whether the channel path is defined as shared between one or more logical partitions. Y indicates that the channel path is shared. A blank indicates it is not shared.
PHYSICAL NETWORK ID(S)	Physical-network identifiers (PNET IDs) of an Ethernet network that is accessible from the ports of the channel path.

PP - CHAN

Table 141. Fields in the Channel Path Activity report (continued)

Field Heading	Meaning
UTILIZATION (%) PART	<p>The channel path utilization percentage for an individual logical partition. RMF uses the values provided by CPMF.</p> <p>In LPAR mode, the calculation is:</p> $\text{PART UTILIZATION (\%)} = \frac{\text{Channel Path Busy Time}}{\text{Channel Path Elapsed Time}} * 100$ <p>For channels like FICON, OSA Express, or OSA Direct Express, which are running in extended CPMF mode, the calculation is as follows:</p> $\text{Part Utilization (\%)} = \frac{\text{LPAR \# of Channel Work Units}}{\text{Max \# of Channel Work Units} * \text{Channel Path Elapsed Time}} * 100$ <p>For some channels like OSAEGbE, FICON EXPRESS/EXPRESS2, this value reflects the microprocessor utilization.</p> <p>For hipersockets, this value is not available.</p>
UTILIZATION (%) TOTAL	<p>The channel path utilization percentage for the CPC during an interval.</p> <p>For processors earlier than z990 and shared channels in LPAR mode, where CPMF is not available, or for all channels in BASIC mode with CPMF not available, the calculation is:</p> $\text{Total Utilization (\%)} = \frac{\text{\# SRM Observations of Channel Path Busy}}{\text{\# Samples}} * 100$ <p>For unshared channels in LPAR mode, the value for total utilization is the same as partition utilization.</p> <p>For channels like FICON, OSA Express, or OSA Direct Express, which are running in extended CPMF mode, the calculation is as follows:</p> $\text{Total Utilization (\%)} = \frac{\text{Total \# of Channel Work Units}}{\text{Max \# of Channel Work Units} * \text{Channel Path Elapsed Time}} * 100$ <p>For some channels like OSAEGbE, FICON EXPRESS/EXPRESS2, this value reflects the microprocessor utilization.</p> <p>For hipersockets, this value is not available.</p>
UTILIZATION (%) BUS	<p>Percentage of bus cycles, the bus has been found busy for this channel in relation to the theoretical limit.</p> <p>For OSAEGbE, the value reflects the PCI bus utilization.</p> <p>For hipersockets, this value is not available.</p>
READ(MB/SEC)	<p><b>PART</b> Data transfer rates from the control unit to the channel for this partition.</p> <p><b>TOTAL</b> Data transfer rates from the control unit to the channel for the CPC.</p> <p>For hipersockets, this value is not available.</p>
WRITE(MB/SEC)	<p><b>PART</b> Data transfer rates from the channel to the control unit for this partition.</p> <p><b>TOTAL</b> Data transfer rates from the channel to the control unit for the CPC.</p>

Table 141. Fields in the Channel Path Activity report (continued)

Field Heading	Meaning
FICON OPERATIONS	<p><b>RATE</b> Number of native FICON operations per second.</p> <p><b>ACTIVE</b> The average number of native FICON operations that are concurrently active during the reporting interval.</p> <p><b>DEFER</b> Number of deferred native FICON operations per second that could not be initiated by the channel due to the lack of available resources.</p> <p>This field is reported for the CPC.</p>
ZHPF OPERATIONS	<p><b>RATE</b> Number of zHPF (High Performance FICON) operations per second.</p> <p><b>ACTIVE</b> The average number of zHPF operations that are concurrently active during the reporting interval.</p> <p><b>DEFER</b> Number of deferred zHPF operations per second that could not be initiated by the channel due to the lack of available resources.</p> <p>This field is reported for the CPC.</p>
WRITE(B/SEC)	<p><b>PART</b> Data transfer rates from the channel to the control unit for this partition.</p> <p><b>TOTAL</b> Data transfer rates from the channel to the control unit for the CPC.</p> <p>The values are shown in bytes/second.</p> <p>This field is for HiperSockets.</p>
MESSAGE RATE	<p><b>PART</b> Rate of messages sent by this partition.</p> <p><b>TOTAL</b> Rate of messages sent by the CPC.</p> <p>This field is for HiperSockets.</p>
MESSAGE SIZE	<p><b>PART</b> Average size of messages sent by this partition.</p> <p><b>TOTAL</b> Average size of messages sent by the CPC.</p> <p>This field is for HiperSockets.</p>
SEND FAIL PART	<p>Rate of messages (sent by this partition) that failed.</p> <p>This field is for HiperSockets.</p>
RECEIVE FAIL	<p><b>PART</b> Rate of messages (received by this partition) that failed due to unavailable buffers. The value could indicate, that more receive buffers are required.</p> <p><b>TOTAL</b> Rate of messages (received by the CPC) that failed due to unavailable buffers.</p> <p>This field is for HiperSockets.</p>

## Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following table shows the overview condition names for the Overview report.

Table 142. Overview names in the Channel Path Activity report

Field Heading or Meaning	Overview Name
Use the following overview condition if CPMF is not available or for CPMF compatibility mode:	
TOTAL UTILIZATION (%)	CHPBSY, CHGPBSY
Use the following overview conditions for CPMF extended mode:	

Table 142. Overview names in the Channel Path Activity report (continued)

Field Heading or Meaning	Overview Name
UTILIZATION (%) PART	CHLBSY, CHGLBSY
UTILIZATION (%) TOTAL	CHTBSY, CHGTBSY
UTILIZATION (%) BUS	CHBTOT, CHGBTOT
PART READ RATE	CHLREAD, CHGLREAD
TOTAL READ RATE	CHTREAD, CHGTREAD
PART WRITE RATE	CHLWRITE, CHGLWRITE
TOTAL WRITE RATE	CHTWRITE, CHGTWRITE
FICON OPERATIONS RATE	CHFRATE
FICON OPERATIONS ACTIVE	CHFACTV
FICON OPERATIONS DEFER	CHFDFER
ZHPF OPERATIONS RATE	CHFXRATE
ZHPF OPERATIONS ACTIVE	CHFXACTV
ZHPF OPERATIONS DEFER	CHFXDFER
MESSAGE RATE PART	CHLMSGST
MESSAGE RATE TOTAL	CHTMSGST
MESSAGE SIZE PART	CHLMSGSZ
MESSAGE SIZE TOTAL	CHTMSGSZ
SEND FAIL PART	CHLMSGF
RECEIVE FAIL PART	CHLRECF
RECEIVE FAIL TOTAL	CHTRECF

---

## CPU - CPU Activity report

The report is divided into the following sections:

### CPU Activity

Provides information on the active processors. For further information, see “CPU Activity” on page 341.

### System Address Space and Work Unit Analysis

Provides overall information about address spaces and running or waiting work units. For further information, see “System Address Space and Work Unit Analysis” on page 343.

### Blocked Workload Analysis

Provides information about blocked workloads. For further information, see “Blocked Workload Analysis” on page 344.

### Partition Data Report

If the z/OS system is running in a PR/SM environment in LPAR mode, this section provides data about all configured partitions. If the z/OS system is running as guest under z/VM, and the Monitor I data gatherer option VMGUEST has been set, this section provides data about the z/OS guest system. Otherwise, this section is not available.

This section is described in “Using the information in the Partition Data Report” on page 350.

### LPAR Cluster Report

Provides data about each LPAR cluster. This section is described in “Using the information in the LPAR Cluster Report” on page 356.

**Group Capacity Report**

Provides data about the capacity limit of each defined capacity group and about the MSU consumption and actual capping of each partition within these groups. This section is described in “Using the information in the Group Capacity Report” on page 357.

**Note:** The *LPAR Cluster Report* and *Group Capacity Report* sections are not available if the system is running in a z/VM guest environment.

**How to request this report**

Monitor I gathers data for this report automatically. If you want to suppress gathering, you need to specify NOCPU.

To produce this report, specify  
REPORTS(CPU)

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

**Example URL for the DDS API:**

<http://ddshost:8803/gpm/rmfpp.xml?reports=CPU>

**Contents of the report**

The contents of the CPU Activity Report includes the following parts:

- “CPU Activity”
- “System Address Space and Work Unit Analysis” on page 343
- “Blocked Workload Analysis” on page 344

**CPU Activity**

The CPU Activity section reports on logical core and logical processor activity. For each processor, the report provides a set of calculations that are provided at a particular granularity that depends on whether multithreading is disabled (LOADxx PROCVIEW CPU parameter is in effect) or enabled (LOADxx PROCVIEW CORE parameter is in effect).

If multithreading is disabled for a processor type, all calculations are at logical processor granularity.

If multithreading is enabled for a processor type, some calculations are provided at logical core granularity and some are provided at logical processor (thread) granularity. The CPU Activity section displays exactly one report line per thread showing all calculations at logical processor granularity. Those calculations that are provided at core granularity are only shown in the same report line that shows the core id in the CPU NUM field and which is representing the first thread of a core.

The following calculations are on a per logical processor basis when multithreading is disabled and on a per logical core basis when multithreading is enabled:

- Percentage of the interval time the processor was online
- LPAR view of the processor utilization (LPAR Busy time percentage)
- Percentage of a physical processor the logical processor is entitled to use
- Multithreading core productivity (only reported when multithreading is enabled)
- Multithreading core utilization (only reported when multithreading is enabled)

The following calculations are on a per logical processor basis regardless whether multithreading is enabled or disabled:

- MVS view of the processor utilization (MVS Busy time percentage)
- Percentage of the online time the processor was parked (in HiperDispatch mode only)
- I/O interrupts rate (general purpose processors only)
- Percentage of I/O interrupts handled by the I/O supervisor without re-enabling (general purpose processors only)

If RMF is running as a guest under z/VM<sup>®</sup> and Monitor I Session option NOVMGUEST is active, it only reports the MVS busy time percentage. If you want to measure partition utilization (as well as the individual CPU utilization of the single guests, namely LPAR busy time percentage), you need to use a z/VM monitor. Performance analysts need both views of CPU utilization. The MVS view is a direct indicator to see a CPU bottleneck, while the LPAR view is important with respect to capacity aspects.

The **LPAR view** of the CPU utilization takes the different states that are possible into account:

- WAIT state
- NON WAIT state being dispatched by PR/SM
- NON WAIT state not being dispatched by PR/SM
- WAIT state being dispatched when the LPAR has dedicated processors

The LPAR Busy time is calculated depending on the status of the logical processor:  
**Dedicated and LOADxx PROCVIEW CPU is in effect or hardware does not support multithreading**

$$\text{CPU time} = \text{Online time} - \text{Wait time}$$

**Dedicated and LOADxx PROCVIEW CORE is in effect on hardware that supports multithreading**

$$\text{CPU time} = \text{MT Core LPAR Busy time}$$

**Wait completion = YES (requires multithreading disabled)**

$$\text{CPU time} = \text{Dispatch time} - \text{Wait time}$$

**Wait completion = NO**

$$\text{CPU time} = \text{Dispatch time}$$

The LPAR view of CPU utilization is calculated as:

$$\text{LPAR Busy Time(\%)} = \frac{\text{CPU time}}{\text{Online time}} * 100$$

The MVS view of the CPU utilization considers the following states:

- CPU wait state
- CPU busy state (which means NON WAIT state)

In HiperDispatch mode, logical processors can be parked and are not dispatched by z/OS. The MVS BUSY fields in the RMF report reflect the effective used capacity for the logical processors and the entire logical partition. The values are based on the difference between online time and MVS wait time to provide an operating system perspective of busy time. Parked processors in HiperDispatch mode generally reflect unavailable capacity at high physical processor utilizations. The formula for MVS Busy has been changed with HiperDispatch mode to exclude the parked time to show how busy the logical processor was when not parked.

**HiperDispatch = NO**

$$\text{Time range} = \text{Online time}$$



**HiperDispatch = YES**

Time range = Online time - Parked time

**Note:** In HiperDispatch mode, the Total/Average MVS BUSY TIME % does not consider parked processors. Therefore, do not use Total/Average LPAR BUSY TIME % and Total/Average MVS BUSY TIME % to calculate the MVS to LPAR busy ratio.

The **MVS view** of CPU utilization is:

$$\text{MVS Busy Time(\%)} = \frac{\text{Time range} - \text{Wait time}}{\text{Time range}} * 100$$

If multithreading is enabled for at least one processor type, you can use the multithreading core productivity and multithreading core utilization metrics to determine the effectiveness of the configured logical cores.

When the multithreading core productivity (MT % PROD) equals 100% in multithreading mode, all threads on the core are executing work and all core resources are being used. If MT % PROD is less than 100%, the core resources were dispatched to physical hardware but one or more threads on a core were in a wait because they had no work to run.

If multithreading is enabled, the available core capacity can be calculated using the multithreading core utilization and LOG PROC SHARE %:

Available Core Capacity = LOG PROC SHARE % - MT % UTIL

**System Address Space and Work Unit Analysis**

The *System Address Space and Work Unit Analysis* section of the CPU activity report provides overall address space and work unit information and also provides the minimum, maximum, and average numbers of running or ready to run work units.

The data in this section analyzes the following types of address spaces:

- In storage and ready to execute
- In storage
- Out of storage and ready to execute
- Out of storage and waiting to execute
- Logically out of storage and ready to execute
- Logically out of storage and waiting to execute.

Data is also presented on the number of address spaces used by batch users, started tasks (STC), TSO/E users, APPC/MVS transaction schedulers (ASCH), and z/OS UNIX (OMVS). Examining this data can indicate when a backlog of address spaces are waiting to use the processor.

The work unit statistics (MIN, MAX, AVG) are provided per processor type, that is, per standard CPs, zAAPs, and zIIPs.

The graphical and numeric presentation of the In-Ready work unit queue distribution provides a detailed view on how many work units are running or waiting for a processor. The distribution does not distinguish between the processor types (CPs, zAAPs, and zIIPs).

### **Blocked Workload Analysis**

If the CPU utilization of a system is at 100%, workloads with low importance (low dispatch priority) might not get dispatched anymore. This could cause problems if the work holds a resource and by that holds up more important workloads. Therefore, any address space or enclave which has ready-to-run work units (TCBs or SRBs), but does not get CPU service within a certain time interval due to its low dispatch priority, will be temporarily promoted by WLM to a higher dispatch priority. This helps to complete low priority work in a finite time period, without permanently delaying high priority work.

The *Blocked Workload Analysis* section lists the number of dispatchable work units that are considered to be blocked and eligible for priority promotion. This section also displays the OPT parameters which define the workload promotion. It also displays the average exploitation of the defined promotion rate during the measurement interval. This information helps you to adjust these OPT parameters. To assess the amount of workload still being blocked, the average and peak number of address spaces and enclaves found blocked and waiting for promotion is also listed.

### **Using the information in the CPU Activity report**

High LPAR/MVS BUSY TIME PERC values could indicate contention for CPU. To check this, add the N+1, ... N+150 percentages in the DISTRIBUTION OF IN-READY WORK UNIT QUEUE (where N is the number of online processors). This sum is the percentage of time when at least one task could not be dispatched. A value higher than 60% implies contention for CPU.

Low LPAR/MVS BUSY TIME PERC values can indicate that other bottlenecks in the system are preventing work from being processed.

An OUT READY average value of more than 1 could reflect processor storage constraints.



PP - CPU

Table 143. Fields in the CPU Activity Report (continued)

Field Heading	Meaning
HIPERDISPATCH	<p>HiperDispatch mode:  <b>YES</b> Active  <b>NO</b> Not active  <b>N/A</b> Not supported by the hardware</p> <p>If the mode changed during the reporting interval, an '*' is appended (for example: NO* indicates a switch from YES to NO).</p>
CPU NUM/TYPE	The logical core identification and the processor type.
TIME % ONLINE	The percentage of time the logical core was online.
TIME % LPAR BUSY	<p>The percentage of the online time that the logical core was dispatched by LPAR.</p> <ul style="list-style-type: none"> <li>For a dedicated partition:                      When LOADxx PROCVIEW CORE is in effect on hardware that supports multithreading:  <math display="block">\text{LPAR BUSY TIME \%} = \frac{\text{MT Core LPAR Busy Time}}{\text{Online Time}} * 100</math> </li> <li>Otherwise:  <math display="block">\text{LPAR BUSY TIME \%} = \frac{\text{Online Time} - \text{Wait Time}}{\text{Online Time}} * 100</math> </li> <li>For a non-dedicated partition when Wait Completion is NO:  <math display="block">\text{LPAR BUSY TIME \%} = \frac{\text{Partition Dispatch Time}}{\text{Online Time}} * 100</math> <p>The partition dispatch time is the elapsed time that PR/SM dispatched this logical core during the interval.</p> </li> <li>For a non-dedicated partition when Wait Completion is YES:  <math display="block">\text{LPAR BUSY TIME \%} = \frac{\text{Partition Dispatch Time} - \text{Wait Time}}{\text{Online Time}} * 100</math> </li> </ul>
TIME % MVS BUSY	<p>The percentage of the online time that the logical processor was busy.</p> $\text{MVS BUSY TIME \%} = \frac{\text{Online Time} - (\text{Wait Time} + \text{Parked Time})}{\text{Online Time} - \text{Parked Time}} * 100$ <p>The MVS view of CPU time is not meaningful if the logical processor is parked during the entire reporting interval. In this case, '----' is shown.</p>
TIME % PARKED	The percentage of time that the logical processor was parked. In HiperDispatch mode, processors with a low amount of physical processor share can be parked. That is, they are not dispatched by z/OS and do not attempt to run work. Without HiperDispatch, processors are not parked and '----' is shown.
MT % PROD	<p>The percentage of the maximum core capacity that was used in the reporting interval while the logical core was dispatched to physical hardware.</p> <p>When MT % PROD equals 100% and the LOADxx PROCVIEW CORE parameter is in effect, all threads on the core are executing work and all core resources are being used. If MT % PROD is less than 100%, the core resources were dispatched to physical hardware but one or more threads on a logical core were in a wait because they had no work to run.</p> <p>If a core was reconfigured offline/online during the reporting interval, no multithreading core productivity is calculated and '-----' is shown. If the LOADxx PROCVIEW CPU parameter is in effect, this field is not displayed.</p>

Table 143. Fields in the CPU Activity Report (continued)

Field Heading	Meaning
MT % UTIL	<p>The percentage of the maximum core capacity that was used in the reporting interval.</p> $\text{MT \% UTIL} = \text{MT Core Productivity} * \text{TIME \% LPAR BUSY}$ <p>If a core was reconfigured offline/online during the reporting interval, no multithreading core utilization is calculated and '-----' is shown. If the LOADxx PROCVIEW CPU parameter is in effect, this field is not displayed.</p>
LOG PROC SHARE %	<p>Percentage of the physical processor that the logical processor is entitled to use.</p> <p>Without HiperDispatch, the processing weight is equally divided between the online logical processors.</p> <p>In HiperDispatch mode, logical processors have a high, medium or low share of the physical processor. The share percentage is the average value for the reporting interval, whereas HIGH, MED or LOW indicates the HiperDispatch priority at the end of the reporting interval. When the priority changed during the interval, an '*' is appended.</p> <p>N/A is displayed if the HiperDispatch priority is not indicated by the hardware at the end of the reporting interval.</p>
I/O INTERRUPTS RATE	<p>The total rate per second that this processor handled I/O interrupts. The rate reflects the processing for the entire interval. This might include periods of time when the SRM enabled or disabled this processor for I/O interrupts. The rate includes interrupts handled by the second level interrupt handler (SLIH), as well as those handled by the Test Pending Interrupt (TPI) instruction.</p> $\text{RATE} = \frac{\text{SLIH} + \text{TPI}}{\text{INT}}$ <p><b>SLIH</b>    Interrupts that the second level interrupt handler handled  <b>TPI</b>      Interrupts that the Test Pending Interrupt instruction handled  <b>INT</b>      Interval time (seconds)</p>
I/O INTERRUPTS % VIA TPI	<p>The percentage of the total interrupts for this processor during the RMF interval that are handled by the I/O supervisor without re-enabling.</p> $\% \text{ VIA TPI} = \frac{\text{TPI}}{\text{SLIH} + \text{TPI}} * 100$ <p><b>TPI</b>      Interrupts that the Test Pending Interrupt instruction handled  <b>SLIH</b>    Interrupts that the second level interrupt handler handled</p>
For the following three TOTAL/AVERAGE values, the logical processors that are parked during the entire interval are not considered in the calculation of the average TIME % MVS BUSY.	
TOTAL/AVERAGE (CP)	The average or total value for general purpose processors (standard CPs).
TOTAL/AVERAGE (zAAP)	The average value for zAAPs. Only visible if zAAPs are configured online.
TOTAL/AVERAGE (zIIP)	The average value for zIIPs. Only visible if zIIPs are configured online.
<b>Multi-Threading Analysis:</b> This information is only displayed when the LOADxx PROCVIEW CORE parameter is in effect. multithreading information is only shown for those processor types for which at least one logical core was configured online for the complete interval.	
CPU TYPE	Processor type CP, IIP, or AAP.
MODE	The multithreading mode of a processor type designates the number of active threads for each online logical core of this type. If MT MODE is greater than 1, multithreading becomes effective for this processor type.
MAX CF	<p>Multithreading maximum capacity factor for a processor type. The multithreading maximum capacity factor represents the ratio of the maximum amount of work that can be accomplished using all active threads to the amount of work that would have been accomplished within this reporting interval when multithreading was disabled.</p> <p>'-----' is shown when the multithreading maximum capacity factor cannot be calculated.</p>

Table 143. Fields in the CPU Activity Report (continued)

Field Heading	Meaning
CF	<p>Multithreading capacity factor for a processor type. The multithreading capacity factor represents the ratio of the amount of work that has been accomplished within this reporting interval to the amount of work that would have been accomplished with multithreading disabled.</p> <p>'-----' is shown when the multithreading capacity factor cannot be calculated.</p>
AVG TD	<p>Average thread density for a processor type. This value represents the average number of active threads for those cores that were dispatched to physical hardware.</p> <p>'-----' is shown when the average thread density cannot be calculated.</p>
<p><b>System Address Space and Work Unit Analysis:</b> contains information about the NUMBER OF ADDRESS SPACES categorized by the QUEUE TYPES, in which they have been waiting, and categorized by the ADDRESS SPACE TYPES. Furthermore, the MIN, MAX, and AVG numbers of work units are categorized by the CPU TYPES for which they have been dispatched (that is, for standard CPs, zAAPs and zIIPs). This section also shows how many work units have been waiting in the IN-READY queue (DISTRIBUTION OF IN-READY WORK UNIT QUEUE).</p>	
QUEUE TYPES	<p>Shows the number of address spaces that are waiting in the different queues. For each queue type, the MIN, MAX and AVG numbers of address spaces are displayed.</p> <p>The following queue types are analyzed:</p> <p><b>IN READY</b> Address spaces that are in central storage and ready to execute or currently in execution.</p> <p><b>IN</b> Address spaces that are in central storage (corresponds to SRM in queue). This count includes the IN READY count.</p> <p><b>OUT READY</b> Address spaces on the SRM out queue that are physically swapped out of central storage and ready to execute. <b>Note:</b> Some address spaces on the SRM out queue might represent those TSO/E users that the SRM intentionally delayed to meet an installation's response time objective. Because these address spaces do not represent a potential performance problem, they are not included in the value reported for OUT READY.</p> <p><b>OUT WAIT</b> Address spaces on the SRM wait queue that are physically swapped out of central storage and not ready to execute.</p> <p><b>LOGICAL OUT RDY</b> Address spaces on the SRM out queue that are physically in central storage but logically swapped out of central storage and ready to execute.</p> <p><b>LOGICAL OUT WAIT</b> Address spaces on the SRM wait queue that are physically in central storage but logically swapped out of central storage and not ready to execute.</p>
ADDRESS SPACE TYPES	<p>Shows the total number of address spaces detected during the report interval, categorized by address space types. For each address space type, the MIN, MAX and AVG numbers of active address spaces are displayed.</p> <p>The following ADDRESS SPACE TYPES are analyzed:</p> <p><b>BATCH</b> Address spaces used for batch jobs.  <b>STC</b> Address spaces used for started task controls.  <b>TSO</b> Address spaces used for TSO/E users.  <b>ASCH</b> APPC/MVS transaction scheduler (ASCH) address spaces.  <b>OMVS</b> Address spaces for z/OS UNIX System Services.</p>
DISTRIBUTION OF IN-READY WORK UNIT QUEUE	<p>The percentaged and graphical distribution of SRM samples when the number of work units on the IN-READY queue is within a certain range. The correlation is based on N, which is the number of online logical processors when the sample is taken. In HiperDispatch mode, N is the number of online logical processors that are not parked.</p> <p>For example, NUMBER OF WORK UNITS = N + 10 with a percentage of 4.3 (see Figure 186 on page 345) indicates that in 4.3 % of the samples ten work units were waiting for a logical processor.</p>

Table 143. Fields in the CPU Activity Report (continued)

Field Heading	Meaning
NUMBER OF WORK UNITS by CPU type	The minimum, maximum and average numbers of running and waiting work units categorized by CPU type (standard CPs, zAAPs and zIIPs).
<b>Blocked Workload Analysis:</b> provides information about blocked address spaces and enclaves.	
OPT PARAMETERS	<p>Lists the OPT parameters which define the workload promotion:</p> <p><b>BLWLTRPCT (%)</b> Specifies how much of the CPU capacity is to be used to promote blocked workloads.</p> <p>This parameter does not influence the amount of CPU service that a single blocked address space or enclave is given. Instead, this parameter influences how many different address spaces or enclaves can be promoted at the same point in time. If the value specified with this parameter is not large enough, blocked workloads might need to wait longer than the time interval defined by BLWLINTHD.</p> <p>This value is specified as a number between 0 and 200 where 200 accounts for 20.0%.</p> <p><b>BLWLINTHD</b> Specifies the threshold time interval in seconds for which a swapped-in address space or enclave must wait before being considered to be blocked and eligible for promotion.</p> <p>If the parameters have been changed during the reporting interval, the values are followed by an '*'.</p>
PROMOTE RATE	<p><b>DEFINED</b> Number of blocked dispatchable work units which may get promoted in their dispatching priority per second. This value is derived from OPT parameter BLWLTRPCT.</p> <p><b>USED (%)</b> The utilization of the defined promote rate during the reporting interval.</p>
WAITERS FOR PROMOTE	<p>Number of address spaces and enclaves found blocked according to OPT parameter BLWLINTHD:</p> <p><b>AVG</b> Average number found blocked during the report interval. <b>PEAK</b> Highest number found blocked during the report interval.</p>

## Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following table shows the overview condition names for the Overview report.

Table 144. Overview names in the CPU Activity Report

Field Heading or Meaning	Overview Name
CPC CAPACITY	NOMCAPAC, EFFCAPAC
ONLINE TIME PERC for general purpose processors	CONTPER
LPAR BUSY TIME PERC	CPUBSY (LPAR mode only)
LPAR BUSY TIME PERC for zAAPs	AAPBSY
LPAR BUSY TIME PERC for zIIPs	IIPBSY
MVS BUSY TIME PERC	MVSBSY, CPUBSY
MVS BUSY TIME PERC for zAAPs	AAPMBSY
MVS BUSY TIME PERC for zIIPs	IIPMBSY
TYPE (IN READY)	AVGIARDY

Table 144. Overview names in the CPU Activity Report (continued)

Field Heading or Meaning	Overview Name
TYPE (other)	MXBATCH, AVGBATCH, MXSTC, AVGSTC, MXTSO, AVGTSO, MXASCH, AVGASCH, MXOMVS, AVGOMVS, AVGOARDY, AVGUIN, AVGUOWT, AVGULRDY, AVGULWT
Number of general purpose processors online	NUMPROC
Number of zAAPs online	NUMAAP
Number of zIIPs online	NUMIIP
Percentage of the report interval during which at least <i>n</i> jobs could not be dispatched (with <i>n</i> =1,2,3,4,5,10,15,20,30,40,60,80)	OCPU1, OCPU2, OCPU3, OCPU4, OCPU5, OCPU10, OCPU15, OCPU20, OCPU30, OCPU40, OCPU60, OCPU80
Number of CPs/zAAPs/zIIPs with high/medium/low HiperDispatch share for the partition	HDCPHIGH, HDAPHIGH, HDIPHIGH, HDCPMED, HDAPMED, HDIPMED, HDCPLow, HDAPLOW, HDIPLow
Percentage of time that the general purpose processor was parked	CPARKPER
Maximum number of in-ready work units for general purpose processors	MXWUCP
Maximum number of in-ready work units for zAAPs	MXWUAAAP
Maximum number of in-ready work units for zIIPs	MXWUIIP
Average number of in-ready work units for general purpose processors	AVGWUCP
Average number of in-ready work units for zAAPs	AVGWUAAAP
Average number of in-ready work units for zIIPs	AVGWUIIP
Percentage of the report interval during which at least <i>n</i> work units could not be dispatched (with <i>n</i> =1,2,3,4,5,10,15,20,30,40,60,80,100,120,150)	WCPU1, WCPU2, WCPU3, WCPU4, WCPU5, WCPU10, WCPU15, WCPU20, WCPU30, WCPU40, WCPU60, WCPU80, WCPU100, WCPU120, WCPU150
Percent multithreading core productivity for general purpose processors	MTPROD
Percent multithreading core productivity for zIIPs	IIPPROD
Percent multithreading core utilization for general purpose processors	MTUTIL
Percent multithreading core utilization for zIIPs	IIPUTIL

## Using the information in the Partition Data Report

When RMF is running in a Processor Resource/Systems Manager (PR/SM) environment in LPAR mode, the *Partition Data Report* section of the *CPU Activity* report provides data about all configured partitions active at the end of the reporting interval, independent of the operating system running in each partition.

When RMF is running on a z/OS guest in a z/VM guest environment, and the Monitor I data gatherer option VMGUEST has been set when the SMF record was collected, then the report section provides data about the z/OS guest system. If you want information about another z/OS guest system, you can run RMF separately on that guest system.

The report contains the following information:

- Header information
- Partition data
- Logical partition processor data



- Average processor utilization percentages

The header information gives an overview of the LPAR mode characteristics:

- MVS partition name
- Image capacity — information related to software pricing
- Number of configured partitions
- Number of physical processors in total and per type
- Wait completion indicator
- Dispatch interval
- If a group of LPARs on the same CEC is managed towards a combined capacity limit, the name of the group and the common capacity limit is displayed
- Capping information

The section PARTITION DATA is grouped by general purpose and special purpose processor types and provides the following information:

- Name
- Status
- Weighting share of resources
- Defined and consumed service units
- Capping information

The section LOGICAL PARTITION PROCESSOR DATA provides the following information about the partition's processors:

- Number and type of processors assigned to this partition
- The partition's effective dispatch time
- The partition's total dispatch time

The section AVERAGE PROCESSOR UTILIZATION PERCENTAGES provides the following information about the partition's processors:

- Logical constraint percentages. If multithreading is enabled, the percentages shown for logical processor resources can be applied to logical core resources.
  - The partition's average effective utilization of the logical processor resource
  - The partition's average total utilization of the logical processor resource
- Physical constraint percentages. If multithreading is enabled, the percentages shown for physical processor resources can be applied to physical core resources.
  - The average LPAR Management utilization of the physical processor resource on behalf of the partition
  - The partition's average effective utilization of the physical processor resource
  - The partition's average total utilization of the physical processor resource

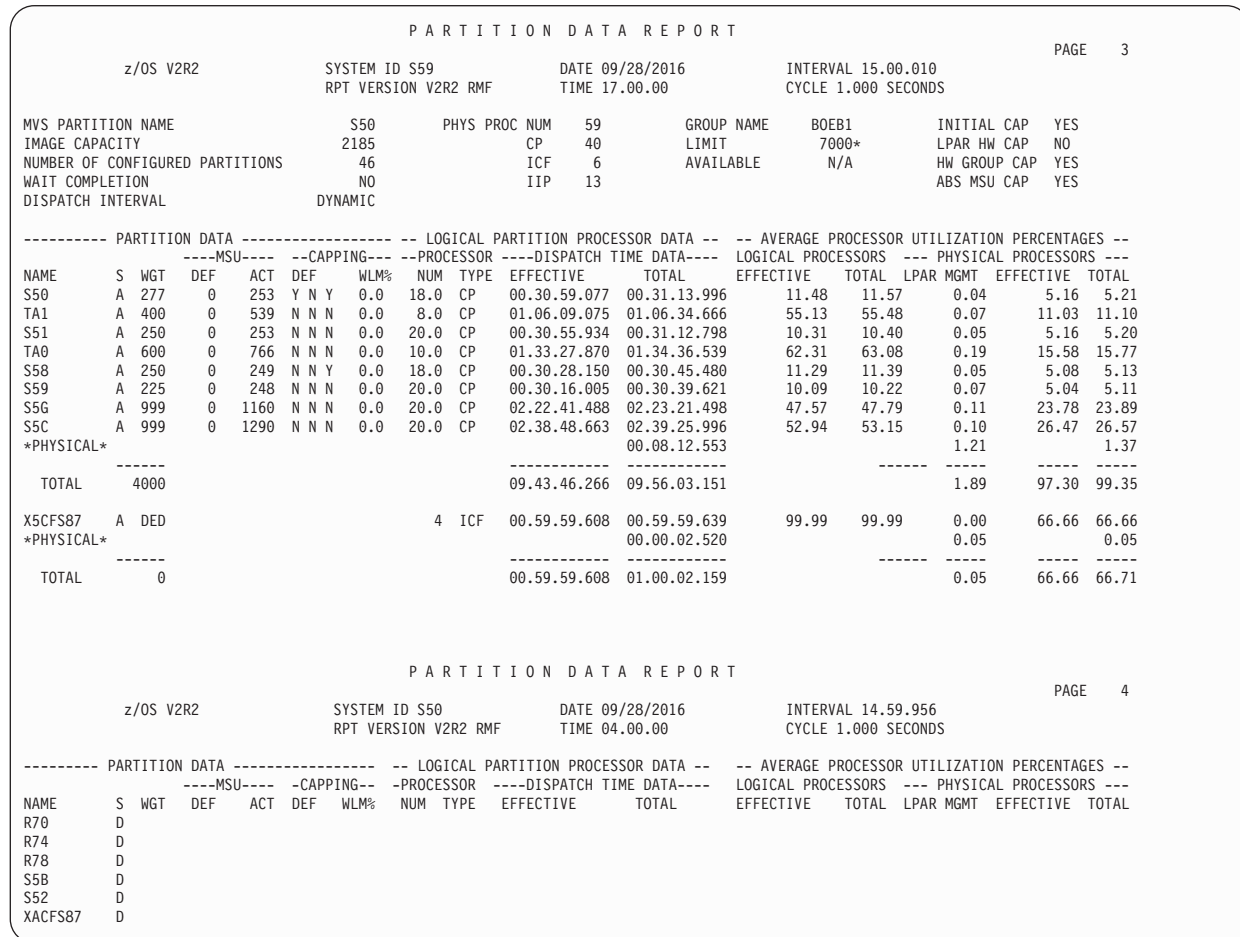


Figure 187. CPU Activity Report - Partition Data Report

**Note:** An asterisk (\*) next to any value indicates a change to this value during the measurement interval.

Table 145. Fields in the Partition Data Report

Field Heading	Meaning
<b>Header Information</b>	
MVS PARTITION NAME	The partition running the z/OS system which requested this report. <i>VMSystem</i> is displayed if the report was requested by a z/OS system running in a z/VM guest environment.
IMAGE CAPACITY	CPU capacity available to the MVS image measured in MSUs (millions of service units) per hour. The field is calculated as minimum of the following capacities: <ul style="list-style-type: none"> <li>the capacity based on the partition's logical CP configuration (includes online and standby [can be configured online] CPs)</li> <li>the defined capacity limit of the partition, if available (image softcap)</li> <li>the capacity limit of the related WLM capacity group, if the partition belongs to a capacity group</li> <li>the absolute physical hardware capping limit</li> <li>the capacity based on the hardware group capping limit.</li> </ul> For z/OS systems running as z/VM guests, the field displays the CPU capacity available to the z/VM partition.
NUMBER OF CONFIGURED PARTITIONS	The total number of activated and deactivated configured partitions. This number does not include the partition reported by the name *PHYSICAL*.

Table 145. Fields in the Partition Data Report (continued)

Field Heading	Meaning
PHYS PROC NUM or VM PROC NUM	Number of physical processors in total and per processor type. Starting with IBM System z9 processors, IFLs and zAAPs are reported separately, and no longer as ICFs.  If the data is reported for a z/OS system running as a z/VM guest, the field presents the number of processors that are assigned to the z/VM partition.
WAIT COMPLETION	The wait completion option of the partition: YES, NO, or MIX <b>YES</b> Implies that the processors assigned to each partition will remain dispatched to the partition until the time slice period has ended. <b>NO</b> Implies that the processors assigned to each partition become available to other partitions when the work for this partition is completed. The time slice period might or might not have ended. This field has no meaning for a dedicated partition. <b>MIX</b> Indicates that a mix of YES and NO is used for processors in the partition where RMF is running.
DISPATCH INTERVAL	Time (in milliseconds) a processor can be used when dispatched. This value is specified on the Logical Partition Control (LPCTO) frame on the Processor Controller Element (PCE).  DYNAMIC appears in this field if a value is not specified and implies that the length of time a processor is assigned to a partition is dynamically allocated. See <i>PR/SM Planning Guide</i> for more information.
GROUP NAME	Name of the capacity group to which the partition belongs, if it is managed towards a common group capacity limit.
LIMIT	Capacity limit (in MSUs) defined for the partition's capacity group.  An '*' following the limit value indicates that this partition started to be a member of this capacity group less than four hours ago. This partition will have a different view of unused group capacity and, therefore, may cap differently than existing group members.
AVAILABLE	Long-term average of CPU service in MSUs/h which would be allowed by the limit of the capacity group but is not used by its members. If the value is negative, the group is capped.
INITIAL CAP	Indicates whether the operator has set 'Initial Capping ON' in the logical partition controls of the Hardware Management Console (HMC) for the partition.
LPAR HW CAP	Indicates whether an absolute physical hardware capping limit has been defined in the logical partition controls of the HMC for any processor type of the partition.
HW GROUP CAP	Indicates whether an absolute hardware group capping limit has been defined in the logical partition group controls of the HMC for any processor type of the partition.
ABS MSU CAP	Indicates whether the ABSMSUCAPPING parameter has been set in the active IEAOPTxx parmlib member for the partition.
<b>Partition Data</b>	
NAME	The name that identifies a partition D <b>Note:</b> 1. The partition identified by the name *PHYSICAL* is not a configured partition. Data reported for *PHYSICAL* is shown only in columns DISPATCH TIME DATA - TOTAL, PHYSICAL PROCESSORS - LPAR MGMT, and PHYSICAL PROCESSORS - TOTAL. 2. When data about a z/OS system in a z/VM guest environment is reported, the *VMSystem* line reports the time used by z/VM itself.
S	The current status of the partition: <b>A</b> Activated <b>D</b> Deactivated. The LPAR is configured but there are currently no logical CPUs online for this partition.  If a partition is deactivated, the rest of the report line is blank.
WGT	Either the partition's current weighting of the shared processor resources or one of the following indicators: <b>DED</b> Indicates that the partition is dedicated. <b>DMX</b> Indicates that a mix of dedicated and non-dedicated processors is used in this partition. <b>WMX</b> Indicates that different share values are assigned to processors used in this partition.

Table 145. Fields in the Partition Data Report (continued)

Field Heading	Meaning
MSU	Shows capacity information for a partition in terms of MSUs per hour. This information is shown for general purpose processors only. <b>DEF</b> Defined capacity limit of the partition.  For the partition which is gathering the RMF data, this value is equal to the image capacity which is shown in the header of the report. <b>ACT</b> Actual consumption.
CAPPING	Shows capping information for a partition. <b>DEF</b> The hardware capping option of the partition. Each DEF value is a three position character string denoting which hardware capping mechanisms have or have not been applied in the logical partition controls of the HMC for the partition. The values in the first, second and third position of the string are either Y (Yes) or N (No) and have the following meaning: The first character (Y or N) indicates whether "Initial Capping ON" has been set. The second character (Y or N) indicates whether an absolute physical hardware capping limit (maximal number of CPUs) has been defined. The third character (Y or N) indicates whether an absolute hardware group capping limit (maximal number of CPUs) has been defined.  For example, a DEF value of "Y N N" indicates that "Initial Capping ON" has been applied to this partition, but the other two options have not.  The information provided in this field is useful only for logical partitions with shared processors. <b>WLM%</b> Percentage of time when WLM capped the partition. This information is shown only for general purpose processors.
<b>Logical Partition Processor Data</b>	
PROCESSOR NUM TYPE	The number of physical processors assigned to this partition and its processor type.
EFFECTIVE DISPATCH TIME	The sum of all processors' effective dispatch times for this partition during the measurement interval; expressed in the form HH.MM.SS.TTT. Partition effective dispatch time is the time, excluding LPAR management time, that a processor was assigned to this partition during the measurement interval.
TOTAL DISPATCH TIME	The sum of all processors' dispatch times for this partition during the measurement interval, including LPAR management time.  It is possible that the total dispatch time is smaller than the effective dispatch time. This situation occurs when partitions get "overruns" in their dispatch intervals caused by machine delays. The most typical form of this is caused by an MVS partition trying to talk to a coupling facility but getting significant delays or time-outs. It is sometimes symptomatic of recovery problems on the machine.  For *PHYSICAL*, this value includes the time during which a physical CPU was busy, but the time could not be attributed to a specific logical partition. This time includes the time PR/SM was controlling the physical processor (LPAR management time), as well as any other time the processor was busy for any reason such as managing coupling facility traffic.
<b>Average Processor Utilization Percentages</b>	
The average utilization of logical processors is based on the total online time of all logical processors assigned to the partition. The average utilization of physical processors is based on the total interval time of all physical processors. <b>Note:</b> If the z/OS system is running as guest under z/VM, and the Monitor I data gatherer option VMGUEST is active, the physical processor utilization represents the logical processor utilization of the z/VM LPAR.	
LOGICAL PROCESSORS - EFFECTIVE	The average partition effective dispatch time percentage. $\frac{\text{Effective Dispatch Time}}{\sum \text{Online Times}} * 100$
LOGICAL PROCESSORS - TOTAL	The average partition total dispatch time percentage. $\frac{\text{Total Dispatch Time}}{\sum \text{Online Times}} * 100$

Table 145. Fields in the Partition Data Report (continued)

Field Heading	Meaning
PHYSICAL PROCESSORS - LPAR MGMT	<p>The average LPAR management time on behalf of the partition reported as a percentage of the measurement interval.</p> $\frac{\text{Total Dispatch Time} - \text{Effective Dispatch Time}}{\# \text{ Physical Processors} * \text{Interval Time}} * 100$ <p>If the total dispatch time is smaller than the effective dispatch time, **** is shown in this column.</p> <p>The calculation for the *PHYSICAL* partition is:</p> $\frac{\text{Total Time} * \text{PHYSICAL} *}{\# \text{ Physical Processors} * \text{Interval Time}} * 100$ <p>Time *PHYSICAL* is the time that could not be attributed to a specific logical partition, but was used by PR/SM to control the physical processor (LPAR management time).  <b>Note:</b> # Physical Processors denotes the number of physical processors of a certain processor group, either general purpose processors or special purpose processors.</p>
PHYSICAL PROCESSORS - EFFECTIVE	<p>The effective utilization of the physical processor resource by the partition.</p> $\frac{\text{Effective Dispatch Time}}{\# \text{ Physical Processors} * \text{Interval Time}} * 100$
PHYSICAL PROCESSORS - TOTAL	<p>The total utilization of the physical processor resource by the partition.</p> $\frac{\text{Total Dispatch Time}}{\# \text{ Physical Processors} * \text{Interval Time}} * 100$
TOTAL	<p>The sum of the current weightings of the shared processor resources.</p> <p>The total amount of time the physical processor resource was assigned to a configured partition and to partition *PHYSICAL*.</p> <p>The sum of the AVERAGE PHYSICAL PROCESSOR UTILIZATION - LPAR MGMT field represents the total utilization of PR/SM physical processor resource by PR/SM.</p> <p>The sum of the AVERAGE PHYSICAL PROCESSOR UTILIZATION - EFFECTIVE field represents the total utilization of PR/SM physical processor resource by the operating systems running in each active partition.</p> <p>The sum of the AVERAGE PHYSICAL PROCESSOR UTILIZATION - TOTAL field represents the total utilization of the PR/SM physical processor resource by all configured partitions and by partition *PHYSICAL*.</p>

### Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following table shows the overview condition names for the Overview report.

Table 146. Overview names in the Partition Data Report

Field Heading or Meaning	Overview Name
PARTITION DATA - MSU DEF	LDEFMSU
PARTITION DATA - MSU ACT	LACTMSU
PARTITION DATA - CAPPING DEF	INICAP, LIMCPU
PARTITION DATA - CAPPING WLM%	WCAPPER
Available long-term average of CPU service (in MSUs/h)	GCMSUAV

## Using the information in the LPAR Cluster Report

Starting with zSeries 900 (z900) servers, the Workload Manager is extended to work with PR/SM to dynamically expand resources that are available across LPARs.

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.

### LPAR CPU Management

Based on workload resource demand, the Workload Manager is able to dynamically adjust the number of logical processors and the weight of a logical partition. This allows the system to distribute the CPU resource in an LPAR cluster to partitions where the CPU demand is high. An LPAR cluster is defined as the set of logical partitions in a single CEC that belong to the same parallel sysplex.

The dynamic adjustment of processor resources within the partitions is reflected in the LPAR Cluster report, which provides LPAR views as well as aggregated views on LPAR cluster level.

LPAR CLUSTER REPORT												PAGE 5		
z/OS V2R2			SYSTEM ID S59			DATE 09/28/2016			INTERVAL 15.00.010					
			RPT VERSION V2R2 RMF			TIME 07.45.00			CYCLE 1.000 SECONDS					
			WEIGHTING STATISTICS					PROCESSOR STATISTICS				STORAGE STATISTICS		
			DEFINED		ACTUAL			NUMBER		TOTAL%		CENTRAL	EXPANDED	
CLUSTER	PARTITION	SYSTEM	INIT	MIN	MAX	AVG	MIN %	MAX %	DEFINED	ACTUAL	LBUSY	PBUSY		
SVPLEXA	TA0	TA0	500	1	999	600	0.0	0.0	10	10.0	63.08	15.77	20480	N/A
	TA1	TA1	500	1	999	400	0.0	0.0	8	8.0	55.48	11.10	20480	N/A
TOTAL			1000						18	118.6	26.87		40960	N/A
SVPLEX5	S5C	S5C	500	1	999	999	0.0	100.0	20	20.0	53.15	26.57	61440	N/A
	S5G	S5G	500	1	999	999	0.0	100.0	20	20.0	47.79	23.89	64512	N/A
	S50	S50	500	1	999	277	0.0	0.0	18	18.0	11.57	5.21	20480	N/A
	S51	S51	500	100	999	250	0.0	0.0	20	20.0	10.40	5.20	51200	N/A
	S58	S58	500	100	999	250	0.0	0.0	18	18.0	11.39	5.13	20480	N/A
	S59	S59	500	100	999	225	0.0	0.0	20	20.0	10.22	5.11	51200	N/A
TOTAL			3000						116	144.5	71.11		269312	N/A

Figure 188. CPU Activity Report - LPAR Cluster Report

Table 147. Fields in the LPAR Cluster Report

Field Heading	Meaning
CLUSTER	This field identifies a sysplex name associated with the partition. All partitions that have the same cluster name are grouped together.
PARTITION	Name of the logical partition.
SYSTEM	z/OS system name.
<b>Weighting Statistics</b>	
All MIN/MAX-related fields are blank for partitions which are not under control of LPAR CPU management.	
DEFINED INIT / MIN / MAX	Defined initial, minimum, and maximum weighting of the shared processor resources.  A value of zero in fields MIN/MAX indicates that the partition is under control of LPAR CPU management, but no MIN/MAX values have been specified.
ACTUAL AVG	Actual weighting of the shared processor resources.  The contents of this field is equal to field WGT in the Partition Data report.
ACTUAL MIN% / MAX%	Percentage of time when the partition was within a bandwidth of 10% above the defined minimum weighting, or 10% below the defined maximum weighting.
<b>Processor Statistics</b>	

Table 147. Fields in the LPAR Cluster Report (continued)

Field Heading	Meaning
NUMBER DEFINED / ACTUAL	Defined and average actual number of general purpose processors assigned to this partition.  The actual number might be different from the defined number because of WLM goal achievement reasons.
TOTAL % LBUSY	Total dispatch time reported as a percentage of the logical processor online time: $\frac{\text{Partition Total Dispatch Time}}{\sum \text{Logical Processor Online Times}} * 100$  The contents of this field is equal to the LOGICAL PROCESSOR UTILIZATION - TOTAL column in the Partition Data report.
TOTAL % PBUSY	Total utilization of the physical processor resource by the partition: $\frac{\text{Partition Total Dispatch Time}}{\# \text{ Physical Processors} * \text{Interval Time}} * 100$  The contents of this field is equal to the PHYSICAL PROCESSOR UTILIZATION - TOTAL column in the Partition Data report.
<b>Storage Statistics</b>	
CENTRAL	The defined size of central storage (in MB) for this partition.
EXPANDED	The defined size of expanded storage (in MB) for this partition.

### Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following table shows the overview condition names for the Overview report.

Table 148. Overview names in the LPAR Cluster Report

Field Heading or Meaning	Overview Name
WEIGHTING - DEFINED INIT (Cluster)	WDEFC
WEIGHTING - DEFINED INIT (general purpose processors)	WDEFL
WEIGHTING - DEFINED MIN	WMINL
WEIGHTING - DEFINED MAX	WMAXL
WEIGHTING - ACTUAL AVG	WACTL
WEIGHTING - ACTUAL MIN%	WMIPL
WEIGHTING - ACTUAL MAX%	WMAPL
PROCESSOR - NUMBER DEFINED (Cluster)	NLDEFC
PROCESSOR - NUMBER DEFINED (Partition)	NLDEFL
PROCESSOR - NUMBER ACTUAL	NLACTL
PROCESSOR - TOTAL% LBUSY (Cluster)	LBUSYC
PROCESSOR - TOTAL% LBUSY (Partition)	LBUSYL
PROCESSOR - TOTAL% PBUSY (Cluster)	PBUSYC
PROCESSOR - TOTAL% PBUSY (general purpose processors)	PBUSYL

### Using the information in the Group Capacity Report

You can apply a defined capacity limit not only to one logical partition, but to a group of LPARs on the same CEC and manage this group considering the combined defined capacities of all members of the group.

With the group capacity limit, a third restriction to an LPAR is added. Even when an LPAR is not limited by its weight or its defined capacity, it can be limited by the group capacity. The minimum of the following limitations is applied to any partition:

1. Defined capacity
2. LPAR weights
3. Group capacity limit
4. Absolute physical hardware capping limit
5. Hardware group capping limit (refer to "Using the information in the Hardware Group Report" on page 359 for more information).

The Group Capacity Report monitors the available capacity of each defined capacity group and the MSU consumption and actual capping of these groups and of each partition within such a group. It helps you to exploit the flexibility to use as much CPU as needed for short periods of time until the 4 hour rolling MSU average exceeds the defined capacity limit for the whole group.

GROUP CAPACITY REPORT										
z/OS V2R2		SYSTEM ID TRX2			DATE 09/28/2016		INTERVAL 05.00.000		PAGE 3	
		RPT VERSION V2R2 RMF			TIME 21.55.00		CYCLE 1.000 SECONDS			
GROUP-CAPACITY NAME	PARTITION	SYSTEM	MSU DEF	MSU ACT	WGT	CAPPING DEF	CAPPING WLM%	CAPPING ACT%	ENTITLEMENT MINIMUM	ENTITLEMENT MAXIMUM
RMFGRP	400	TRX1	0	5	40	NO	0.0	0.0	29	400
		TRX2	0	11	500	NO	0.0	0.0	370	400
-----			TOTAL							
			16		540					

Figure 189. CPU Activity Report - Group Capacity Report

Table 149. Fields in the Group Capacity Report

Field Heading	Meaning
GROUP-CAPACITY NAME	Name of the capacity group.
GROUP-CAPACITY LIMIT	MSU limit defined for the capacity group.
PARTITION	Name of the logical partition.
SYSTEM	Name of the z/OS system.
MSU DEF	User defined capacity limit.
MSU ACT	Actual MSU consumption of this partition.
WGT	The partition's weighting of the shared processor resources which is used for WLM Group Capacity decisions. In case of hard capped partitions (see field CAPPING DEF), dashes (---) are displayed.
CAPPING DEF	The initial capping option of the partition: YES/NO  Initially capped partitions (also referred to as hard capped) are excluded from WLM group capacity management.
CAPPING WLM%	Percentage of time when WLM considers to cap the partition.  For more information on WLM capping, refer to topic <i>Workload Management and Workload License Charges</i> in <i>z/OS MVS Planning: Workload Management</i> .
CAPPING ACT%	Percentage of time when capping actually limited the usage of processor resources for the partition.
MINIMUM ENTITLEMENT	The minimum share of the MSU limit defined for the capacity group that the partition receives, even if all other partitions within the capacity group are running high workload. N/A is displayed for hard capped partitions.



Table 149. Fields in the Group Capacity Report (continued)

Field Heading	Meaning
MAXIMUM ENTITLEMENT	The maximum share of the MSU limit defined for the capacity group that a partition can receive if all other partitions within the capacity group are running without workload. N/A is displayed for hard capped partitions.

### Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following table shows the overview condition names for the Overview report.

Table 150. Overview names in the Group Capacity Report

Field Heading or Meaning	Overview Name
MSU - ACT	GCMSUACT
WGT	GCWEIGHT
MINIMUM ENTITLEMENT	MINENT
MAXIMUM ENTITLEMENT	MAXENT

## Using the information in the Hardware Group Report

With the hardware group capping limit, which can be set in the logical partition controls of the Hardware Management Console (HMC), another restriction to an LPAR is added. This limit enforces an absolute capping for each type of processor defined to every partition in the hardware group. Even when an LPAR is not limited by its weight, its defined capacity, or a group capacity limit, it can be limited by this hardware group capping value. The minimum of the following limitations is applied to any partition:

- Defined capacity
- LPAR weights
- Group capacity limit
- Absolute physical hardware capping limit
- Hardware group capping value

The Hardware Group Report displays the settings of the hardware groups and their partitions.

H A R D W A R E   G R O U P   R E P O R T							PAGE 7
z/OS V2R2		SYSTEM ID R74		DATE 09/28/2016		INTERVAL 04.59.999	
		RPT VERSION V2R2 RMF		TIME 06.21.00		CYCLE 1.000 SECONDS	
HW GROUP NAME	PARTITION	SYSTEM	----- HW GROUP LIMIT -----	CP	IIP	ICF	IFL
BOEB1	R74	R74	1.50	2.00	0.00	0.00	
	R75	R75					
BOEB2	S74	S74	1.00	2.85	0.00	0.00	
	S75	S75					

Figure 190. CPU Activity Report – Hardware Group Report

Table 151. Fields in the Hardware Group Report

Field Heading	Meaning
HW GROUP NAME	Name of the hardware group.
PARTITION	Name of the logical partition.
SYSTEM	Name of the z/OS system.

Table 151. Fields in the Hardware Group Report (continued)

Field Heading	Meaning
HW GROUP LIMIT	Absolute limit on partition usage of all CPs / zIIPs / ICFs / IFLs that are members of the same hardware group, in terms of numbers of CPUs. If the hardware group name or the limit changed during the reporting interval, an '*' is appended.

## Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. Table 152 shows the overview condition names for the Overview report.

Table 152. Overview names in the Hardware Group Report

Field Heading or Meaning	Overview Name
HW GROUP LIMIT - CP	HGCCP
HW GROUP LIMIT - IIP	HGCIIIP
HW GROUP LIMIT - ICF	HGCICF
HW GROUP LIMIT - IFL	HGCIFL

## Duration report

The following aspects have to be considered for a duration report.

The Postprocessor accumulates only similar SMF record types when the CPU activity report is requested. The first record determines the type of records to be accumulated. For example, if the first SMF record RMF encounters is a PR/SM SMF record, RMF accumulates only PR/SM SMF records. Non-PR/SM SMF records are skipped.

The following hierarchy exists when the Postprocessor encounters SMF records that RMF writes while running in different PR/SM environments:

1. If the SMF records contain different system identifiers then the records are processed separately.  
For example, if SMF records written in partition 1 and partition 2 have different system identifiers, and the SYSID control statement is not used, then two separate reports will be generated. If the SYSID control statement is specified, only the system identified in the statement will be reported. See *z/OS RMF User's Guide* for more information on this processing.
2. If the SMF records contain the same system identifiers, but differs in MVS partition name, number of configured partitions or partition name, the records are processed selectively.  
The first SMF record encountered defines the type of records to be accumulated. Any subsequent records that do not have the same characteristics are skipped.
3. If the SMF records contain the same system identifier but differs in number of physical processors, status, wait completion or number of logical processors, the records are processed as if they were from the same system. All records are processed. When a new value is encountered, an asterisk (\*) appears next to the changed value on the report. The new value is reported.

## CRYPTO - Crypto Hardware Activity report

The *Crypto Hardware Activity* report provides information about the activities in the various cryptographic hardware functions. Most cryptographic hardware functions can only be used through Cryptographic Support for z/OS (ICSF). ICSF is a standard component of z/OS. It provides cryptographic services in the z/OS environment. The report provides the following sections:

- **Cryptographic CCA coprocessors**

This section provides measurements about secure cryptographic functions executed on Common Cryptographic Architecture (CCA) coprocessors, use of secure encrypted key values, clear key and secure PKA operations, and special user cryptographic functions (using the user defined extension (UDX) capability of the card). For cryptographic CCA coprocessors, special attention should be given to RSA key-generation operations because these operations require a high amount of cryptographic processing capacity. Therefore, they are reported in addition to the total number of operations.

- **Cryptographic PKCS11 coprocessors**

This section provides measurements about secure public-key operations executed by cryptographic symmetric- and asymmetric-key functions.

- **Cryptographic accelerators**

This section provides measurements about public key operations (RSA cryptography operations) used with Secure Sockets Layer (SSL) or Transport Layer Security (TLS) protocols which are widely used to help secure e-business applications. The data for cryptographic accelerators is showing details for the two available algorithms, modular exponentiation (ME) and Chinese Remainder Theorem (CRT) for available key lengths (1024, 2048, and 4096 bit). This provides information how the usage of these algorithms affects the utilization of the accelerator.

- **ICSF Services**

The Crypto Hardware Activity report provides performance measurements on selected ICSF activities:

- Using the single and triple Data Encryption Standard (DES) and the Advanced Encryption Standard (AES) to encipher and decipher data.
- Generating and verifying message authentication codes (MAC). The MAC is a value calculated from the message according to a secret shared DES key or AES key and sent to the receiver together with the message. The receiver can recalculate the MAC and compare it with the MAC received. If the MAC values are identical, the message has not been altered during transmission.
- Using public hash functions. A hash is calculated from the transmission data according to a public key or function in cases where it is impossible to share a secret key. If the recalculated hash is identical to the one calculated before transmission, data integrity is ensured.
- Translating and verifying PINs.
- Digital signature generation and verification. A digital signature is created using the data to be signed and a private key, either using the ECC (Elliptic Curve Cryptography) or the RSA (Ron Rivest, Adi Shamir and Leonard Adleman) algorithm. The digitally signed data is sent to the receiver. The receiver can verify that the signature is valid, using the signer's public key.
- Format preserving encryption (FPE) to encipher, decipher, and translate data while preserving the original formatting of the data.

## How to request this report

Monitor I gathers data for this report automatically. If you want to suppress gathering, you need to specify NOCRYPTO.

To produce this report, specify  
REPORTS(CRYPTO)

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

## Example URL for the DDS API

<http://ddshost:8803/gpm/rmfpp.xml?reports=CRYPTO>

## Contents of the report

The data shown for cryptographic coprocessors and accelerators always reflects the total activity in your CPC, while the data shown for ICSF services is for the partition. If measurement data for one of the cryptographic features is not available, the corresponding report section is omitted.

CRYPTO HARDWARE ACTIVITY													PAGE	1
z/OS V2R2		SYSTEM ID TRX2			DATE 09/28/2016			INTERVAL 14.59.974						
		RPT VERSION V2R2 RMF			TIME 06.59.35			CYCLE 1.000 SECONDS						
----- CRYPTOGRAPHIC CCA COPROCESSOR -----														
TOTAL				KEY-GEN										
TYPE	ID	RATE	EXEC TIME	UTIL%	RATE									
CEX5C	4	0.19	46.20	0.9	0.00									
	5	0.17	47.90	0.8	0.00									
----- CRYPTOGRAPHIC PKCS11 COPROCESSOR -----														
TOTAL				OPERATIONS DETAILS										
TYPE	ID	RATE	EXEC TIME	UTIL%	FUNCTION	RATE	EXEC TIME	UTIL%						
CEX5P	6	0.09	35.95	0.3	ASYM FAST	0.00	0.000	0.0						
					ASYM GEN	0.00	0.000	0.0						
					ASYM SLOW	0.08	36.41	0.3						
					SYMM COMPLETE	0.00	0.000	0.0						
					SYMM PARTIAL	<0.01	0.896	0.0						
----- CRYPTOGRAPHIC ACCELERATOR -----														
TOTAL				ME-FORMAT RSA OPERATIONS				CRT-FORMAT RSA OPERATIONS						
TYPE	ID	RATE	EXEC TIME	UTIL%	KEY	RATE	EXEC TIME	UTIL%	RATE	EXEC TIME	UTIL%			
CEX5A	7	1574	54.66	100.0	1024	0.00	0.000	0.0	1332	35.45	100.0			
					2048	0.00	0.000	0.0	242.1	160.4	100.0			
					4096	0.00	0.000	0.0	0.00	0.000	0.0			
----- ICSF SERVICES -----														
ENCIPHER			DECIPHER			HASH			PIN					
SDES		TDES	AES	SDES	TDES	AES	SHA-1	SHA-256	SHA-512	TRANSLATE	VERIFY			
RATE	0.00	0.07	0.04	0.00	0.04	0.04	0.11	0.11	0.07	0.04	0.07			
SIZE	0.00	28.00	32.00	0.00	40.00	32.00	159.5	106.7	128.0					
MAC			AES MAC			RSA DSIG			ECC DSIG			FORMAT PRESERVING ENCRYPTION		
GENERATE		VERIFY	GENERATE		VERIFY	GENERATE		VERIFY	GENERATE		VERIFY	ENCIPHER	DECIPHER	TRANSLATE
RATE	0.39	0.39	0.13		0.13	0.26		0.26	0.13		0.13	0.29	0.29	0.29
SIZE	13.32	32.67	256.0		256.0							16.00	16.00	16.00

Figure 191. Crypto Hardware Activity Report

Table 153. Fields in the CRYPTO Hardware Activity Report

Field Heading	Meaning
Cryptographic CCA Coprocessor	

Table 153. Fields in the CRYPTO Hardware Activity Report (continued)

Field Heading	Meaning														
TYPE	Type that defines the cryptographic CCA coprocessor:  <table border="0"> <thead> <tr> <th>Type</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>PCICC</td> <td>PCI Cryptographic Coprocessor.</td> </tr> <tr> <td>PCIXCC</td> <td>PCI-X Cryptographic Coprocessor.</td> </tr> <tr> <td>CEX2C</td> <td>Crypto Express2 Coprocessor.</td> </tr> <tr> <td>CEX3C</td> <td>Crypto Express3 Coprocessor.</td> </tr> <tr> <td>CEX4C</td> <td>Crypto Express4S Coprocessor.</td> </tr> <tr> <td>CEX5C</td> <td>Crypto Express5S Coprocessor.</td> </tr> </tbody> </table>	Type	Meaning	PCICC	PCI Cryptographic Coprocessor.	PCIXCC	PCI-X Cryptographic Coprocessor.	CEX2C	Crypto Express2 Coprocessor.	CEX3C	Crypto Express3 Coprocessor.	CEX4C	Crypto Express4S Coprocessor.	CEX5C	Crypto Express5S Coprocessor.
Type	Meaning														
PCICC	PCI Cryptographic Coprocessor.														
PCIXCC	PCI-X Cryptographic Coprocessor.														
CEX2C	Crypto Express2 Coprocessor.														
CEX3C	Crypto Express3 Coprocessor.														
CEX4C	Crypto Express4S Coprocessor.														
CEX5C	Crypto Express5S Coprocessor.														
ID	Index that specifies the cryptographic CCA coprocessor.														
TOTAL	<b>RATE</b> Rate of all operations on this cryptographic coprocessor. <b>EXEC TIME</b> Average execution time (milliseconds) of all operations on this cryptographic coprocessor. <b>UTIL%</b> Total utilization percentage of this coprocessor.														
KEY-GEN RATE	Rate for RSA-key-generation operations.														
Cryptographic PKCS11 Coprocessor															
TYPE	Type that defines the cryptographic PKCS11 coprocessor:  <table border="0"> <thead> <tr> <th>Type</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>CEX4P</td> <td>Crypto Express4S PKCS11 Coprocessor.</td> </tr> <tr> <td>CEX5P</td> <td>Crypto Express5S PKCS11 Coprocessor.</td> </tr> </tbody> </table>	Type	Meaning	CEX4P	Crypto Express4S PKCS11 Coprocessor.	CEX5P	Crypto Express5S PKCS11 Coprocessor.								
Type	Meaning														
CEX4P	Crypto Express4S PKCS11 Coprocessor.														
CEX5P	Crypto Express5S PKCS11 Coprocessor.														
ID	Index that specifies the cryptographic PKCS11 coprocessor.														
TOTAL	Rate, average execution time (in milliseconds) and utilization percentage for all operations executed on this cryptographic PKCS11 coprocessor.														
OPERATIONS DETAILS	Rate, average execution time (in milliseconds) and utilization percentage for executed operations, categorized by cryptographic function type:  <table border="0"> <thead> <tr> <th>Type</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>ASYM FAST</td> <td>Fast asymmetric-key function.</td> </tr> <tr> <td>ASYM GEN</td> <td>Asymmetric-key generation function.</td> </tr> <tr> <td>ASYM SLOW</td> <td>Slow asymmetric-key function.</td> </tr> <tr> <td>SYMM COMPLETE</td> <td>Symmetric-key function that returns a complete or final result.</td> </tr> <tr> <td>SYMM PARTIAL</td> <td>Symmetric-key function that returns partial or incremental results.</td> </tr> </tbody> </table>	Type	Meaning	ASYM FAST	Fast asymmetric-key function.	ASYM GEN	Asymmetric-key generation function.	ASYM SLOW	Slow asymmetric-key function.	SYMM COMPLETE	Symmetric-key function that returns a complete or final result.	SYMM PARTIAL	Symmetric-key function that returns partial or incremental results.		
Type	Meaning														
ASYM FAST	Fast asymmetric-key function.														
ASYM GEN	Asymmetric-key generation function.														
ASYM SLOW	Slow asymmetric-key function.														
SYMM COMPLETE	Symmetric-key function that returns a complete or final result.														
SYMM PARTIAL	Symmetric-key function that returns partial or incremental results.														
Cryptographic Accelerator															
TYPE	Type that defines the cryptographic accelerator:  <table border="0"> <thead> <tr> <th>Type</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>PCICA</td> <td>PCI Cryptographic Accelerator.</td> </tr> <tr> <td>CEX2A</td> <td>Crypto Express2 Accelerator.</td> </tr> <tr> <td>CEX3A</td> <td>Crypto Express3 Accelerator.</td> </tr> <tr> <td>CEX4A</td> <td>Crypto Express4S Accelerator.</td> </tr> <tr> <td>CEX5A</td> <td>Crypto Express5S Accelerator.</td> </tr> </tbody> </table>	Type	Meaning	PCICA	PCI Cryptographic Accelerator.	CEX2A	Crypto Express2 Accelerator.	CEX3A	Crypto Express3 Accelerator.	CEX4A	Crypto Express4S Accelerator.	CEX5A	Crypto Express5S Accelerator.		
Type	Meaning														
PCICA	PCI Cryptographic Accelerator.														
CEX2A	Crypto Express2 Accelerator.														
CEX3A	Crypto Express3 Accelerator.														
CEX4A	Crypto Express4S Accelerator.														
CEX5A	Crypto Express5S Accelerator.														
ID	Index that specifies the cryptographic accelerator.														
TOTAL	Rate, average execution time (in milliseconds) and utilization for all operations on this cryptographic accelerator.														
KEY	Used RSA key length for each cryptographic accelerator and for each available RSA operation format (ME or CRT).														
ME-FORMAT RSA OPERATIONS	Rate, average execution time (in milliseconds) and utilization for all operations in ME-format (one line for each used RSA key length).														
CRT-FORMAT RSA OPERATIONS	Rate, average execution time (in milliseconds) and utilization for all operations in CRT-format (one line for each used RSA key length).														
ICSF services															

## PP - CRYPTO

Table 153. Fields in the CRYPTO Hardware Activity Report (continued)

Field Heading	Meaning
ENCRYPTION	<p><b>SDES RATE</b> Rate of encipher service calls using single DES.</p> <p><b>SDES SIZE</b> Average number of bytes per service call that have been enciphered using single DES.</p> <p><b>TDES RATE</b> Rate of encipher service calls using double and triple DES.</p> <p><b>TDES SIZE</b> Average number of bytes per service call that have been enciphered using double and triple DES.</p> <p><b>AES RATE</b> Rate of encipher service calls using AES.</p> <p><b>AES SIZE</b> Average number of bytes per service call that have been enciphered using AES.</p>
DECRYPTION	<p><b>SDES RATE</b> Rate of decipher service calls using single DES.</p> <p><b>SDES SIZE</b> Average number of bytes per service call that have been deciphered using single DES.</p> <p><b>TDES RATE</b> Rate of decipher service calls using double and triple DES.</p> <p><b>TDES SIZE</b> Average number of bytes per service call that have been deciphered using double and triple DES.</p> <p><b>AES RATE</b> Rate of decipher service calls using AES.</p> <p><b>AES SIZE</b> Average number of bytes per service call that have been deciphered using AES.</p> <p><b>Note:</b> For AES, only service calls and bytes sent to a coprocessor are reported.</p>
MAC	<p><b>GENERATE RATE</b> Rate of requests to generate MACs.</p> <p><b>GENERATE SIZE</b> Average number of bytes per request for which MAC has been generated.</p> <p><b>VERIFY RATE</b> Rate of requests to verify MACs.</p> <p><b>VERIFY SIZE</b> Average number of bytes per request for which MAC has been verified.</p>

Table 153. Fields in the CRYPTO Hardware Activity Report (continued)

Field Heading	Meaning
HASH	<p><b>SHA-1 RATE</b> Rate of requests to hash using the SHA-1 hash algorithm.</p> <p><b>SHA-1 SIZE</b> Average number of bytes to be hashed per request using the SHA-1 hash algorithm.</p> <p><b>SHA-256 RATE</b> Rate of requests to hash using the SHA-224 or the SHA-256 hash algorithm.</p> <p><b>SHA-256 SIZE</b> Average number of bytes to be hashed per request using the SHA-224 or the SHA-256 hash algorithm.</p> <p><b>SHA-512 RATE</b> Rate of requests to hash using the SHA-384 or the SHA-512 hash algorithm.</p> <p><b>SHA-512 SIZE</b> Average number of bytes to be hashed per request using the SHA-384 or the SHA-512 hash algorithm.</p>
PIN	<p><b>TRANSLATE RATE</b> Rate of requests to translate PIN.</p> <p><b>VERIFY RATE</b> Rate of requests to verify PIN.</p>
AES MAC	<p><b>GENERATE RATE</b> Rate of requests to generate AES MACs.</p> <p><b>GENERATE SIZE</b> Average number of bytes per request for which AES MACs have been generated.</p> <p><b>VERIFY RATE</b> Rate of requests to verify AES MACs.</p> <p><b>VERIFY SIZE</b> Average number of bytes per request for which AES MACs have been verified.</p>
RSA DSIG	<p><b>GENERATE RATE</b> Rate of requests to generate RSA digital signatures.</p> <p><b>VERIFY RATE</b> Rate of requests to verify RSA digital signatures.</p>
ECC DSIG	<p><b>GENERATE RATE</b> Rate of requests to generate ECC digital signatures.</p> <p><b>VERIFY RATE</b> Rate of requests to verify ECC digital signatures.</p>
FORMAT PRESERVING ENCRYPTION	<p><b>ENCIPHER RATE</b> Rate of requests to encipher data using FPE.</p> <p><b>ENCIPHER SIZE</b> Average number of bytes per request that have been enciphered using FPE.</p> <p><b>DECIPHER RATE</b> Rate of requests to decipher data using FPE.</p> <p><b>DECIPHER SIZE</b> Average number of bytes per request that have been deciphered using FPE.</p> <p><b>TRANSLATE RATE</b> Rate of requests to translate data using FPE.</p> <p><b>TRANSLATE SIZE</b> Average number of bytes per request that have been translated using FPE.</p>

## Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following table shows the overview condition names for the Overview report.

Table 154. Overview names in the CRYPTO Hardware Activity Report

Field Heading or Meaning	Overview Name
Cryptographic coprocessor TOTAL RATE	CRYCTR
Cryptographic coprocessor TOTAL EXEC TIME	CRYCTE
Cryptographic coprocessor TOTAL UTIL%	CRYCTU
Cryptographic coprocessor KEY-GENERATION RATE	CRYCKR
Cryptographic accelerator ME(1024) RATE	CRYAM1R
Cryptographic accelerator ME(1024) EXEC TIME	CRYAM1E
Cryptographic accelerator ME(1024) UTIL%	CRYAM1U
Cryptographic accelerator ME(2048) RATE	CRYAM2R
Cryptographic accelerator ME(2048) EXEC TIME	CRYAM2E
Cryptographic accelerator ME(2048) UTIL%	CRYAM2U
Cryptographic accelerator ME(4096) RATE	CRYAM3R
Cryptographic accelerator ME(4096) EXEC TIME	CRYAM3E
Cryptographic accelerator ME(4096) UTIL%	CRYAM3U
Cryptographic accelerator CRT(1024) RATE	CRYAC1R
Cryptographic accelerator CRT(1024) EXEC TIME	CRYAC1E
Cryptographic accelerator CRT(1024) UTIL%	CRYAC1U
Cryptographic accelerator CRT(2048) RATE	CRYAC2R
Cryptographic accelerator CRT(2048) EXEC TIME	CRYAC2E
Cryptographic accelerator CRT(2048) UTIL%	CRYAC2U
Cryptographic accelerator CRT(4096) RATE	CRYAC3R
Cryptographic accelerator CRT(4096) EXEC TIME	CRYAC3E
Cryptographic accelerator CRT(4096) UTIL%	CRYAC3U
ENCRYPTION SDES RATE	CRYISDER
ENCRYPTION SDES SIZE	CRYISDES
ENCRYPTION TDES RATE	CRYITDER
ENCRYPTION TDES SIZE	CRYITDES
ENCRYPTION AES RATE	CRYIAER
ENCRYPTION AES SIZE	CRYIAES
Average number of coprocessor calls for AES encipher services	CRYIAEO
DECRYPTION SDES RATE	CRYISDDR
DECRYPTION SDES SIZE	CRYISDDS
DECRYPTION TDES RATE	CRYITDDR
DECRYPTION TDES SIZE	CRYITDDS
DECRYPTION AES RATE	CRYIADR
DECRYPTION AES SIZE	CRYIADS
Average number of coprocessor calls for AES decipher services	CRYIADO
MAC GENERATE RATE	CRYIMGR
MAC GENERATE SIZE	CRYIMGS
MAC VERIFY RATE	CRYIMVR
MAC VERIFY SIZE	CRYIMVS
HASH SHA-1 RATE	CRYIHAR



Table 154. Overview names in the CRYPTO Hardware Activity Report (continued)

Field Heading or Meaning	Overview Name
HASH SHA-1 SIZE	CRYIHAS
HASH SHA-256 RATE	CRYIH2R
HASH SHA-256 SIZE	CRYIH2S
PIN TRANSLATE RATE	CRYIPTR
PIN VERIFY RATE	CRYIPVR
AES MAC GENERATE RATE	CRYIAMGR
AES MAC GENERATE SIZE	CRYIAMGS
AES MAC VERIFY RATE	CRYIAMVR
AES MAC VERIFY SIZE	CRYIAMVS
RSA DIGITAL SIGNATURE GENERATE RATE	CRYIDRGR
RSA DIGITAL SIGNATURE VERIFY RATE	CRYIDRVR
ECC DIGITAL SIGNATURE GENERATE RATE	CRYIDEGR
ECC DIGITAL SIGNATURE VERIFY RATE	CRYIDEVVR
FPE ENCIPHER RATE	CRYIFPER
FPE ENCIPHER SIZE	CRYIFPES
FPE DECIPHER RATE	CRYIFPDR
FPE DECIPHER SIZE	CRYIFPDS
FPE TRANSLATION RATE	CRYIFPTR
FPE TRANSLATION SIZE	CRYIFPTS
Cryptographic PKCS11 coprocessor TOTAL RATE	CRYPTR
Cryptographic PKCS11 coprocessor TOTAL UTIL%	CRYPTU
Cryptographic PKCS11 coprocessor TOTAL EXEC TIME	CRYPTE
Cryptographic PKCS11 coprocessor SLOW ASYM RATE	CRYP SAR
Cryptographic PKCS11 coprocessor SLOW ASYM UTIL%	CRYP SAU
Cryptographic PKCS11 coprocessor SLOW ASYM EXEC TIME	CRYP SAE
Cryptographic PKCS11 coprocessor FAST ASYM RATE	CRYP FAR
Cryptographic PKCS11 coprocessor FAST ASYM UTIL%	CRYP FAU
Cryptographic PKCS11 coprocessor FAST ASYM EXEC TIME	CRYP FAE
Cryptographic PKCS11 coprocessor SYMM PART RATE	CRYP SPR
Cryptographic PKCS11 coprocessor SYMM PART UTIL%	CRYP SPU
Cryptographic PKCS11 coprocessor SYMM PART EXEC TIME	CRYP SPE
Cryptographic PKCS11 coprocessor SYMM COMPL RATE	CRYP SCR
Cryptographic PKCS11 coprocessor SYMM COMPL UTIL%	CRYP SCU
Cryptographic PKCS11 coprocessor SYMM COMPL EXEC TIME	CRYP SCE
Cryptographic PKCS11 coprocessor ASYM GEN RATE	CRYP AGR
Cryptographic PKCS11 coprocessor ASYM GEN UTIL%	CRYP AGU
Cryptographic PKCS11 coprocessor ASYM GEN EXEC TIME	CRYP AGE

## DEVICE - Device Activity report

The Device Activity report provides information for all devices in one or more device classes (such as TAPE or DASD) or for those devices you specify on the DEVICE option.

When used with the Channel Path Activity and I/O Queuing Activity reports, this report can help you analyze the I/O activity at your installation and identify bottlenecks caused by a particular device.

## How to request this report

Monitor I gathers data for this report automatically with the default option `DEVICE(DASD)`. If you want to suppress gathering, you need to specify `NODEVICE`.

To produce this report, specify  
`REPORTS(DEVICE(type))`

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

*Example URL for the DDS API:*

`http://ddshost:8803/gpm/rmfpp.xml?reports=DEVICE(NMBR(2000,3FFF))&sysid=SYSA`

## Contents of the report

Each Device Activity report begins on a new page, and the class of devices included in the report is indicated by one of the following titles:

Report Title	What you specified
CHARACTER READER DEVICE ACTIVITY	DEVICE(CHRDR)
COMMUNICATION EQUIPMENT ACTIVITY	DEVICE(COMM)
DIRECT ACCESS DEVICE ACTIVITY	DEVICE(DASD)
GRAPHICS DEVICE ACTIVITY	DEVICE(GRAPH)
MAGNETIC TAPE DEVICE ACTIVITY	DEVICE(TAPE)
UNIT RECORD DEVICE ACTIVITY	DEVICE(UNITR)

The devices included in the report are grouped by logical control unit. The logical control unit provides a way to identifying a related set of devices. Moreover, this organization makes it easier to compare the data in the Device Activity report with the data in the I/O Queuing Activity report.

RMF follows the individual device data lines in each group with a summary line that provides a weighted average or total values for the entire logical control unit.

### NMBR and SG suboptions

If the `NMBR` and `SG` suboptions are specified together, the device report is divided into two parts.

- The first part of the report contains the devices specified by the `NMBR` suboption. The devices are sorted by LCU and device number.
- The second part of the report contains the devices specified for the `SG` suboption. The devices are sorted by storage group and by device numbers within the group.

**Note:** Some devices might be reported twice, since you can specify a device on the `NMBR` suboption that is part of a storage group specified on the `SG` suboption.

### Byte-multiplexor-channel-attached device

For any device attached to a byte multiplexor channel, the only measurement data available is the start subchannel (`SSCH`) + resume subchannel (`RSCH`) instruction count.

## Direct Access Device Activity report

For the DASD Activity report, the information can be sorted by LCU, or storage group, or both. When the storage group (SG) option is specified, the DASD Activity report is sorted by device number within each storage group. The storage group name that a volume is assigned to is always reported, even when the SG option was not selected. If a volume does not belong to a storage group, the STORAGE GROUP field for that volume is blank.

RMF follows the individual device data lines in each group with a summary line that provides average or total values for the entire storage group.

**Note:** When comparing I/O rates in the DASD Activity report and in the Cache Subsystem Activity report, you may see differences due to different ways how I/Os are counted:

- In the DASD Activity report, one I/O is counted for one SSCH or RSCH instruction. There can be record chaining, for example for paging I/O, which is not reflected in the SSCH count.
- In the Cache Subsystem Activity report, one I/O is counted for each cache request, and one I/O chain may cause several cache requests.

These two ways can lead to higher I/O rates in the Cache Subsystem Activity report than in the DASD Activity report.

Figure 192 on page 372 shows a DASD Activity report.

## Device data incomplete or missing

Device data can be incomplete or missing because:

- Device not available during entire interval
- Device changed or deleted
- Hardware data not available
- Required data not available
- Device in use
- Average cannot be calculated

**Device not available during entire interval:** Data lines are included for each device that has been online at least once since Monitor I session initialization. However, data is not reported for devices that were offline at the end of the reporting interval, that came online during the interval, or that were affected by dynamic device reconfiguration during the interval. One of the following messages will appear in the data line indicating the reason why data was not presented:

### **NOW ONLINE**

Brought online during this interval and still online at the end of the interval

### **NOW OFFLINE**

Taken offline during this interval and still offline at the end of the interval

### **OFFLINE**

Offline for the entire interval.

### **DEVICE DYNAMICALLY DELETED**

Device dynamically deleted during the interval.

### **DEVICE DYNAMICALLY CHANGED**

A device changed from static to dynamic during the interval, or a device deleted and a new device added with the same device number during the interval.

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

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

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

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

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

**Device in use:** RMF prints the status message **NO H/W DATA, DEVICE IN USE BY SYSTEM** when it cannot initialize the channel subsystem interface needed in order to gather the measurement data from the channel subsystem. This may occur for CTC devices that are being used by applications using protocols with never ending channel programs.

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

### **Overflow condition occurred**

Depending on the processor model you have, the hardware measurement data might be incomplete because of an overflow in the measurement timer. Any I/O request that exceeds that maximum time limit causes overflow. For example, chain scheduling, which the system uses when accessing page data sets or printing a SYSOUT data set, results in long channel programs and can cause timer overflow.

For shared DASD, pending times in excess of 8.3 seconds can occur due to RESERVE activity on the sharing system. Overflow conditions in pending time, however, are *not detected*.

**Overflow in non-Monitor II reports:** In Monitor II reports, the overflow counts are not reported. A value affected by overflow, however, is marked by an asterisk (\*). In exception reports and summary reports, overflows are neither detected nor

identified, because Monitor II device activity reports, as well as some exception reports and some fields in the summary report, are based on data the Monitor I session collects. The same inaccuracies apply to these reports.

**Connect/Disconnect time overflow:** Overflow conditions in connect time and disconnect time are detected by the hardware, counted by z/OS, and reported by RMF.

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

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

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

The report of overflows for paging devices does not necessarily indicate a problem. Consult the Page Data Set Activity report and the Workload Activity report to determine whether or not paging delays are a problem. If they are, the device data can be used in conjunction with the two reports to analyze the problem.

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

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

**PP - DEVICE**

DIRECT ACCESS DEVICE ACTIVITY																		PAGE 2	
z/OS V2R2		SYSTEM ID SYSF		DATE 09/28/2016		INTERVAL 14.59.998													
		RPT VERSION V2R2 RMF		TIME 06.00.00		CYCLE 1.000 SECONDS													
TOTAL SAMPLES = 900		IODF = 00		CR-DATE: 09/14/2016		CR-TIME: 10.31.31		ACT: POR											
STORAGE GROUP	DEV NUM	DEVICE TYPE	VOLUME OF	LCU	ACTIVITY RATE	AVG RESP TIME	AVG IOSQ	AVG CMR DLY	AVG DB DLY	AVG INT DLY	AVG PEND TIME	AVG DISC TIME	AVG CONN TIME	DEV CONN	DEV UTIL	DEV RESV	AVG NUMBER	AVG ANY	
XTEST	2208	33903	3339	TRXSX9	1 0032	0.001	.384	.000	.128	.000	.123	.256	.000	.128	0.00	0.00	0.0	0.0	100.0
	2209	33903	3339	TRXSXA	1 0032	0.001	.256	.000	.000	.000	.135	.256	.000	.000	0.00	0.00	0.0	0.0	100.0
	220A	33909	10017	TRXT01	1 0032	0.000	.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0
	220B	33909	10017	TRXT02	1 0032	0.000	.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0
	220C	33909	10017	TRXT03	1 0032	0.000	.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0
	220D	33909	10017	TRXT04	1 0032	0.000	.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0
	220E	33909	10017	TRXT05	1 0032	0.000	.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0
	220F	33909	32760	TRXT06	1 0032	0.000	.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0
	2210	33909	32760																
	2211	33909	32760																

Figure 192. Direct Access Device Activity Report

The reports for communication equipment, character reader devices, graphic devices, and unit record devices have the same format. The Communication Equipment Activity report is shown as example in Figure 193.

COMMUNICATION EQUIPMENT ACTIVITY																		PAGE 1	
z/OS V2R2		SYSTEM ID SYSF		DATE 09/28/2016		INTERVAL 15.00.000													
		RPT VERSION V2R2 RMF		TIME 06.15.00		CYCLE 1.000 SECONDS													
TOTAL SAMPLES = 900		IODF = 00		CR-DATE: 09/14/2016		CR-TIME: 10.31.31		ACT: ACTIVATE											
DEV NUM	DEVICE TYPE	VOLUME SERIAL	LCU	ACTIVITY RATE	AVG RESP TIME	AVG IOSQ	AVG CMR DLY	AVG DB DLY	AVG INT DLY	AVG PEND TIME	AVG DISC TIME	AVG CONN TIME	DEV CONN	DEV UTIL	DEV RESV	NUMBER ALLOC	AVG ANY	MT	NOT RDY
0120			0001	0.129	.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	100.0			0
0121			0001	0.129	.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	100.0			0
905D			007C	1.482	.291	.000	.017	.000	.001	.177	.086	.028	0.00	0.02	0.0	100.0			0
9E5D			007C	1.702	.588	.000	.018	.000	.003	.171	.587	.057	0.01	99.97	0.0	100.0			0
			LCU 007C	3.184	.314	.000	.018	.000	.002	.173	.314	.044	0.01	50.00	0.0	100.0			0

Figure 193. Communication Equipment Activity Report

The following figure shows the Magnetic Tape Device Activity report.

MAGNETIC TAPE DEVICE ACTIVITY																		PAGE 1	
z/OS V2R2		SYSTEM ID SYSE		DATE 09/28/2016		INTERVAL 15.00.027													
		RPT VERSION V2R2 RMF		TIME 23.45.00		CYCLE 1.000 SECONDS													
TOTAL SAMPLES = 810		IODF = 00		CR-DATE: 09/14/2016		CR-TIME: 12.03.30		ACT: ACTIVATE											
DEV NUM	DEVICE TYPE	VOLUME SERIAL	LCU	ACTIVITY RATE	AVG RESP TIME	AVG IOSQ	AVG CMR DLY	AVG DB DLY	AVG INT DLY	AVG PEND TIME	AVG DISC TIME	AVG CONN TIME	DEV CONN	DEV UTIL	DEV RESV	NUMBER OF MOUNTS	AVG MOUNT TIME	TIME DEVICE ALLOC	
0660	3590		0015	0.000	.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0	0	0	
0661	3590		0015	0.000	.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0	0	0	
0662	3590		0015	0.000	.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0	0	0	
0663	3590		0015	0.000	.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0	0	0	
			LCU 0015	0.000	.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0	0	0	

Figure 194. Magnetic Tape Device Activity Report

Table 155. Fields in the Device Activity Reports

Field Heading	Meaning
IODF = xx	The IODF number where xx is the suffix of the IODF data set name.
CR-DATE: mm/dd/yyyy	The creation date of the IODF.
CR-TIME: hh.mm.ss	The creation time of the IODF.
ACT: text	The configuration state where text indicates how the IODF was activated.
STORAGE GROUP	The name of the storage group to which the device belongs. Your storage administrator assigns the names. These names are available on the direct access device report only.
DEV NUM	The four-digit hexadecimal device number of a physical I/O device.
DEVICE TYPE	The device type on which the data set resides.

Table 155. Fields in the Device Activity Reports (continued)

Field Heading	Meaning
NUMBER OF CYL	The DASD volume capacity (in cylinders).
VOLUME SERIAL	The volume serial number (for direct access and magnetic tape reports) of the volume mounted on the device at the end of the reporting interval.
PAV	<p>The number of parallel access volumes (base and alias) which were available at the end of the reporting interval.</p> <p>If the number has changed during the reporting interval, it is followed by an '*'. If the device is a HyperPAV base device, the number is followed by an 'H', for example, 5.4H. The value is the average number of HyperPAV volumes (base and alias) in that interval.</p> $\text{Average \# of HPAV devices} = \frac{\text{Accumulated \# of HPAV devices}}{\text{Number of Samples}}$
LCU	<p>The number of the logical control unit (LCU) to which the device belongs.</p> <p>An LCU is a set of devices attached to the same physical control unit (or a group of physical control units with one or more devices in common.) The IOP, which is part of the channel subsystem, manages and schedules I/O work requests.</p> <p>There are two reasons that this field is blank:</p> <ul style="list-style-type: none"> <li>• RMF encountered an error while gathering data, check the operator console for messages.</li> <li>• This is a non-dedicated device in a z/VM guest system environment.</li> </ul>
DEVICE ACTIVITY RATE	<p>The rate at which start subchannel (SSCH) instructions to the device completed successfully.</p> $\text{ACTV RATE} = \frac{\text{\# Successful SSCH Instructions}}{\text{Interval Time}}$ <p>For devices using suspended channel programs, resume I/O requests are included in the SSCH counts.</p> <p>In the LCU summary line, this field contains the sum of the rates for each individual device.</p> <p>If the device has been deleted during the last interval, DEVICE DYNAMICALLY DELETED appears in the field instead of the measurement data.</p> <p>If the device has changed from static to dynamic, or was deleted and a new device added with the same device number, DEVICE DYNAMICALLY CHANGED appears in the field instead of the measurement data.</p>
AVG RESP TIME	<p>The average number of milliseconds the device required to complete an I/O request. This value reflects the total hardware service time and the front end software queuing time involved for the average I/O request to the device. The channel measures active time, which starts at the acceptance of a SSCH instruction (indicated by a condition code 0) and ends at the acceptance of the channel end (primary status pending). It does not, however, include the time required to process the interruption. The IOS queue length is factored in to reflect the front end queuing time.</p> $\text{AVG ACT TIME} = \frac{\text{Device Active Time}}{\text{Measurement Event Count}}$ $\text{AVG RESP TIME} = \text{AVG ACT TIME} + \text{AVG IOSQ TIME}$ <p>The active time is the sum of connect, disconnect, and pending time as described later.</p> <p>In the LCU summary line, this field contains the weighted average of the individual average response times for each device.</p>
AVG IOSQ TIME	<p>The average number of milliseconds an I/O request must wait on an IOS queue before a SSCH instruction can be issued.</p> $\text{AVG IOSQ TIME} = \frac{\text{IOSQ Count / \# Samples}}{\text{Device Activity Rate}}$



## PP - DEVICE

Table 155. Fields in the Device Activity Reports (continued)

Field Heading	Meaning
AVG CMR DLY	<p>The average number of milliseconds of delay that a successfully initiated start or resume function needs until the first command is indicated as accepted by the device. It allows to distinguish between real H/W errors versus workload spikes (contention in the fabric and at the destination port).</p> $\text{AVG CMR DLY} = \frac{\text{Initial Command Response Time}}{\# \text{ I/O Operations Accepted on that Path}}$
AVG DB DLY	<p>The average number of milliseconds of delay that I/O requests to this device encountered because the device was busy. Device busy might mean:</p> <ul style="list-style-type: none"> <li>• Another system is using the volume</li> <li>• Another system reserved the device</li> <li>• Head of string busy conditions caused contention</li> <li>• Some combination of these three conditions has occurred.</li> </ul> $\text{AVG DB DLY} = \frac{\text{Device Busy Delay Time}}{\text{Measurement Event Count}}$
AVG INT DLY	<p>The average interrupt delay time in units of milliseconds encountered for I/O requests to this device. For each I/O request, the time is measured from when the I/O operation is complete to when the operating system begins to process the status.</p> $\text{AVG INT DLY} = \frac{\text{Device Interrupt Delay Time}}{\text{Measurement Event Count}}$
AVG PEND TIME	<p>The average number of milliseconds an I/O request must wait in the hardware. This value reflects the time between acceptance of the SSCH function by the channel subsystem (SSCH-function pending) and acceptance of the first command associated with the SSCH function at the device (subchannel active). This value also includes the time waiting for an available channel path and control unit as well as the delay due to shared DASD contention.</p> <p>If the value is high, refer to the device's LCU entry in the I/O queuing activity report for an indicator of the major cause of the delay.</p> $\text{PEND TIME} = \frac{\text{Device Pending Time}}{\text{Measurement Event Count}}$
AVG DISC TIME	<p>The average number of milliseconds the device was disconnected while processing an SSCH instruction. This value reflects the time when the device was in use but not transferring data. It includes the overhead time when a device might disconnect to perform positioning functions such as SEEK/SET SECTOR, as well as any reconnection delay.</p> $\text{AVG DISC TIME} = \frac{\text{Device Disconnect Time}}{\text{Measurement Event Count}}$ <p>The measurement event count is the same as the number of SSCH instructions issued, unless there has been a timer overflow error in the channel.</p>
AVG CONN TIME	<p>The average number of milliseconds the device was connected to a channel path and actually transferring data between the device and central storage. Typically, this value, measures data transfer time but also includes the search time needed to maintain channel path, control unit, and device connection.</p> $\text{AVG CONN TIME} = \frac{\text{Device Connect Time}}{\text{Measurement Event Count}}$
% DEV CONN	<p>The percentage of time during the interval when the device was connected to a channel path.</p> $\% \text{ DEV CONN} = \frac{\text{Device Connect Time}}{\text{Interval Time}} * 100$



Table 155. Fields in the Device Activity Reports (continued)

Field Heading	Meaning
% DEV UTIL	<p>The percentage of time during the interval when the device was in use. This percentage includes both the time when the device was involved in I/O operations (connect and disconnect time) and the time when it was reserved but not involved in an I/O operation.</p> <p>The percentage reported represents the time during the interval when the device is "tied up" when it could not be used to service a request from another system. Some small portion of device busy (reserved) time is missed when the device is reserved but the I/O request is pending in the channel.</p> $\% \text{ DEV UTIL} = \left( \frac{\text{CON} + \text{DISC}}{\text{INT}} + \frac{\text{RSV}}{\text{N}} \right) * 100$ <p><b>CON</b> Device connect time  <b>DISC</b> Device disconnect time  <b>PAV</b> Number of parallel access volumes (base and alias); in case of non-PAV devices, PAV is set to 1  <b>RSV</b> Number of samples when the device was reserved but not involved in an I/O operation  <b>INT</b> Interval time (seconds)  <b>N</b> Total number of samples</p>
% DEV RESV	<p>The percentage of time during the interval when a shared device was reserved by the processor on which RMF was started.</p> <p>At each RMF cycle, RMF checks to see if a device is reserved, and a counter is kept of all such samples. At the end of the interval, the percentage is computed.</p> $\% \text{ DEV RESV} = \frac{\# \text{ Device-reserved Samples}}{\# \text{ Samples}} * 100$
AVG NUMBER ALLOC	<p>The average number of data control blocks (DCBs) and access method control blocks (ACBs) concurrently allocated for each volume. This field is reported only for direct access storage devices.</p> <p>At each RMF cycle, a counter is increased to reflect the number of data sets concurrently allocated. At the end of the interval, the average is calculated by dividing the total number of allocated data sets for all samples by the total number of samples.</p>
% ANY ALLOC	<p>The percentage of time during the reporting interval when the device was allocated to one or more data sets. Permanently mounted direct access devices show a 100% allocation, regardless of whether or not a data set was actually allocated.</p> <p>To determine the value, RMF keeps a count of whether or not the device was allocated or permanently resident at each cycle. At the end of the interval, the percentage is computed.</p> $\% \text{ ANY ALLOC} = \frac{\# \text{ Samples when the Device was Allocated}}{\# \text{ Samples}} * 100$
% MT PEND	<p>The percentage of time during the interval when a mount was pending for the device. This field is reported only for direct access devices and magnetic tape devices.</p> <p>At each cycle, RMF updates a counter when it detects a mount pending condition. At the end of the interval, the percentage is computed.</p> $\% \text{ MT PEND} = \frac{\text{Counter for Mount-Pending Condition}}{\# \text{ Samples}} * 100$
%NOT RDY	<p>The percentage of time during the reporting interval when the device was not ready for use. For example, when a tape has just been mounted but is not yet ready to be used to the system. This field is not reported for direct access devices. However, the value is recorded in the corresponding field of the SMF record, should your installation need the information.</p> <p>At each RMF cycle, a counter is updated when the status of the device indicates that it is not ready. At the end of the interval, the percentage is computed.</p> $\% \text{ NOT RDY} = \frac{\# \text{ Samples when the Device was not Ready}}{\# \text{ Samples}} * 100$

## PP - DEVICE

Table 155. Fields in the Device Activity Reports (continued)

Field Heading	Meaning
NUMBER OF MOUNTS	<p>The number of tape mounts, shown as an integer value, detected by RMF.</p> <p>If the tape mount was pending at the first cycle of the interval, an asterisk is placed before the numerical value of the tape mount. If the tape mount was pending at the last cycle of the interval, an asterisk is placed immediately following the numerical value of the tape mount.</p> <p>If a mount-pending condition is detected at the first cycle of the interval, the mount count for the interval increments by one.</p> <p>In the LCU summary line, this field contains the sum of all mount counts.</p> <p>This field is reported only for magnetic tape devices.  <b>Note:</b> Due to the fact that the tape mount count is a sampled value, it might happen that it does not contain all subsecond mounts of VTS devices.</p>
AVG MOUNT TIME	<p>The average mount time pending for every device, expressed in the form of HH:MM:SS.</p> $\text{AVG MOUNT TIME} = \frac{\# \text{ Samples Tape Mount was Pending} * \text{Interval}}{\# \text{ Samples}}$ <p>If the mount count or the sample count is zero, the result is zero.</p> <p>This field is reported only for magnetic tape devices.</p>
TIME DEVICE ALLOC	<p>The total time the device was allocated during the interval, expressed in the form of HH:MM:SS.</p> $\text{TIME DEVICE ALLOC} = \frac{\# \text{ Samples Tape Device was Allocated} * \text{Interval}}{\# \text{ Samples}}$ <p>If the sample count is zero, the result is zero.</p> <p>This field is reported only for magnetic tape devices.</p>

## Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following table shows the overview condition names for the Overview report.

Table 156. Overview names in the DASD Activity Report

Field Heading or Meaning	Overview Name
NUMBER OF CYL	DVCAP
DEVICE ACTIVITY RATE	DART
AVG RESP TIME	DRTAVG
AVG IOSQ TIME	DQTAVG
AVG CMR DLY	CMRDL
AVG INT DLY	INTDL
AVG DB DLY	DBDL
AVG PEND TIME	DPTAVG
AVG DISC TIME	DDTAVG
AVG CONN TIME	DCTAVG
% DEV UTIL	DVUTL
% DEV RESV	DR
% MT PEND	DMTPEND

### Data inaccuracies in duration report

When you request a duration report for I/O device activity, the identifier of each I/O device that had any VARY activity during the duration interval is followed by a single asterisk (\*). The data recorded for such a device is partial; that is, no data was collected during one or more of the measurement intervals included in the duration report because the device was varied during a measurement interval. The calculations the Postprocessor performs to generate a duration report make no adjustments for RMF measurement intervals when no data was collected for a device. Thus, the data for a device that moved online or offline might appear to be inconsistent. For example, the percentages reported for the QLENGTH DISTRIBUTION field do not equal 100% when data for the device is partial.

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

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

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

The following conditions can occur for storage group reporting:

- The STORAGE GROUP field shows **\*\*CHGD\*\*** for the volume if the storage group name changes in an SMF record for the duration period.
- The STORAGE GROUP field shows **\*\*CHGD\*\*** and the device identifier is followed by an asterisk if a volume is not reported in all SMF records of the duration and has changed the storage group name at least once.
- STORAGE GROUP DATA NOT AVAILABLE is reported between the TOTAL SAMPLES field and the report headings if the storage management subsystem is not available in one of the reports during the duration period.
- SMS INTERFACE ERROR, NEW STORAGE GROUP INFORMATION CANNOT BE OBTAINED is reported between the TOTAL SAMPLES = field and the report headings if a system-managed storage interface error occurs in one of the reports during the duration period.

---

## DOMINO - Lotus Domino Server report

The Domino Server family is an integrated messaging and Web application software platform. The Domino Server enables Web clients to communicate with Notes<sup>®</sup> servers.

The Lotus Domino Server report provides information about the activities of a server. The information can be used to analyze the activities of the server in case of problems.

### How to request this report

The Postprocessor is using SMF records type 108 as input for the Lotus Domino Server report. These records are not gathered by an RMF monitor, but are written by Domino servers. See the *z/OS RMF User's Guide* for details.

To produce this report, specify  
 REPORTS(DOMINO)

### Contents of the report

The report consists of two parts:

- Lotus Domino Server Summary  
 The summary contains one line for each server which is part of the report.
- Lotus Domino Server Details

This part consists of the following sections:

- Definition data (provided by record type 108-3)
- Performance data (provided by record type 108-3)
- Load data (provided by record type 108-1)

LOTUS DOMINO SERVER SUMMARY											PAGE 1
z/OS V2R2	SYSTEM ID LN21	DATE 09/28/2016	INTERVAL 05.00.000								
	RPT VERSION V2R2 RMF	TIME 18.40.00									
SERVER NAME	AVAILABLE HHH.MM.SS	---- USERS CONNECTED	----- ACTIVE	TASKS	TRANSACTION RATE	ASYNC READS	I/O RATE WRITES	MAIL RATE DELIVERED	SENT	SMTP RATE READS	WRITES
SUT1/COCPOK	000.30.00	2036	18	2136	62.50	137.0	101.9	4.42	0.46	0.00	0.00
BLUED1/BIGBLUE	000.30.00	5034	32	3532	119.31	207.4	199.3	9.14	1.04	0.00	0.00

Figure 195. Lotus Domino Server Report - Summary

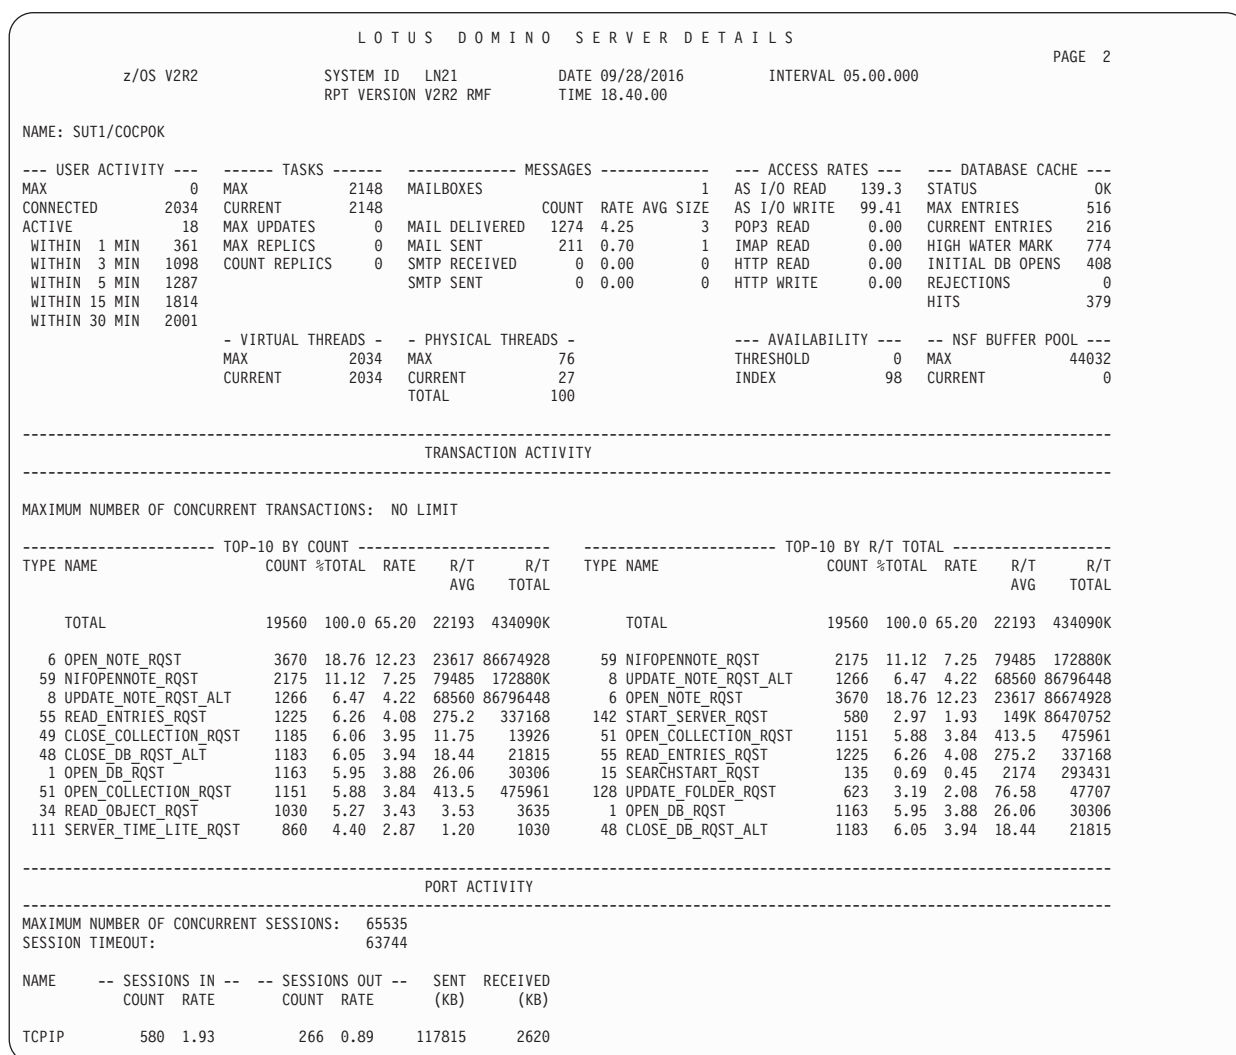


Figure 196. Lotus Domino Server Report - Details

Table 157. Fields in the Domino Server Summary Report

Field Heading	Meaning
SERVER NAME	Server name.
AVAILABLE	Total time (hhh.mm.ss) the server was available during the interval.
USERS CONNECTED	Average number of currently connected users.
USERS ACTIVE	Average number of currently active users.
TASKS	Average number of tasks currently in use.
TRANSACTION RATE	Rate of all transactions processed during the interval.
ASYNC I/O RATE - READS	Rate of asynchronous reads.
ASYNC I/O RATE - WRITES	Rate of asynchronous writes.
MAIL RATE - DELIVERED	Rate of Domino mail messages delivered to local users.
MAIL RATE - SENT	Rate of Domino mail messages sent to other servers.
SMTP RATE - READS	Rate of SMTP messages received from other servers.
SMTP RATE - WRITES	Rate of SMTP messages sent to other servers.

## PP - DOMINO

Table 158. Fields in the Domino Server Details Report

Field Heading	Meaning
NAME	Server name
<b>User Activity</b>	
MAX	Maximum number of users that are allowed to access the server. The value 0 means that there is no limit.
CONNECTED	Number of current users (connections).
ACTIVE	Number of active users.
WITHIN n MIN	Number of currently connected users that have been active within the last <b>1, 3, 5, 15,</b> and <b>30</b> minutes.
<b>Tasks</b>	
MAX	Maximum number of tasks in use.
CURRENT	Number of tasks currently in use.
MAX UPDATES	Maximum number of concurrent update tasks.
MAX REPLICS	Maximum number of concurrent replicator tasks.
COUNT REPLICS	Number of replications initiated by this server.
<b>Messages</b>	
MAILBOXES	Number of mail boxes.
MAIL DELIVERED	The number, rate and average size of Domino mail messages delivered to local users.
MAIL SENT	Domino mail messages sent to other servers.
SMTP RECEIVED	SMTP messages received from other servers.
SMTP SENT	SMTP messages sent to other servers.
<b>Access Rates</b>	
AS I/O READ	Rate of asynchronous I/O reads.
AS I/O WRITE	Rate of asynchronous I/O writes.
POP3 READ	Rate of POP3 reads.
IMAP READ	Rate of IMAP reads.
DOMINO READ	Rate of Domino reads.
DOMINO WRITE	Rate of Domino writes.
<b>Database Cache</b>	
STATUS	Status of the database cache: either <b>OK</b> or ? (=undefined).
MAX ENTRIES	Maximum number of database entries allowed in cache at any one time.
CURRENT ENTRIES	Number of current entries.
HIGH WATER MARK	High water mark.
INITIAL DB OPENS	Number of initial database opens.
REJECTIONS	Number of overcrowding rejections.
HITS	Hits in database cache.
<b>Virtual Threads</b>	
MAX	Maximum number of virtual thread pool threads.
CURRENT	Number of virtual thread pool threads currently in use.
<b>Physical Threads</b>	
MAX	Maximum number of physical thread pool threads in use.
CURRENT	Number of physical thread pool threads currently in use.
TOTAL	Total number of physical thread pool threads.
<b>Availability</b>	
THRESHOLD	Server availability threshold.
INDEX	Server availability index.

Table 158. Fields in the Domino Server Details Report (continued)

Field Heading	Meaning
<b>NSF Buffer Pool</b>	
MAX	Maximum size (in bytes) of the NSF (Notes Storage Facility) buffer pool.
CURRENT	Number of bytes of the NSF buffer pool currently in use.
<b>Transaction Activity</b>	
MAXIMAL CONCURRENT	Limit for number of concurrent transactions on a server.
<b>Top-10 List of Transaction Types</b> — Sorted by COUNT and by R/T TOTAL.	
TYPE	Transaction type.
NAME	Transaction name.
COUNT	Number of transactions processed during interval.
%TOTAL	%Percentage based on all transactions.
RATE	Rate of processed transactions.
R/T AVG	Average response time (milliseconds).
R/T TOTAL	Total response time (milliseconds) of all transactions that completed during the interval.
<b>Port Activity</b>	
MAX CONCURRENT SESSIONS	Maximum number of sessions that can run concurrently on the server.
SESSION TIMEOUT	Time limit (minutes) after which idle connections are terminated.
NAME	Port name.
SESSIONS IN	Count and rate of incoming sessions (from clients to the server) established during the interval.
SESSIONS OUT	Count and rate of outgoing sessions established during the interval.
SENT (KB)	Number of K bytes sent to the network.
RECEIVED (KB)	Number of K bytes received from the network.

## ENQ - Enqueue Activity report

The Enqueue Activity report provides information about resources that periodically build up queues of one or more requestors waiting to use the resource. Contention is reported for those resources where access is controlled by jobs that issue ENQ and DEQ macro instructions. RMF records related resource contention status changes signalled by GRS. When contention detail data (such as resource owner or numbers of exclusive/shared waiters) is passed with a signalled contention, RMF attributes these data to queue length buckets and reports individually on related contention measurements. To complement the picture about resource contentions, the portion of contention change events about which RMF does not obtain detail information is visible as percentage of the total number of contention change events.

### Using the information given in the report

Because the amount of time that a requestor must spend waiting for a resource can seriously affect system throughput, the information in this report can be very helpful in locating resources that consistently cause bottlenecks.

Once you have defined a critical resource, such as a serially-reusable resource that can be requested on either an exclusive or shared basis, your installation can improve the situation in a variety of ways. You could change the hardware configuration to release device bottlenecks, change data set placement, or

reschedule jobs to improve throughput, or re-specify the installation tuning parameter ERV (enqueue residence value) to give more processor time to the holder of the resource.

The information in the detail report can help you to balance your workload to minimize resource contention.

## How to request this report

To gather data for this report, specify as a Monitor I gatherer option  
ENQ(SUMMARY | DETAIL[,majorname,[minorname]])

To produce this report, specify  
REPORTS(ENQ)

**Note:** The ENQ report is only available as an interval report, not as a duration report.

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

### Example URL for the DDS API

<http://ddshost:8803/gpm/rmfpp.xml?reports=ENQ>

## Different report levels

The contents of the report depends on the gathering options:

- Summary report - ENQ(SUMMARY)
- Detail report - ENQ(DETAIL) or ENQ(DETAIL,majorname [,minorname])

The **Summary Activity report** includes:

- All resources for which contention has occurred during the reporting interval
- A description of the contention time for each resource
- A queue length distribution and average queue length for each resource
- Information on the type of requests made (either exclusive or shared)
- The total number of enqueue contention events that occurred.

An enqueue contention event is defined as the period from the time when the resource first has contention until the resource no longer has contention.

- The total number of contention status change events.

Contention status change events are events such as an incident where at least one waiter gets queued for a given resource, or an incident where the number of waiters or the contention owner changes, and also the contention-completion event at the end of an enqueue contention event.

- The percentage of status change events that did not provide contention detail data and therefore can not be attributed uniquely to queue length buckets.

A resource for which contention is still occurring at the end of the interval will be indicated by an asterisk following the TOT field, which is under the CONTENTION TIME field.

The **Detail Activity report** shows several lines of data for all resources for which contention occurs.



- The total number of jobs that own the resource and the names of one or two jobs that own the resource
- The total number of jobs that are waiting for the resource and the names of one or two jobs that are waiting for the resource.
- The identifier of the system on which the job is running following each job name
- An E if the request is exclusive or an S if the request is shared

RMF selects the job names shown in the detail report during the period of maximum contention in the interval by determining the longest contention event in the interval. For that event RMF reports the owners and waiters at the point when the event queue is the longest.

When there are several occurrences of the same length queue, the latest queue is reported. RMF reports the job names that were active at maximum contention even though those jobs might have been processed and flushed from the system by the time the contention no longer exists.

You can request data for a specific resource by specifying a *major* name, with or without a *minor* name. Various combinations of the reporting options can give you a complete picture of both critical resources and the jobs that are impacting system throughput by monopolizing a specific resource.

## Contents of the report

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

ENQUEUE ACTIVITY															PAGE 1	
z/OS V2R2				SYSTEM ID SYS1		DATE 09/28/2016		INTERVAL 14.59.946								
				RPT VERSION V2R2 RMF		TIME 16.30.00		CYCLE 1.000 SECONDS								
ENQUEUE DETAIL ACTIVITY																
-NAME- ----- CONTENTION TIME -----																
MAJOR MIN MAX TOT AVG																
MINOR																
GRS MODE: RING																
-- JOBS AT MAXIMUM CONTENTION--																
----- OWN ----- WAIT -----																
TOT NAME TOT NAME																
SYSNAME SYSNAME																
SYSZJES2																
SJB.2087F970																
0.000 0.000 0.000 0.000																
1 *MASTER*(E) 1 WEID (S)																
100 0.0 0.0 0.0 1.00 0 0 1 1 2 17 3.1																
RMFG RMFG																

Figure 197. Enqueue Detail Activity Report

Table 159. Fields in the Enqueue Activity Report

Field Heading	Meaning
GRS MODE	Shows the GRS mode in which the system is running.
NAME (MAJOR MINOR)	The name of a resource that has one or more requestors waiting. The major name is one to eight characters in length; the minor name can be from 1 to 255 characters, but only 44 characters will be printed. When the name exceeds 44 characters, it is truncated in the report. An asterisk (*) following the resource name indicates that it has been truncated. A resource with a scope of "SYSTEMS" will be followed by (SYSTEMS); a resource with a scope of "SYSTEM" will have no indication; and a resource with a scope of "STEP" will not be included in the report.

Table 159. Fields in the Enqueue Activity Report (continued)

Field Heading	Meaning
CONTENTION TIME (MIN MAX TOT AVG)	<p>The contention time observed for the resource during the RMF reporting interval. The maximum, minimum, total, and average contention times are reported in seconds. The time reported can be 0.000; this indicates a contention time of less than one-thousandth of a second and is most likely to appear as a minimum value.</p> <p>The contention time is calculated by subtracting the time the delay began (when the first ENQHOLD was issued) from the time the contention was ended (when the last ENQRLSE was issued) by freeing the resource. An asterisk(*) following the total contention time indicates that the contention extended beyond the end of the measurement interval.</p> <p>The calculation used to determine the average contention time is:</p> $\text{AVG CONT TIME} = \frac{\text{Contention Time for the Resource}}{\# \text{ Contention Events}}$
JOBS AT MAXIMUM CONTENTION	<p>The total number of resource owners and the total number of jobs waiting to use the resource. In addition, the names of one or two owners and one or two names of waiting jobs are reported.</p> <p>The reported counts refer to the period of maximum contention for a resource in the RMF reporting interval.</p> <p>RMF selects the names during the period of maximum contention for each resource. Within this period of maximum contention, RMF determines the point when the queue of waiting jobs was longest and reports the names of the first two jobs on the queue. Each name is followed by an (E) if that job requested exclusive use of the resource or an (S) if that job requested shared use of the resource. Under SYSNAME, RMF reports the name of the system on which the job is executing in a global resource serialization complex. This information can help you to determine which jobs were contributing most heavily to the contention for the resource.</p> <p>The field is reported only when the enqueue activity detail report is requested.</p>
% QLEN DISTRIBUTION (1 2 3 4+)	<p>The percentage of contention status change events during the interval when the number of requestors queued to the resource was one, two, three, four or more. The samples are taken for each contention status change where RMF receives contention detail data relatable uniquely to one of these four queue length buckets. Examples for a contention status change event are the change of the contention owner or the number of waiters. At each such sample, an accumulator for the observed length is updated. At the end of the measurement interval, the percentage for each queue length is computed.</p> <p>The calculation used for each queue length is:</p> $\% \text{ QLEN} = \frac{\text{Accumulator for that Queue Length}}{\text{Sum of Accumulators for all Queue Lengths}} * 100$
AVG Q LNGTH	<p>The average length of the queue of requestors that is waiting for the resource over the duration of the reporting interval. A consistently high number here indicates that the use of the resource is seriously out of balance.</p> $\text{AVG Q LNGTH} = \frac{\# \text{ Requestors Waiting}}{\text{Sum of Accumulators for all Queue Lengths}}$
-REQUEST TYPE- -EXCL -- SHARE - (MIN MAX MIN MAX)	<p>The type of the requests, either exclusive or shared that is waiting for use of the resource. The requestor would require exclusive use of the resource if the job expects to modify the resource or if the resource is by nature only serially reusable. Other requests would be for shared use of the resource. Both the minimum number and maximum number of waiting shared requests and waiting exclusive requests are reported.</p>
CONTENTION EVENT TOTAL	<p>The total number of resource contention events that occur during the measurement interval. This is the total number of periods each starting from the time when the resource has contention until the resource no longer has contention.</p>

Table 159. Fields in the Enqueue Activity Report (continued)

Field Heading	Meaning
CONTENTION STAT CHNG	TOTAL  The total number of contention status change events. This includes status change events relatable uniquely to specific queue length buckets about which RMF received contention information, as well as those status change events about which RMF has no specific contention detail data.  %NODET  The percentage of contention status change events which did not provide contention detail data.

## Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following table shows the overview condition names for the Overview report.

Table 160. Overview names in the Enqueue Activity Report

Field Heading or Meaning	Overview Name
CONTENTION TIME - TOT	ENQT
CONTENTION TIME - AVG	ENQAVG
CONTENTION TIME - MAX	ENQMAX
TOTAL EVENT	ENQNE
% CONTENTION STATUS CHANGE EVENTS NO DETAIL	ENQPNOD

## Messages

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

### NO CONTENTION OCCURRED

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

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

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

When the message occurs, you could reduce the length of the RMF interval, or, if you want to ensure that the contention activity for a specific resource is reported, you can request enqueue activity reporting for the specific critical resource.

### TERMINATE DUE TO DATA EVENT ERROR - TRY RERUN

Explanation: During the interval, the enqueue measurement routines encountered invalid data while processing a contention event. All enqueue measurement activity is terminated for the session; that is, the recovery

from the error includes modifying the enqueue activity option to NOENQ. Because the error encountered might not be a permanent error, you can modify the session options to re-specify enqueue measurement and reporting. If the message occurs again and there are no other indicators of a system problem, report the message to the RMF license holder at your installation.

#### **BAD CPU CLOCK OCCURRED - FIX CPU CLOCK AND RERUN**

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

---

## **ESS - Enterprise Disk Systems report**

The Enterprise Disk Systems report provides measurements about the activities of an enterprise disk system. RMF monitors the activity on an enterprise disk system independently from the source of the activity. Activity may be caused by the z/OS system on which RMF is running or from any other system using the enterprise disk system. You can use the data contained in this report for checking your current disk configuration, for bottleneck analysis and for capacity planning.

If the Monitor I data gatherer is set up to collect all available data, the report contains three sections described in the following:

- ESS Link Statistics
- ESS Extent Pool Statistics<sup>1</sup>
- ESS Rank Statistics<sup>1</sup>

### **ESS Link Statistics**

For each adapter of an ESS, this section contains statistics about the occurred I/O operations. One adapter supports one or more type of I/O (link type). The following link types are reported:

- ECKD read and write<sup>1</sup>
- SCSI read and write
- PPRC send and receive

For each link type, this section provides the average number of transferred bytes and the average number of operations per second as well as their average response time. The I/O intensity shows the utilization of the adapter during the report interval. Use this section for analysis of the external link usage and for capacity planning of the peer-to-peer remote copy (PPRC) links.

### **ESS Extent Pool Statistics**

This section provides capacity and performance information about allocated disk space. For each extent pool, it shows the real and virtual capacity, the number of real and virtual extents and the number of conversions from a virtual into a real

---

<sup>1</sup>. available for the IBM TotalStorage DS family

extent and vice versa. Such a conversion occurs when an application writes to a virtual extent. A conversion from a real into a virtual extent occurs if an extent is freed or migrated.

Use this section to check the available disk capacity and if required, change the capacity. For example, if you use the FlashCopy function of the enterprise disk system to create an instant point-in-time backup copy of your application data or data base, then you can read from the conversion information provided in this section, whether your provisioning strategy is successful. If it is dissatisfactory, you may decide to add real storage to the subsystem.

## ESS Rank Statistics

This section provides activity statistics about read and write operations in each rank of an extent pool. It also shows the number of arrays and the array width of all ranks. These values show the current configuration. The wider the rank, the more performance capability it has. By changing these values in your configuration, you can influence the throughput of your work. Use this section to detect and resolve performance problems and also for disk space capacity planning.

## How to request this report

The default option for Monitor I data gathering is NOESS. Therefore, you must specify the ESS Monitor I gatherer option if you want to get data for this report.

To produce this report, specify  
REPORTS(ESS)

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

### Example URL for the DDS API:

<http://ddshost:8803/gpm/rmfpp.xml?reports=ESS>

## Contents of the report

Depending on your Monitor I gatherer options and available data, this report consists of up to three sections. The following fields are common for all sections:

Table 161. Common Fields in the Enterprise Disk Systems Report

Field Heading	Meaning
SERIAL NUMBER	Serial number of the primary control unit.
TYPE-MODEL	ESS type and model.
CDATE	Date when the cache interval started.
CTIME	Time when the cache interval started.
CINT	Cache interval time.  In interval reports, the format is <i>mm.ss</i> , while in duration reports the format is <i>hh.mm.ss</i> .
<p><b>Note:</b> Device reserve activity can cause a data gatherer interface to wait until a reserve has been released. This in turn can cause the cache interval to be much longer than a regular RMF interval.</p> <p>Therefore, CDATE, CTIME and CINT have been introduced to show the actual point in time to which the cache interval start is related, and the actual cache interval length. All rates shown in the report are based on CINT, not on INTERVAL.</p>	

ESS LINK STATISTICS							PAGE 1
z/OS V2R2	SYSTEM ID VSL1	DATE 09/28/2016	INTERVAL 15.00.000				
	RPT VERSION V2R2 RMF	TIME 08.30.00	CYCLE 1.000 SECONDS				
SERIAL NUMBER 0000002471	TYPE-MODEL 002107-922	CDATE 09/28/2016	CTIME 08.29.30	CINT 15.00			
-----ADAPTER-----	--LINK TYPE--	BYTES	BYTES	OPERATIONS	RESP TIME	I/O	
SAID TYPE		/SEC	/OPERATION	/SEC	/OPERATION	INTENSITY	
0004 FIBRE 2Gb	ECKD READ	162.1K	13.7K	11.8	0.3	3.9	
	ECKD WRITE	2.4M	26.5K	92.5	0.8	76.2	
						-----	
						80.1	
0011 FIBRE 1Gb	NO DATA TO REPORT OR ZERO						
0024 FIBRE 2Gb	SCSI READ	156.0K	13.9K	11.2	0.3	3.6	
	SCSI WRITE	2.5M	26.5K	93.2	0.8	76.8	
						-----	
						80.4	
0088 FIBRE 2Gb	PPRC SEND	8.5M	50.4K	169.2	16.1	2729.9	
	PPRC RECEIVE	0.0	0.0	0.0	0.0	0.0	
						-----	
						2729.9	

Figure 198. ESS Link Statistics

RMF issues the informational message 'NO DATA TO REPORT OR ZERO' if the counters for all link types (both read or write) return 'zero'. This happens in the following cases: either there was no ESS activity in the report interval or the ESS did not deliver any data.

Table 162. Fields in the ESS Link Statistics

Field Heading	Meaning
ADAPTER	Specifies the channel adapter: <b>SAID</b> system adapter identifier <b>TYPE</b> adapter type, for example, FIBRE 2Gb; "Undefined", if RMF could not determine the type.
LINK TYPE	Type of I/O operation performed by the adapter, which can be one of the following: <ul style="list-style-type: none"> <li>ECKD READ or ECKD WRITE: designates extended count key data I/O</li> <li>SCSI READ or SCSI WRITE: designates small computer system interface I/O</li> <li>PPRC SEND or PPRC RECEIVE: designates peer-to-peer remote copy traffic</li> </ul>
BYTES /SEC	The average number of bytes transferred per second for all operations of the indicated link type during the reporting interval.
BYTES /OPERATION	The average number of bytes transferred per operation for all operations of the indicated link type during the reporting interval.
OPERATIONS /SEC	The average number of operations of the indicated link type per second during the reporting interval.
RESP TIME /OPERATION	The average response time of operations of the indicated link type during the report interval. This is the entire time from sending out a data block until the notice of receipt from the receiver arrives. This value is measured in milliseconds.
I/O INTENSITY	The portion of the reporting interval during which an adapter was active. It is the product of OPERATIONS/SEC times RESP TIME/OPERATIONS. The I/O intensity is provided as a total for each adapter as well as for each link type.  It is measured in milliseconds/second. That is, a value of 1000 for a link type indicates that this link was busy all the time during the report interval.  On a Fiber Channel, multiple data blocks can be sent concurrently without waiting for the notices of receipt. Therefore, an I/O intensity greater than 1000 ms per second for an adapter may occur if such concurrent operations had been active.

ESS EXTENT POOL STATISTICS								PAGE 2
z/OS V2R2		SYSTEM ID VSL1		DATE 09/28/2016		INTERVAL 15.00.000		
		RPT VERSION V2R2 RMF		TIME 08.30.00		CYCLE 1.000 SECONDS		
SERIAL NUMBER 0000022399	TYPE-MODEL	2107-921	CDATE	09/28/2016	CTIME	08.29.30	CINT 15.00	
--EXTENT POOL--		REAL			VIRTUAL			
ID	TYPE	CAPACITY	EXTENTS	CONVERSIONS	CAPACITY	EXTENTS	CONVERSIONS	
0000	CKD 1Gb	7579	8501	0	641	720	0	
0001	CKD 1Gb	7579	8501	0	641	720	0	
0002	FIBRE 1Gb	1542	1542	0	0	0	0	
0003	FIBRE 1Gb	1542	1542	0	0	0	0	
0004	CKD 1Gb	19978	22407	0	0	0	0	
0005	CKD 1Gb	19978	22407	0	0	0	0	
0006	CKD 1Gb	17187	19277	0	0	0	0	

Figure 199. ESS Extent Pool Statistics

The **ESS Extent Pool Statistics** section presents overview information on the defined disk capacity of extent pools.

Table 163. Fields in the ESS Extent Pool Statistics

Field Heading	Meaning
EXTENT POOL	Pool of allocation units for logical volumes. <b>ID</b> extent pool identifier <b>TYPE</b> extent pool type, for example, FIBRE 1Gb or CKD 1Gb.
REAL CAPACITY	The capacity of physical storage in gigabytes for real extents in an extent pool. This is available capacity for the operating system.
REAL EXTENTS	Number of real extents in an extent pool. A discrete number of extents can be used to create volumes.
REAL CONVERSIONS	Number of real extent conversions. A virtual extent is converted to a real extent, if an application writes to that extent.
VIRTUAL CAPACITY	Virtual extent pool capacity. The capacity of physical storage in gigabytes for virtual extents in an extent pool.
VIRTUAL EXTENTS	Number of virtual extents in an extent pool. A virtual extent is an extent that has not yet been converted to a real extent.
VIRTUAL CONVERSIONS	Number of virtual extent conversions. A real extent gets converted via migration or by freeing an extent.

ESS RANK STATISTICS															PAGE	8		
z/OS V2R2		SYSTEM ID SYSF				DATE 09/28/2016				INTERVAL 30.00.026								
		RPT VERSION V2R2 RMF				TIME 13.30.00				CYCLE 1.000 SECONDS								
SERIAL NUMBER	0000DKA61	TYPE-MODEL	002107-961	CDATE	09/28/2016	CTIME	13.30.00	CINT	30.00									
--EXTENT POOL--		---- READ OPERATIONS ----				---- WRITE OPERATIONS ----				----ARRAY----			MIN	RANK	RAID			
ID	TYPE	RRID	ADAPT	ID	OPS	BYTES	BYTES	RTIME	OPS	BYTES	BYTES	RTIME	SSD	NUM	WDTH	RPM	CAP	TYPE
/SEC	/OP	/SEC	/OP	/SEC	/OP	/SEC	/OP	/OP	/SEC	/OP	/SEC	/OP						
0000	CKD 1Gb	0000	0000		0.0	0.0	0.0	16.0	0.0	0.0	0.0	96.0	1	6	15	1800G	RAID 5	
		0004	0000		0.0	65.5K	72.8	0.0	0.0	1.3M	2.8K	100.0	1	7	15	2100G	RAID 5	
		0010	000A		190.0	57.2K	10.9M	2.2	8.0	1.1M	8.9M	9.3	Y	1	6	N/A	2400G	RAID 5
		0012	000A		180.6	57.3K	10.3M	2.3	8.3	1.1M	9.0M	9.5	Y	1	6	N/A	2400G	RAID 5
		POOL			370.6	57.2K	21.2M	2.2	16.3	1.1M	17.9M	9.4	Y	4	25	0	8700G	RAID 5
0001	CKD 1Gb	0001	0000		0.0	0.0	0.0	0.0	0.0	1.4M	22.9K	22.9	1	6	15	1800G	RAID 5	
		0005	0000		0.0	0.0	0.0	0.0	0.0	1.6M	22.9K	39.4	1	7	15	2100G	RAID 5	
		000F	000A		82.9	57.3K	4.7M	2.4	4.3	1.0M	4.5M	7.8	Y	1	6	N/A	2400G	RAID 5
		0011	000A		82.9	57.3K	4.7M	2.4	5.2	1.1M	5.6M	7.5	Y	1	6	N/A	2400G	RAID 5
		POOL			165.8	57.3K	9.5M	2.4	9.5	1.1M	10.1M	7.7	Y	4	25	0	8700G	RAID 5
0002	FIBRE 1Gb	0002	0000		0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1	6	15	1800G	RAID 5	
0003	FIBRE 1Gb	0003	0000		0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1	6	15	1800G	RAID 5	
0004	CKD 1Gb	0006	0002		0.8	25.2K	20.5K	8.0	23.2	1.3M	30.6M	16.9	1	6	10	7200G	RAID 5	
		0008	0002		0.7	25.2K	17.9K	12.7	19.8	1.3M	26.1M	34.6	1	6	10	7200G	RAID 5	
		000A	0002		0.5	18.3K	8.3K	11.5	17.3	1.5M	26.3M	23.9	1	7	10	8400G	RAID 5	
		POOL			2.0	23.7K	46.7K	10.5	60.4	1.4M	82.9M	24.7	3	19	10	22800G	RAID 5	
0005	CKD 1Gb	0007	0002		0.1	32.8K	3.8K	10.2	0.0	0.0	0.0	0.0	1	6	10	7200G	RAID 5	
		0009	0002		0.1	32.3K	3.6K	8.9	0.0	0.0	0.0	0.0	1	6	10	7200G	RAID 5	
		000B	0002		0.2	32.6K	5.0K	11.7	0.0	0.0	0.0	0.0	1	7	10	8400G	RAID 5	
		POOL			0.4	32.6K	12.5K	10.4	0.0	0.0	0.0	0.0	3	19	10	22800G	RAID 5	
0006	CKD 1Gb	000E	0003		0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1	5	7	20000G	RAID 6	

Figure 200. ESS Rank Statistics

Table 164. Fields in the ESS Rank Statistics

Field Heading	Meaning
EXTENT POOL	Pool of allocation units for logical volumes. <b>ID</b> extent pool identifier <b>TYPE</b> extent type, for example, FIBRE 1Gb or CKD 1Gb.
RRID	RAID rank identifiers in the extent pool. <b>Note:</b> The line where RRID = POOL contains the average for all rank values of the entire extent pool.
ADAPT ID	Adapter Pair ID.
READ OPERATIONS OPS/SEC	Number of read operations per second.
READ OPERATIONS BYTES/OP	Average number of bytes per read operation.
READ OPERATIONS BYTES/SEC	Average bandwidth of a read operation.
READ OPERATIONS RTIME/OP	Average response time of read operations in milliseconds.
WRITE OPERATIONS OPS/SEC	Number of write operations per second.
WRITE OPERATIONS BYTES/OP	Average number of bytes per write operation.
WRITE OPERATIONS BYTES/SEC	Average bandwidth of a write operation.
WRITE OPERATIONS RTIME/OP	Average response time of write operations in milliseconds.



Table 164. Fields in the ESS Rank Statistics (continued)

Field Heading	Meaning
ARRAY	<p><b>SSD</b> If a Y is displayed, then there is at least one solid state drive in the rank array.</p> <p><b>NUM</b> Number of arrays on the rank.</p> <p><b>WDTH</b> Sum of DDMs (disk drive modules) of a rank excluding spares of the rank. For example, if you have a RAID-5 array with 6 data disks and 1 parity disk, ARRAY WDTH is 7, or for a RAID-10 with 3 mirrored disks, ARRAY WDTH is 6.</p>
MIN RPM	The slowest drive of the rank in units of 1000 RPM (rounds per minute).
RANK CAP	The sum of bytes of a rank.
RAID TYPE	<p>RAID type found for the rank, for example,</p> <ul style="list-style-type: none"> <li>• RAID-5</li> <li>• RAID-10</li> </ul> <p>In the line displaying the average values for the entire extent pool (where RRID = POOL), 'MIXED' is shown if different RAID types have been encountered for the individual ranks in the extent pool.</p>

### Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following tables show the overview condition names for the Overview report, divided according to the sections of the enterprise disk systems report.

Table 165. Overview names in the ESS Link Statistics section

Field Heading or Meaning	Overview Name
BYTES/SEC for SCSI READ	ESTRSRD
BYTES/SEC for SCSI WRITE	ESTRSWR
BYTES/OPERATION for SCSI READ	ESPSSRD
BYTES/OPERATION for SCSI WRITE	ESPSSWR
OPERATIONS/SEC for SCSI READ	ESARSRD
OPERATIONS/SEC for SCSI WRITE	ESARSWR
RESP TIME/OPERATION for SCSI READ	ESRTSRD
RESP TIME/OPERATION for SCSI WRITE	ESRTSWR
I/O INTENSITY for SCSI READ	ESIOISRD
I/O INTENSITY for SCSI WRITE	ESIOISWR
I/O INTENSITY for SCSI TOTAL	ESIOIST
BYTES/SEC for ECKD READ	ESTRERD
BYTES/SEC for ECKD WRITE	ESTREWR
BYTES/OPERATION for ECKD READ	ESPSERD
BYTES/OPERATION for ECKD WRITE	ESPSEWR
OPERATIONS/SEC for ECKD READ	ESARERD
OPERATIONS/SEC for ECKD WRITE	ESAREWR
RESP TIME/OPERATION for ECKD READ	ESRTERD
RESP TIME/OPERATION for ECKD WRITE	ESRTEWR
I/O INTENSITY for ECKD READ	ESIOIERD
I/O INTENSITY for ECKD WRITE	ESIOIEWR
I/O INTENSITY for ECKD TOTAL	ESIOIET
BYTES/SEC for PPRC SEND	ESTRPSD
BYTES/SEC for PPRC RECEIVE	ESTRPRV

Table 165. Overview names in the ESS Link Statistics section (continued)

Field Heading or Meaning	Overview Name
BYTES/OPERATION for PPRC SEND	ESPSPSD
BYTES/OPERATION for PPRC RECEIVE	ESPSPRV
OPERATIONS/SEC for PPRC SEND	ESARPSD
OPERATIONS/SEC for PPRC RECEIVE	ESARPRV
RESP TIME/OPERATION for PPRC SEND	ESRTPSD
RESP TIME/OPERATION for PPRC RECEIVE	ESRTPRV
I/O INTENSITY for PPRC SEND	ESIOIPSD
I/O INTENSITY for PPRC RECEIVE	ESIOIPRV
I/O INTENSITY for PPRC TOTAL	ESIOIPT

Table 166. Overview names in the ESS Extent Pool Statistics section

Field Heading	Overview Name
REAL CAPACITY	ESXRCAP
REAL EXTENTS	ESXRNSG

Table 167. Overview names in the ESS Rank Statistics section

Field Heading	Overview Name
READ OPERATIONS OPS/SEC	ESRROP
READ OPERATIONS BYTES/OP	ESRRBOP
READ OPERATIONS BYTES/SEC	ESRRBD
READ OPERATIONS RTIME/OP	ESRRRT
WRITE OPERATIONS OPS/SEC	ESRWOP
WRITE OPERATIONS BYTES/OP	ESRWBOP
WRITE OPERATIONS BYTES/SEC	ESRWBD
WRITE OPERATIONS RTIME/OP	ESRWRT

---

## FCD - FICON Director Activity report

With the Fibre Channel architecture and Fibre Channel switches (referred to in the following as FICON directors), link busy conditions are not returned. Instead, the FICON director queues the frames internally and sends them through when the port becomes available. This switch latency can grow as contention for ports increases. Therefore, it is important to report this switch latency (per port), this helps for the following tasks:

- Capacity planning
- Analysis of performance problems and bottlenecks
- Identification of contributors to device pending and disconnect times
- Understanding the contention for reconnection status

### How to request this report

The default option for Monitor I data gathering is NOFCD. Therefore, it is required that you specify FCD if you want to get this report.

To produce this report, specify  
 REPORTS(FCD(option))

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

### Example URL for the DDS API

<http://ddshost:8803/gpm/rmfpp.xml?reports=FCD>

## Contents of the report

The measurements provided for a port in the FCD report do not only comprise the I/O for the system on which the report is taken, but include all I/O that is directed through this port, regardless of which LPAR requests the I/O.

F I C O N D I R E C T O R A C T I V I T Y									
z/OS V2R2		SYSTEM ID SYS1		DATE 09/28/2016		INTERVAL 15.00.000		PAGE 1	
		RPT VERSION V2R2 RMF		TIME 07.15.00		CYCLE 1.000 SECONDS			
IODF = 99 CR-DATE: 09/14/2016 CR-TIME: 08.45.00 ACT: POR									
SWITCH DEVICE: 0414 SWITCH ID: 01 TYPE: 005000 MODEL: 001 MAN: MCD PLANT: 01 SERIAL: 00000MK00109									
PORT ADDR	UNIT	CONNECTION ID	SERIAL NUMBER	AVG FRAME PACING	AVG FRAME READ	AVG FRAME WRITE	PORT BANDWIDTH (MB/SEC)		ERROR COUNT
							-- READ --	-- WRITE --	
05	CHP	FA	000000099802	0	808	285	50.04	10.50	0
07	CHP	4A	00000009F230	0	149	964	20.55	5.01	0
08	CHP-H	F4	000000070B82	0	568	965.5T	70.56	4.02	1
09	CHP	FC	00000003C03F	0	558	1424	50.07	10.53	0
0B	CHP	F4	00000004C057	0	872	896	50.00	10.56	0
12	CHP	D5	00000005C86D	0	73	574	20.51	5.07	0
13	CHP	C8	00000008C1DF	0	868	1134	70.52	2.08	1
14	SWITCH	----	0000013124DA	0	962	287	50.03	10.59	0
15	CU	C800	0000000CF811	0	1188	731	20.54	5.00	0
	CU	CA00	0000000CD111				70.55	3.01	1
16	CHP	CB	00000009CE35	0	740	1185	70.50	2.06	1

Figure 201. FICON Director Activity Report

Table 168. Fields in the FICON Director Activity Report.

Field Heading	Meaning
IODF = xx	The IODF number where xx is the suffix of the IODF data set name.
CR-DATE: mm/dd/yyyy	The creation date of the IODF.
CR-TIME: hh.mm.ss	The creation time of the IODF.
ACT: text	The configuration state where text indicates how the IODF was activated.
SWITCH DEVICE	The hexadecimal number of the switch device of the FICON director for which measurements are being reported.
SWITCH ID	The hexadecimal switch identifier of the FICON director which is associated with this switch device. In case of cascaded switches, '**' may be shown.
TYPE, MODEL, MAN, PLANT, SERIAL	The hardware description of the switch device.
PORT ADDR	The hexadecimal address of the port.

Table 168. Fields in the FICON Director Activity Report. (continued)

Field Heading	Meaning
CONNECTION	<p>Provides information about the connected unit.</p> <p><b>UNIT</b></p> <ul style="list-style-type: none"> <li>• CHP: denotes a channel path</li> <li>• CHP-H: denotes a channel path of the system which requested this report</li> <li>• CU: denotes a control unit</li> <li>• SWITCH: denotes a switch</li> </ul> <p>If the unit is not unique, dashes are displayed. For example, for CTC channels, there might be a CU and a CHP connected to the same port.</p> <p><b>ID</b> The hexadecimal identifier of the connector.</p> <p>For connection unit SWITCH, dashes are provided.</p> <p>Dashes are also displayed in this field for UNIT = CU, if the system with the FCD data gathering option ON is not connected to that control unit.</p> <p><b>SERIAL NUMBER</b> The serial number of the connected unit.</p>
AVG FRAME PACING	The average time (in microseconds) a frame had to wait before it could be transmitted.
AVG FRAME SIZE READ	The average frame size (in bytes) used to receive data during the interval.
AVG FRAME SIZE WRITE	The average frame size (in bytes) used to transmit data during the interval.
PORT BANDWIDTH READ	The rate (in MB/sec) of data which was received during the interval.
PORT BANDWIDTH WRITE	The rate (in MB/sec) of data which was transmitted during the interval.
ERROR COUNT	The number of errors which were encountered during the interval.

### Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following table shows the overview condition names for the Overview report.

Table 169. Overview names in the FICON Director Activity Report

Field Heading or Meaning	Overview Name
AVG FRAME PACING	FDAFPT
PORT BANDWIDTH (READ - MB/SEC)	FDMBREAD
PORT BANDWIDTH (WRITE - MB/SEC)	FDMBWRT
ERROR COUNT	FDNERR

## HFS - Hierarchical File System Statistics report

The Hierarchical File System Statistics report provides information about activities and storage usage within your z/OS UNIX environment. This data can be used to analyze whether storage and buffer pool definitions are correct, or whether some adjustments should be performed to improve the performance of I/O activities for HFS files.

### How to request this report

Monitor III gathers global data for this report as SMF record type 74.6. If you want to get information about specific hierarchical file systems, you have to activate the Monitor III gatherer option `HFSNAME(ADD(hfsname))`.

To produce this report, specify  
`REPORTS(HFS)`

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

### Example URL for the DDS API

<http://ddshost:8803/gpm/rmfpp.xml?reports=HFS>

## Contents of the report

The report consists of two parts.

### HFS Global Statistics Report

The first part of the HFS report provides overall data about I/O activities of HSF files and gives statistics about the various buffer pools which have been defined.

The report can be used as an entry point for performance investigation and capacity planning.

### HFS File System Statistics Report

The second part of the report is based on data gathering for specific file systems. You get data about I/O activities and about the internal structure (index) of the HFS files.

Both parts of the report can help you

- in getting a general understanding of the throughput recognized and achieved by HFS to optimally use your resources,
- in identifying potential problems and bottlenecks within HFS and taking corrective actions.

## HFS Global Statistics

HFS GLOBAL STATISTICS										PAGE 1
z/OS V2R2		SYSTEM ID SYS1		DATE 09/28/2016	INTERVAL 15.00.000					
		RPT VERSION V2R2 RMF		TIME 12.00.00	CYCLE 1.000 SECONDS					
--- STORAGE LIMITS (MB) ---				----- FILE I/O -----		--- METADATA I/O ---				
VIRTUAL	MAX	1000		COUNT	RATE	COUNT	RATE			
	USE	766	CACHE	3543	3.937	800	0.889			
FIXED	MIN	200	DASD	200	0.222	173	0.192			
	USE	400	HIT RATIO	94.66		82.22				
----- BUFFER POOL STATISTICS -----										
POOL NUMBER	NUMBER BUFFERS	BUFFER SIZE	PAGES	POOL SIZE BYTES	%FIXED	DATA SPACES	TOTAL	I/O ACTIVITY FIXED	%FIXED	
1	39353	1	39353	154M	0	1	115K	0	0	
2	0	4	0	0K	0	1	16	0	0	
3	2	16	32	128K	0	1	15	0	0	
4	2	64	128	512K	0	1	4	0	0	

Figure 202. HFS Global Statistics Report

Table 170. Fields in HFS Global Statistics Report

Field Heading	Meaning
<b>Storage Limits</b> - All fields are given in megabytes and show the values at interval end.	
VIRTUAL MAX	Value of VIRTUAL(MAX) parameter.
VIRTUAL USE	Total amount of virtual storage assigned to I/O buffers.
FIXED MIN	Value of FIXED(MIN) parameter.
FIXED USE	Total amount of permanently fixed storage assigned to I/O buffers. This number is included in the VIRTUAL USE field.

## PP - HFS

Table 170. Fields in HFS Global Statistics Report (continued)

Field Heading	Meaning
<b>File I/O</b> - The fields are given as COUNT and RATE (count per second).	
CACHE	The first page of a data file was requested and found in virtual storage (cache).
DASD	The first page of a data file was requested and not found in virtual storage, and an I/O was necessary.
HIT RATIO	Percentage of cache-found requests based on total number of requests.
<b>Metadata I/O</b> - The fields are given as COUNT and RATE (count per second).	
CACHE	The metadata for a file was found in virtual storage during file lookup.
DASD	The metadata for a file was not found in virtual storage during file lookup, and an index call was necessary which may result in an I/O.
HIT RATIO	Percentage of cache-found requests based on total number of requests.
<b>Buffer Pool Statistics</b>	
POOL NUMBER	HFS defines up to four buffer pools for processing. This number is used to refer to one of these pools.
NUMBER BUFFERS	Number of buffers in this buffer pool currently residing in virtual storage.
BUFFER SIZE	Size of each buffer in this pool (in pages).
POOL SIZE - PAGES	Size of this buffer pool currently in virtual storage (in pages).
POOL SIZE - BYTES	Size of this buffer pool currently in virtual storage (in bytes).
POOL SIZE - %FIXED	Percentage of the size of the buffers which are permanently fixed.
DATA SPACES	Number of data spaces comprising this buffer pool.
I/O ACTIVITY - TOTAL	Total number of buffers in this buffer pool for which I/Os were issued. This is not necessarily the number of actual I/Os issued since multiple buffers can be written in a single I/O request.
I/O ACTIVITY - FIXED	Number of times a buffer was already fixed prior to an I/O request in this buffer pool.
I/O ACTIVITY - %FIXED	Percentage of fixed I/Os.

## HFS File System Statistics

HFS FILE SYSTEM STATISTICS										
z/OS V2R2		SYSTEM ID	SYS1	DATE	09/28/2016	INTERVAL	15.00.000	PAGE 2		
		RPT VERSION	V2R2 RMF	TIME	12.00.00	CYCLE	1.000 SECONDS			
--- ALLOCATION (MB) ---		----- FILE I/O ----		--- METADATA I/O ---		---- INDEX I/O ----		---- INDEX EVENTS ---		
SIZE		COUNT	RATE	COUNT	RATE	COUNT	RATE	COUNT		
FILE SYSTEM NAME: OMVS.SYS4.ROOT										
MOUNT DATE: 04/03/2016 TIME: 07:58:21										
SYSTEM	172	CACHE	0	0.000	15	0.017	75	0.083	NEW LEVEL	0
DATA	50	DASD	20	0.022	0	0.000	0	0.000	SPLITS	0
ATTR. DIR	4.714	HIT RATIO	0.00		100.00		100.00		JOINS	0
		SEQUENTIAL	20	0.022						
CACHED	0.000	RANDOM	0	0.000						
FILE SYSTEM NAME: OMVS.SYS4.S670D05.USR										
FILE SYSTEM DATA IS NOT AVAILABLE. BPX1PCT RC= 81, RS= 105.										
FILE SYSTEM NAME: OMVS.SYS4.USERS										
MOUNT DATE: 04/03/2016 TIME: 07:58:24										
SYSTEM	563	CACHE	3550	3.944	3257	3.619	122K	135.092	NEW LEVEL	0
DATA	562	DASD	1340	1.489	10	0.011	10	0.011	SPLITS	278
ATTR. DIR	221	HIT RATIO	72.60		99.69		99.99		JOINS	0
		SEQUENTIAL	0	0.000						
CACHED	0.000	RANDOM	0	0.000						

Figure 203. HFS File System Statistics Report

Table 171. Fields in the HFS File System Statistics Report

Field Heading	Meaning
FILE SYSTEM NAME	The name of the HFS file system which has been selected for reporting.
MOUNT DATE and TIME	Date and time when the selected file system was mounted.
<b>Allocation</b> - All fields are given in megabytes.	
SYSTEM	Amount of storage allocated to this file system.
DATA	Amount of storage internally used within HFS for data files, directories and HFS internal structures like the attribute directory (AD).
ATTR. DIR	Amount of storage used for the attribute directory (AD). This number is included in the DATA field.  The attribute directory is the internal HFS structure (index) which contains attribute information about individual file system objects as well as attributes of the file system itself.
CACHED	Amount of data buffer storage cached by this file system.
<b>File I/O</b> - The fields are given as COUNT and RATE (count per second).	
CACHE	The first page of a data file was requested and found in virtual storage (cache).
DASD	The first page of a data file was requested but was not found in virtual storage (cache) and an I/O was necessary.
HIT RATIO	Percentage of cache-found requests based on total number of requests.
SEQUENTIAL	Sequential file data I/O requests.  A sequential I/O is one of a series of I/Os to read or write a data file, where the first I/O started at the first byte of the file and each subsequent I/O was for the next sequential set of bytes.
RANDOM	Random file data I/O requests.  A random I/O is an I/O that does not read or write the start of a file, and was not preceded by an I/O that read or wrote the immediately preceding set of bytes.
<b>Metadata I/O</b> - The fields are given as COUNT and RATE (count per second).	
CACHE	The metadata for a file was found in virtual storage (cache) during file lookup.
DASD	The metadata for a file was not found in virtual storage during file lookup and an index call was necessary which may result in an I/O.

Table 171. Fields in the HFS File System Statistics Report (continued)

Field Heading	Meaning
HIT RATIO	Percentage of cache-found requests based on total number of requests.
<b>Index I/O</b> - The fields are given as COUNT and RATE (count per second).	
CACHE	Index page read/write hits.
DASD	Index page read/write misses.
HIT RATIO	Percentage of cache-found requests based on total number of requests.
<b>Index Events</b>	
NEW LEVEL	Number how often HFS added a new level to its index structure.  The index statistics are relative to all of the indices in the HFS data set. The attribute directory (AD) is one index (the largest), but each directory (including the root) is also an index.
SPLITS	Number how often an index page was split into two pages because new records were inserted. This gives an idea of how much insertion activity there has been for the index structure.
JOINS	Number how often HFS was able to combine two index pages into one, because enough index records had been deleted in the two pages.

## Special considerations

It might be possible that some data is not available during data gathering. This will result in an incomplete report containing one of the following error messages:

OMVS KERNEL NOT READY

BUFFER LIMIT DATA IS NOT AVAILABLE. BPX1PCT RC= rc, RS= rs.

GLOBAL HFS DATA IS NOT AVAILABLE. BPX1PCT RC= rc, RS= rs.

GLOBAL HFS DATA IS PARTIALLY AVAILABLE.

FILE SYSTEM DATA IS NOT AVAILABLE. BPX1PCT RC= rc, RS= rs.

MOUNT TIME CHANGED DURING INTERVAL.

FILE SYSTEM NOW MOUNTED.

Please, refer to *z/OS UNIX System Services Messages and Codes* for an explanation of the return and reason code.

---

## HTTP - HTTP Server report

The HTTP Server is the Web server for the family of WebSphere® application servers which provide the run-time environment for e-business applications.

The HTTP Server report provides information about the activities of a server. The information can be used to analyze the activities of the server in case of problems.

## How to request this report

The Postprocessor requires type 103 subtypes 1 and 2 SMF records as input for the HTTP Server report. These records are not gathered by an RMF monitor, but were written by the IBM HTTP Server (IHS) powered by Domino, which is no longer supported in z/OS V2R2. SMF type 103 records created on a prior release of z/OS can still be used to generate a RMF Postprocessor HTTP report.

The IBM HTTP server powered by Apache does not write type 103 subtypes 1 or 2 SMF records, which means that no RMF Postprocessor HTTP report can be generated for that HTTP server.



To produce this report, specify  
REPORTS(HTTP)

## Contents of the report

The report consists of two parts:

- HTTP Server Summary

The summary contains one line for each server which is part of the report.

- HTTP Server Details

This part consists of two sections:

- Configuration data (provided by record type 103-1)
- Performance data (provided by record type 103-2)

Configuration data is reported together with performance data. Configuration data is not reported, if there is not at least one corresponding performance data record.

If there is no configuration data available, the line

```
*** NO CONFIGURATION DATA AVAILABLE WITHIN GIVEN RECORD INTERVAL ***
```

is shown.

H T T P S E R V E R S U M M A R Y													
z/OS V2R2		SYSTEM ID	SYS1	DATE 09/28/2016		INTERVAL 30.00.000							PAGE 1
SERVER NAME		AVAILABLE	REQUEST	RESPONSE	THROUGHPUT RATE		THREADS	CACHE SIZE		CACHE FILES		TIMEOUTS	
SERVER TOKEN		HHH.MM.SS	RATE	RATE	IN	OUT	MAX	USED	MAX	USED	MAX	USED	
MVS071	4F37-CA73-0005-678E	000.24.35	0.03	0.03	5.46	30.15	39	0.00	5120	0.90	0	0.00	1
mvs047.tcp.raleigh.ibm.com	3D29-CA40-0006-997C	000.01.10	2.14	0.07	8.67	41.37	23	0.00	5120	1.06	0	0.00	0
s390server17.wscilab.washington.ibm.com	3D29-CA4C-000B-7E86	000.07.37	0.00	0.00	0.00	0.00	150	0.00	5120	0.00	NO	0.00	0
s390server18.idelab.boeblingen.ibm.com	3D29-CA40-0006-997C	000.15.00	3.57	3.57	990.6	12838	40	2.00	5120	2.59	NO	1.00	42

Figure 204. HTTP Server Report - Summary

# PP - HTTP

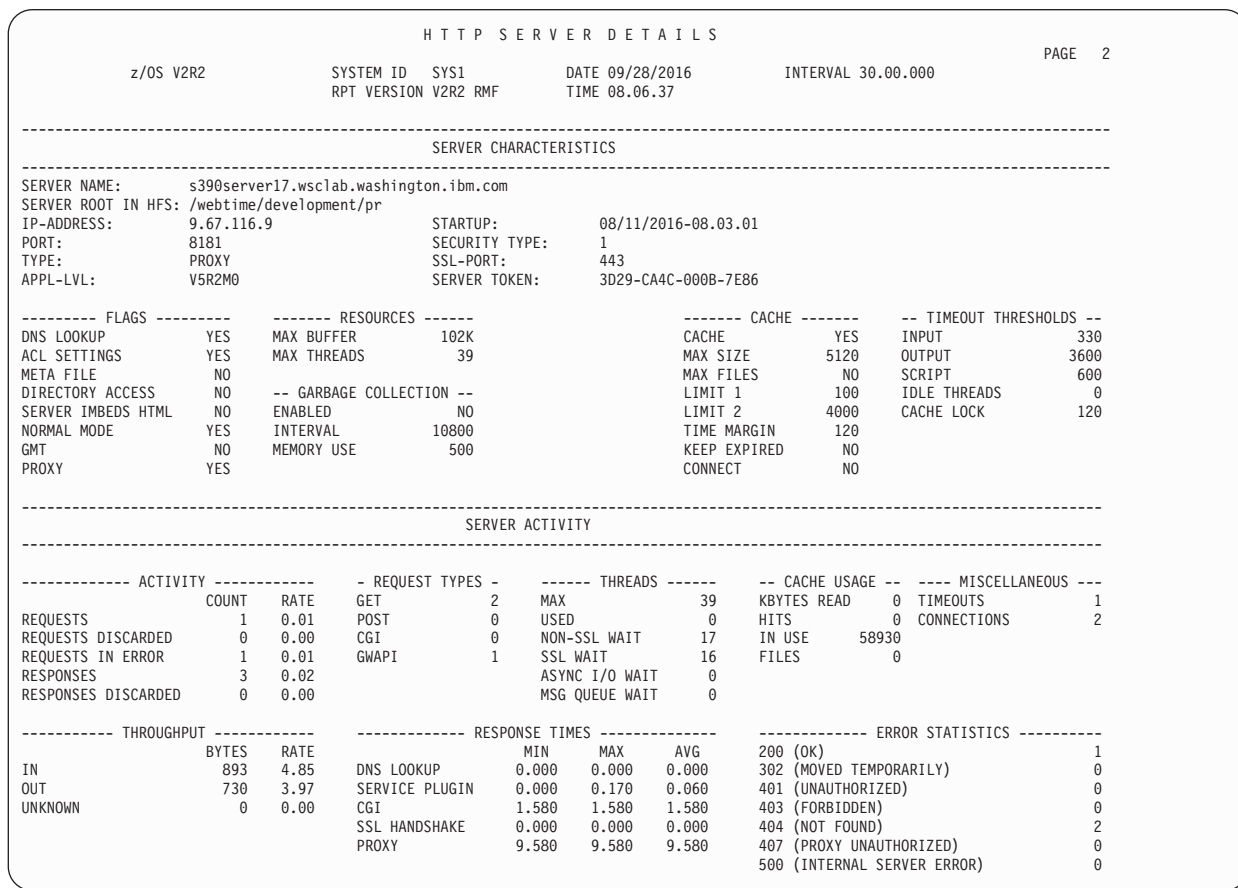


Figure 205. HTTP Server Report - Details

Table 172. Fields in the HTTP Server Summary Report

Field Heading	Meaning
SERVER NAME	Server name. If the server name is longer than 32 characters, the line is broken after the name and the value are displayed in the line below, as shown in Figure 204 on page 399.
SERVER TOKEN	When running multiple HTTP servers or operating in scalable server mode, multiple instances of the HTTP server have the same server name. The server token provides a unique identification of each server instance. If a server is restarted, it keeps its token. Thus, identical tokens may appear in the summary and detail section. If the token cannot be built from SMF record type 103, N/A is shown instead.
AVAILABLE	Total time (hhh.mm.ss) the server was available during the interval.
REQUEST RATE	Number of requests that the HTTP server has successfully served per second.
RESPONSE RATE	Number of successful responses sent per second.
THROUGHPUT RATE	Number of bytes received or sent by this server per second.
THREADS	<b>MAX</b> Maximum number of threads the server can have in the thread pool (or N0 if no limit has been specified). <b>USED</b> Number of currently active threads of the server.
CACHE SIZE	<b>MAX</b> Maximum cache size (KB) for this server. <b>USED</b> Used cache size of this server.
CACHE FILES	<b>MAX</b> Maximum number of files to be in the cache of this server. <b>USED</b> Number of files in the cache of this server.
TIMEOUTS	Number of timeouts on the server.

Table 173. Fields in the HTTP Server Details Report

Field Heading	Meaning
<b>Server Characteristics - Configuration Data</b>	
SERVER NAME	Server name
IP-ADDR	IP address of the host this HTTP server runs on.
PORT	Port number this HTTP server listens to.
TYPE	Server role. <b>HTTP</b> Simple or normal HTTP server <b>PROXY</b> Proxy server <b>CACHING</b> Caching server <b>CACHING PROXY</b> Caching proxy <b>UNKNOWN</b> Unknown server role
APPL-LVL	Version of software the server is running.
SERVER ROOT IN HFS	Directory for server_root.
STARTUP	Server startup date/time.
SECURITY TYPE	Security type.
SSL-PORT	Security port.
<b>Server Characteristics - Flags</b>	
DNS LOOKUP	DNS lookup flag.
ACL SETTINGS	ACL settings.
META FILE	Meta file flag.
DIRECTORY ACCESS	Directory access flag.
SERVER IMBEDS HTML	Server imbeds HTML flag.
NORMAL MODE	Normal mode flag.
GMT	GMT flag.
PROXY	Proxy flag.
<b>Server Characteristics - Resources</b>	
MAX BUFFER	Maximum size of content buffer.
MAX THREADS	Maximum number of threads the server can have in the thread pool.
<b>Server Characteristics - Garbage Collection</b>	
ENABLED	Indication whether garbage collection is enabled.
INTERVAL	Garbage collection interval in seconds.
MEMORY USE	Garbage collection memory usage.
<b>Server Characteristics - Cache</b>	
CACHE	Cache flag.
MAX SIZE	Maximal cache size (KB).
MAX FILES	Maximal number of files in cache. NO is indicating that there is no maximum defined.
LIMIT 1	Cache limit 1.
LIMIT 2	Cache limit 2.
TIME MARGIN	Cache time margin (seconds).
KEEP EXPIRED	Keep expired flag.
CONNECT	Cache connect flag.
<b>Server Characteristics - Timeout Thresholds (in seconds)</b>	
INPUT	Input timeout.
OUTPUT	Output timeout.
SCRIPT	Script timeout.

## PP - HTTP

Table 173. Fields in the HTTP Server Details Report (continued)

Field Heading	Meaning
IDLE THREADS	Timeout for idle threads.
CACHE LOCK	Cache lock timeout.
<b>Server Activity - Requests</b> - The fields are given as COUNT and RATE (COUNT per second).	
REQUESTS	Requests that the HTTP server has successfully served.
REQUESTS DISCARDED	Requests sent to the HTTP server that are not valid.
REQUESTS IN ERROR	Requests that the HTTP server responded to with an error.
RESPONSES	Number of responses successfully sent.
RESPONSES DISCARDED	Responses the HTTP server was not able to send back to the client.
<b>Server Activity - Request Types</b>	
GET	Number of GET requests received by this server.
POST	Number of POST requests received by this server.
CGI	Number of CGI requests received by this server.
GWAPI	Number of GWAPI requests received by this server.
<b>Server Activity - Threads</b>	
MAX	Maximum number of threads as specified in the HTTP server configuration file on the MaxActiveThreads directive.
USED	Number of threads currently used.
NON-SSL WAITING	Number of non-Secure Sockets Layer (SSL) threads available for use. If this value is 0, all non-SSL threads are allocated.
SSL WAITING	Number of Secure Sockets Layer (SSL) threads available for use. If this value is 0, all SSL threads are allocated.
ASYNC I/O WAITING	If the HTTP server is running in Scalable Server mode, number of asynchronous I/O threads available for use. If this value is 0, all asynchronous I/O threads are allocated.
MSG QUEUE WAITING	If the HTTP server is running in Scalable Server mode, number of message queue threads available for use. If this value is 0, all message queue threads are allocated.
<b>Server Activity - Cache Usage</b>	
KBYTES READ	Number of kilobytes read from the cache of this server.
HITS	Number of requests for files stored in the cache of this server.
IN USE	Number of kilobytes of RAM used by the cache of this server.
FILES	Average number of files in the cache of this server.
<b>Server Activity - Miscellaneous</b>	
TIMEOUTS	Number of timeouts on the server. This value is not affected by any changes to the configuration of the server.
CONNECTIONS	Number of connections this server has provided.
<b>Server Activity - Throughput</b> - The fields are given as BYTES and RATE (BYTES per second).	
IN	Number of bytes sent to the HTTP server in requests.
OUT	Number of bytes sent by the HTTP server in responses.
UNKNOWN	Bytes that are not identified as part of a request.
<b>Server Activity - Response Times</b> - The values are given as minimum, maximum and average response time (in seconds).	
<b>Note:</b> These values refer to the complete server run time, not only to the current interval.	
DNS LOOKUP	Time it takes to complete the search for a domain name in the Domain Name Server (DNS).
SERVICE PLUGINS	Time it takes to complete customized application functions.
CGI	Time it takes to complete Common Gateway Interface (CGI) programs.
SSL HANDSHAKE	Time it takes to complete the exchange of security information between the HTTP server and browser.
PROXY RESPONSE	If configured as a Proxy Web server: time it takes to complete a transaction between a browser, this proxy server, and the destination server.

Table 173. Fields in the HTTP Server Details Report (continued)

Field Heading	Meaning																
Server Activity - Error Statistics - The number of responses with a specific error code.																	
ERROR	<table border="1"> <thead> <tr> <th>Code</th> <th>Meaning</th> </tr> </thead> <tbody> <tr> <td>200</td> <td>OK</td> </tr> <tr> <td>302</td> <td>Moved temporarily</td> </tr> <tr> <td>401</td> <td>Unauthorized</td> </tr> <tr> <td>403</td> <td>Forbidden</td> </tr> <tr> <td>404</td> <td>Not found</td> </tr> <tr> <td>407</td> <td>Proxy unauthorized</td> </tr> <tr> <td>500</td> <td>Internal server error</td> </tr> </tbody> </table>	Code	Meaning	200	OK	302	Moved temporarily	401	Unauthorized	403	Forbidden	404	Not found	407	Proxy unauthorized	500	Internal server error
Code	Meaning																
200	OK																
302	Moved temporarily																
401	Unauthorized																
403	Forbidden																
404	Not found																
407	Proxy unauthorized																
500	Internal server error																

## IOQ - I/O Queuing Activity report

The I/O Queuing Activity report provides information on the I/O configuration and activity rate, queue lengths, and percentages when one or more I/O components, grouped by a logical control unit (LCU), were busy.

For all channels that are managed by **Dynamic Channel Path Management (DCM)**, additional information is available. DCM allows an installation to identify channels which they wish to be managed dynamically. These channels are not assigned permanently to a specific control unit, but belong to a pool of channels. Based on workload requirements in the system, these channels will be assigned dynamically by DCM. For each LCU with DCM managed channels, a summary line displays the minimum and maximum number of connected DCM managed channels, the number of defined DCM managed channels and accumulated activity data.

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 - System Assist Processor), which is part of the channel subsystem, manages and schedules I/O work requests to the various devices within the LCU. If an I/O request is unsuccessful because the control unit is busy, the request is queued on the control unit header (CU-HDR) queue. Once the busy condition is resolved, the CU-HDR is then placed in the initiative queue.

PAV base mode is the mode when alias devices are assigned to one PAV base device. An I/O for a PAV base device is executed using aliases assigned to that PAV base device.

HyperPAV mode is the mode when a pool of alias devices is assigned to one LCU. An I/O for a PAV base device can be executed using any alias device of that pool.

SuperPAV mode is the mode when a pool of alias devices is assigned to one LCU and multiple LCUs are grouped into one Alias Management Group (AMG). An I/O for a PAV base device can be executed using any alias device of these multiple alias pools. The favored way is to use the alias device assigned to the same LCU (home LCU) that the PAV base device is assigned to.

Your installation defines your I/O configuration as input to the input/output configuration program (IOCP). The IOCP uses the information you supply to define the relationship between channel paths, control units, and I/O devices. The IOCP generates and assigns LCU identifiers to these groups of channel paths, control units and I/O devices. The IOCP then places this configuration definition

in a configuration data set (IOCDS). RMF uses the configuration definition as well as measurement data gathered during the interval to generate the I/O Queuing Activity report.

## How to request this report

To gather data for this report, specify as a Monitor I gatherer option:

```
IOQ(option)
```

To produce this report, specify

```
REPORTS(IOQ)
```

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

### Example URL for the DDS API

```
http://ddshost:8803/gpm/rmfpp.xml?reports=IOQ
```

## Using the information given in the report

If the Channel Path Activity and I/O Device Activity reports have shown that a problem exists, you can use the information in the I/O Queuing Activity report to pinpoint the reason for contention delays associated with channel paths, control units, and devices. For example, if the I/O Device Activity report shows an unusually large pending time for one or more devices in an LCU, the I/O Queuing Activity report indicates what proportion of the delay is caused by control unit busy and device busy. This proportion indicates which part of the configuration might need adjustment.

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

## Data gathering considerations

The report depends on information in the I/O configuration data set (IOCDS). If RMF cannot read the IOCDS, or if the IOCDS has been updated so that the data might not apply to the present configuration, no report is available. For example, when the operator partitions the system in such a way that RMF cannot read the IOCDS because it appears in another partition of a multi-processing system, RMF terminates the I/O Queuing Activity report and issues a message to the operator, I/O QUEUING ACTIVITY RMF REPORT TERMINATED.

## Missing data in report fields

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

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

the channel path as OFFLINE. If a channel path was taken offline or brought online during an interval, the additional line identifies the channel path as either NOW OFFLINE or NOW ONLINE.

When RMF **cannot obtain valid hardware data** for CONTENTION RATE and DELAY Q LENGTH, it prints the message NO H/W DATA under those headings.

If the **channel measurement facility is inactive** or has been interrupted during the interval, CHANNEL MEASUREMENT FACILITY NOT ACTIVE OR INTERRUPTED appears after the headings where the data normally appears in the report.

If the **diagnosis interface fails** during the interval, DIAGNOSIS INTERFACE FAILURE appears after the headings in the report.

## Messages

During the measurement of I/O Queuing activity, you may see one of the following messages in the data line:

### **LCU DYNAMICALLY CHANGED**

A LCU was dynamically changed during the interval.

### **LCU DYNAMICALLY ADDED**

A LCU was dynamically created during the interval.

### **LCU CHANGE ATTEMPTED**

A configuration change was attempted, but did not complete successfully.

## Contents of the report

I/O QUEUING ACTIVITY															
z/OS V2R2			SYSTEM ID SYS1			DATE 09/28/2016			PAGE 1						
			RPT VERSION V2R2 RMF			TIME 16.30.00			INTERVAL 15.00.036 CYCLE 1.000 SECONDS						
TOTAL SAMPLES = 900 IODF = 26 CR-DATE: 09/14/2016 CR-TIME: 16.09.06 ACT: ACTIVATE															
----- INPUT/OUTPUT PROCESSORS -----															
IOP	- INITIATIVE QUEUE -		IOP UTILIZATION			-- % I/O REQUESTS RETRIED --					----- RETRIES / SSCH -----				
	ACTIVITY RATE	AVG Q LENGH	% IOP BUSY	I/O START RATE	INTERRUPT RATE	ALL	CP BUSY	DP BUSY	CU BUSY	DV BUSY	ALL	CP BUSY	DP BUSY	CU BUSY	DV BUSY
00	2496.504	0.01	7.96	2496.503	1715.900	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
01	2496.866	0.00	4.18	2496.863	7853.703	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
02	2496.654	0.00	2.52	2496.653	839.133	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
03	2496.709	0.00	1.82	2496.708	317.149	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
17	2496.655	0.00	1.61	2496.654	172.826	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
SYS	59919.54	0.00	2.85	59919.66	61455.93	0.7	0.7	0.0	0.0	0.0	0.01	0.01	0.00	0.00	0.00
----- ALIAS MANAGEMENT GROUPS -----															
AMG	DCM GROUP	CHAN	CHPID	% DP	% CU	AVG CUB	AVG CMR	CONTENTION	DELAY Q	AVG CSS	HPAV	AVG OPEN	DATA XFER		
	MIN MAX DEF	PATHS	TAKEN	BUSY	BUSY	DLY	DLY	RATE	LNTH	DLY	WAIT MAX	EXCH	CONC		
00000016		65	0.537	0.00	0.00	0.0	3.6								
		34	0.538	0.00	0.00	0.0	4.5								
		99	0.532	0.00	0.00	0.0	1.3								
		*	1.607	0.00	0.00	0.0	3.1	0.000	0.00	0.8	0.000	0			
00000017		65	0.361	0.00	0.00	0.0	2.1								
		34	0.356	0.00	0.00	0.0	3.8								
		99	0.357	0.00	0.00	0.0	1.2								
		*	1.073	0.00	0.00	0.0	2.4	0.000	0.00	0.7	0.000	0			
----- LOGICAL CONTROL UNITS -----															
LCU/AMG	CU	DCM GROUP	CHAN	CHPID	% DP	% CU	AVG CUB	AVG CMR	CONTENTION	DELAY Q	AVG CSS	HPAV	AVG OPEN	DATA XFER	
		MIN MAX DEF	PATHS	TAKEN	BUSY	BUSY	DLY	DLY	RATE	LNTH	DLY	WAIT MAX	EXCH	CONC	
0049	5100		43	0.026	0.00	0.00	0.0	0.0							
			55	0.024	0.00	0.00	0.0	0.0							
			44	0.024	0.00	0.00	0.0	0.0							
			56	0.026	0.00	0.00	0.0	0.0							
			*	0.100	0.00	0.00	0.0	0.0	0.000	0.00	0.1	0.000	0	0.00 0.00	
0107	B101		65	0.084	0.00	0.00	0.0	3.6							
00000016			34	0.083	0.00	0.00	0.0	4.7							
			99	0.082	0.00	0.00	0.0	1.3							
			*	0.250	0.00	0.00	0.0	3.2	0.000	0.00	0.8	0.000	0		
011A	CF01		65	0.072	0.00	0.00	0.0	3.8							
00000016			34	0.073	0.00	0.00	0.0	4.9							
			99	0.071	0.00	0.00	0.0	1.2							
			*	0.217	0.00	0.00	0.0	3.4	0.000	0.00	0.7	0.000	0		
0106	B001		65	0.084	0.00	0.00	0.0	2.1							
00000017			34	0.082	0.00	0.00	0.0	4.1							
			99	0.083	0.00	0.00	0.0	1.2							
			*	0.250	0.00	0.00	0.0	2.5	0.000	0.00	0.8	0.000	0		

Figure 206. I/O Queuing Activity Report

The I/O Queuing Activity Report contains three sections with these titles:

### INPUT/OUTPUT PROCESSORS

This section shows the measurements accumulated for I/O processors.



**ALIAS MANAGEMENT GROUPS**

For each defined Alias Management Group (AMG), this section shows performance measurements for all channel paths connected to the LCUs grouped into the AMG.

**LOGICAL CONTROL UNITS**

For each LCU having online devices, this section shows performance measurements for all channel paths connected to the LCU.

Table 174. Fields in the I/O Queuing Activity Report

Field Heading	Meaning
IODF = <i>xx</i>	The IODF number where <i>xx</i> is the suffix of the IODF data set name.
CR-DATE: <i>mm/dd/yyyy</i>	The creation date of the IODF.
CR-TIME: <i>hh.mm.ss</i>	The creation time of the IODF.
ACT: <i>text</i>	The configuration state where <i>text</i> indicates how the IODF was activated.
INPUT/OUTPUT PROCESSORS	
IOP	<p>The two-digit hexadecimal identifier of the I/O processor (IOP). The IOP data sections are sorted according ascending IOP numbers.</p> <p>Following the last IOP data line is a line that summarizes the measurement data of the individual IOPs. This summary line starts with the character string <i>SYS</i>, indicating that it contains system wide information.</p>
INITIATIVE QUEUE	<p><b>ACTIVITY RATE</b></p> <p>The rate at which I/O requests are placed on the IOP initiative queue. There is one initiative queue for each IOP, and this value reflects the load of I/O requests on each IOP. This rate may be greater than the actual I/O rate due to potential re-queues.</p> $\text{ACTIVITY RATE} = \frac{\# \text{ I/O Requests on the IOP Queue}}{\text{Interval}}$ <p><b>AVG Q LENGH</b></p> <p>The average number of entries on the initiative queue for this IOP. Each time a request is added to the initiative queue, the new queue length is added to an accumulator.</p> $\text{AVG Q LENGH} = \frac{\text{Accumulated Queue Length}}{\# \text{ I/O Requests on the IOP Queue}} - 1$

Table 174. Fields in the I/O Queuing Activity Report (continued)

Field Heading	Meaning
IOP UTILIZATION	<p><b>% IOP BUSY</b>                      The ratio of the number of times the IOP was found busy to the total number of I/O processor samples.</p> $\% \text{ IOP BUSY} = \frac{\# \text{ Busy samples}}{\# \text{ Busy samples} + \# \text{ Idle samples}} * 100$ <p>A high IOP utilization might be caused by a high level of activity in terms of SSCH, I/O or sysplex operations per second or by contention in the I/O configuration. If contention is caused by CP BUSY or CU BUSY conditions, the request is placed on the IOP queue. This is indicated by an AVG Q LENGTH value greater than zero. If contention is caused by DP BUSY conditions, this is not indicated by the AVG Q LENGTH value, because the requests are kept internally. When the IOP is idle, these requests are processed which is reflected by the %IOP BUSY field.</p> <p><b>I/O START RATE</b>                      The rate at which I/O functions are initially started on this IOP. The value reflects the load of I/O requests on each IOP. It can be compared with the I/O rate in the device activity, or, the CHPID taken rate in the I/O queuing activity reports.</p> $\text{I/O START RATE} = \frac{\# \text{ I/O functions started}}{\text{Interval}}$ <p><b>INTERRUPT RATE</b>                      The rate at which I/O interrupts have been processed on this IOP. This value may be greater than the I/O start rate because it includes also the PCI interrupts.</p> $\text{INTERRUPT RATE} = \frac{\# \text{ Processed I/O interrupts}}{\text{Interval}}$

Table 174. Fields in the I/O Queuing Activity Report (continued)

Field Heading	Meaning
% I/O REQUESTS RETRIED	<p><b>ALL</b> The ratio of the number of retries to the number of I/O functions initially started plus the total number of retries.</p> $\%ALL = \frac{\# \text{ Retries}}{\# \text{ I/O functions started} + \# \text{ Retries}} * 100$ <p><b>CP BUSY</b> The ratio of the number of I/O operations retried on the I/O processor because the selected channel path was busy, to the number of I/O functions initially started plus the total number of retries.</p> $\%CP \text{ BUSY} = \frac{\# \text{ Retries due to channel path busy}}{\# \text{ I/O functions started} + \# \text{ Retries}} * 100$ <p><b>DP BUSY</b> The ratio of the number of times an I/O operation to a device was retried on the I/O processor because a director port on the path to that device was busy to the number of I/O functions initially started plus the total number of retries.</p> $\%DP \text{ BUSY} = \frac{\# \text{ Retries due to director port busy}}{\# \text{ I/O functions started} + \# \text{ Retries}} * 100$ <p><b>CU BUSY</b> The ratio of the number of times an I/O operation was retried on the I/O processor because the control unit of the targeted device was busy to the number of I/O functions initially started plus the total number of retries.</p> $\%CU \text{ BUSY} = \frac{\# \text{ Retries due to control unit busy}}{\# \text{ I/O functions started} + \# \text{ Retries}} * 100$ <p><b>DV BUSY</b> The ratio of the number of times an I/O operation was retried on the I/O processor because the targeted device was busy to the number of I/O functions initially started plus the total number of retries.</p> $\%DV \text{ BUSY} = \frac{\# \text{ Retries due to device busy}}{\# \text{ I/O functions started} + \# \text{ Retries}} * 100$

Table 174. Fields in the I/O Queuing Activity Report (continued)

Field Heading	Meaning
RETRIES / SSCH	<p><b>ALL</b> The ratio of the number of retries on the I/O processor to the number of I/O functions initially started.</p> $ALL = \frac{\text{\# Retries}}{\text{\# I/O functions started}}$ <p><b>CP BUSY</b> The ratio of the number of retries on the I/O processor because the selected channel path was busy to the number of I/O functions initially started.</p> $CP\ BUSY = \frac{\text{\# Retries due to channel path busy}}{\text{\# I/O functions started}}$ <p><b>DP BUSY</b> The ratio of the number of retries on the I/O processor because a director port on the path to that device was busy to the number of I/O functions initially started.</p> $DP\ BUSY = \frac{\text{\# Retries due to director port busy}}{\text{\# I/O functions started}}$ <p><b>CU BUSY</b> The ratio of the number of retries on the I/O processor because the control unit of the targeted device was busy to the number of I/O functions initially started.</p> $CU\ BUSY = \frac{\text{\# Retries due to control unit busy}}{\text{\# I/O functions started}}$ <p><b>DV BUSY</b> The ratio of the number of retries on the I/O processor because the targeted device was busy to the number of I/O functions initially started.</p> $DV\ BUSY = \frac{\text{\# Retries due to device busy}}{\text{\# I/O functions started}}$
<b>ALIAS MANAGEMENT GROUPS</b>	
AMG	The eight-digit hexadecimal system Alias Management Group assigned by I/O Supervisor
<b>LOGICAL CONTROL UNITS</b>	
LCU/ AMG	<p><b>LCU/</b> The four-digit hexadecimal identifier of the Logical Control Unit (LCU).  <b>AMG</b> The eight-digit hexadecimal system Alias Management Group assigned by I/O Supervisor, if the LCU is grouped to an AMG.</p> <p>An LCU is the logical representation of a physical control unit or a group of physical control units with one or more devices in common. Each physical control unit and each device can belong to only one LCU; they cannot be shared between LCUs.</p> <p>To find the LCU number, RMF must access the I/O configuration data set (IOCDs). If RMF cannot read it, or if it has been updated so that the data might not apply to the present configuration, RMF ends the I/O Queuing Activity report. If no activity has occurred during the interval for all selected LCUs, the message NO ACTIVITY FOR SELECTED LCUs appears instead of the data after the headings of the report.</p>
CU	The four-digit hexadecimal identifier of each physical control unit contained in the logical control unit.
<b>Note:</b> The following fields apply for both Alias Management Groups and Logical Control Units.	

Table 174. Fields in the I/O Queuing Activity Report (continued)

Field Heading	Meaning
DCM GROUP MIN - MAX - DEF	<p>The values in columns MIN MAX DEF report the minimum and maximum number of DCM managed channels for one LCU/AMG (in this interval) as well as the installation-specified definition for this LCU/AMG.</p> <p>The line with these values is available only for LCUs/AMGs with DCM managed channels. It contains in addition the accumulated values of the I/O activity rate, the director port contention, and the control unit contention of all DCM managed channels. These values may include also measurements of managed channels which were partially online.</p>
CHAN PATHS	<p>The two-digit hexadecimal channel path identifiers (CHPIDs) of the channel paths that are attached to the physical control units contained in the LCU/AMG. There can be up to eight channel paths in a logical control unit. The channel paths that are offline or moved online or offline during the interval are indicated as follows:</p> <p>OFFLINE NOW OFFLINE NOW ONLINE</p> <p>Channel paths that are online to the system but that might or might not be connected during the interval to any device in an LCU are indicated as follows:</p> <p>PATH OFFLINE PATH NOW OFFLINE PATH NOW ONLINE</p> <p>An '*' in this column indicates a summary line for all channel paths connected to the same LCU/AMG.</p> <p>If the control unit supports channel path attributes, RMF displays them together with the channel path:  <b>PF</b> preferred path  <b>NP</b> non-preferred path  <b>NS</b> path not specified</p> <p>In the following cases, RMF cannot find channel path attributes and therefore only displays the CHPID:</p> <ul style="list-style-type: none"> <li>• for devices residing in control units that do not support path attributes</li> <li>• for offline channels</li> <li>• for summary lines.</li> </ul>
CHPID TAKEN	<p>The rate at which I/O requests to devices of this LCU/AMG are satisfied by each CHPID during the interval. By reviewing the rate at which each channel path of the LCU/AMG satisfies I/O requests, you can see how evenly the work requests are distributed among the available paths and how effectively those paths are arranged for the LCU/AMG.</p> $\text{CHPID TAKEN} = \frac{\# \text{ I/O Operations Accepted on that Path}}{\text{Interval}}$ <p><b>Note:</b> If vary activity has occurred during the interval, this field is blank.</p>
% DP BUSY	<p>The ratio of the number of times an I/O request was deferred because the director port was busy to the number of attempts to service I/O requests during the measurement interval. This field indicates director port contention.</p> $\% \text{ DP BUSY} = \frac{\text{DPB}}{\text{DPB} + \text{CUB} + \text{SUC}} * 100$ <p><b>DPB</b> Number of deferred I/O requests due to director port busy  <b>CUB</b> Number of deferred I/O requests due to control unit busy  <b>SUC</b> Number of successful I/O requests on that path</p>

Table 174. Fields in the I/O Queuing Activity Report (continued)

Field Heading	Meaning
% CU BUSY	<p>The ratio of the number of requests deferred due to control unit busy to the number of attempts to service I/O requests during the measurement interval. This field indicates control unit contention and is reported for each path within the LCU/AMG.</p> $\% \text{ CU BUSY} = \frac{\text{CUB}}{\text{DPB} + \text{CUB} + \text{SUC}} * 100$ <p><b>DPB</b> Number of deferred I/O requests due to director port busy  <b>CUB</b> Number of deferred I/O requests due to control unit busy  <b>SUC</b> Number of successful I/O requests on that path</p> <p>RMF reports a value even if the channel path changes status during the interval.</p>
AVG CUB DLY	<p>The average number of milliseconds of delay that an I/O request encountered for the channel path because the control unit was busy.</p> $\text{AVG CUB DLY} = \frac{\text{Control Unit Busy Time}}{\# \text{ I/O Operations Accepted on that Path}}$
AVG CMR DLY	<p>The average number of milliseconds of delay that a successfully initiated start or resume function needs until the first command is indicated as accepted by the device. It allows to distinguish between real H/W errors versus workload spikes (contention in the fabric and at the destination port).</p> $\text{AVG CMR DLY} = \frac{\text{Initial Command Response Time}}{\# \text{ I/O Operations Accepted on that Path}}$
CONTENTION RATE	<p>The rate at which the I/O processor places delayed I/O requests on the CU-HDR for this LCU/AMG. The IOP places an I/O request on the CU-HDR when all paths to the subchannel are busy and at least one path to the control unit is busy. For devices with only one path or for devices where multiple paths exist and the busy condition is resolved immediately over an alternate path, the IOP does not count the condition.</p> $\text{CONTENTION RATE} = \frac{\# \text{ Enqueued Requests}}{\text{Interval}}$
DELAY Q LENGH	<p>The average number of delayed requests on the control unit header (CU-HDR). Each time a request is enqueued on the CU-HDR, RMF counts the number of requests on the queue and adds that number to the accumulator.</p> $\text{DELAY Q LENGH} = \frac{\text{Accumulated Queue Length}}{\# \text{ Enqueued Requests}} - 1$
AVG CSS DLY	<p>The average number of milliseconds of delay that an I/O request encountered after the acceptance of the start or resume function at the subchannel for the LCU/AMG, until the channel subsystem first attempts to initiate the operation.</p> $\text{AVG CSS DLY} = \frac{\text{Channel Subsystem Time}}{\# \text{ I/O Operations Accepted}}$
HPAV WAIT	<p>The ratio of the number of I/O requests that could not start because no HyperPAV aliases were available, to the total number of I/O requests for an LCU/AMG:</p> $\text{HPAV Wait} = \frac{\text{I/Os that could not start}}{\text{Total I/Os}}$
HPAV MAX	<p>The maximum number of concurrently used HyperPAV alias devices (including borrowed aliases) for that LCU/AMG during the interval.</p>
AVG OPEN EXCH	<p>The estimated average number of concurrently active I/O operations is provided in the LCU/AMG summary line if at least one FICON channel is connected to the LCU/AMG.</p> $\text{AVG OPEN EXCH} = \frac{\text{CMR+CONN+DISC}}{\text{RMF interval}}$ <p><b>CMR</b> initial command response time  <b>CONN</b> connect time  <b>DISC</b> disconnect time</p>

Table 174. Fields in the I/O Queuing Activity Report (continued)

Field Heading	Meaning
DATA XFER CONC	The data transfer concurrency is provided in the LCU/AMG summary line if at least one FICON channel is connected to the LCU/AMG.  CONN DATA XFER CONC = ----- RMF interval

## Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following table shows the overview condition names for the Overview report.

Table 175. Overview names in the I/O Queuing Activity Report

Field Heading or Meaning	Overview Name
ACTIVITY RATE	IOPAC
AVG Q LENGH	IOPQL
% IOP BUSY	IOPIPB
Percent I/O processor idle	IOPIPI
I/O START RATE	IORIFS
INTERRUPT RATE	IORPII
% I/O REQU RETRIED (ALL)	IOPALB
% I/O REQU RETRIED (CP BUSY)	IOPCHB
% I/O REQU RETRIED (DP BUSY)	IOPDPB
% I/O REQU RETRIED (CU BUSY)	IOPCUB
% I/O REQU RETRIED (DV BUSY)	IOPDVB
RETRIES / SSCH (ALL)	IONALB
RETRIES / SSCH (CP BUSY)	IONCHB
RETRIES / SSCH (DP BUSY)	IONDPB
RETRIES / SSCH (CU BUSY)	IONCUB
RETRIES / SSCH (DV BUSY)	IONDVB
CHPID TAKEN	IOART
% DP BUSY	IODPB
% CU BUSY	IOCUB
CONTENTION RATE	IOCTR
DELAY Q LENGH	IODLQ
AVG CUB DLY	IOCBT
AVG CMR DLY	IOCMR
AVG CSS DLY	IOCSS
HPAV WAIT	IOHWAIT
HPAV MAX	IOHMAX

## OMVS - OMVS Kernel Activity report

The OMVS Kernel Activity report provides information about:

- OMVS System Call Activity
- OMVS Process Activity
- OMVS Inter-Process Communication
- OMVS Memory Map - Shared Library Regions - Queued Signals

## How to request this report

Monitor III gathers data for this report automatically. If you want to suppress gathering, you have to disable writing SMF record type 74.3.

To produce this report, specify  
REPORTS(OMVS)

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

### Example URL for the DDS API

<http://ddshost:8803/gpm/rmfpp.xml?reports=OMVS>

## Contents of the report

The OMVS Kernel Activity report has these parts:

- OMVS System Call Activity
- OMVS Process Activity
- OMVS Inter-Process Communication
- OMVS Memory Map - Shared Library Regions - Queued Signals

OMVS KERNEL ACTIVITY												PAGE 1
z/OS V2R2		SYSTEM ID AOTS		DATE 09/28/2016		INTERVAL 30.00.000						
		RPT VERSION V2R2 RMF		TIME 13:00:00		CYCLE 1.000 SECONDS						
TOTAL SAMPLES = 1,800												
OMVS SYSTEM CALL ACTIVITY												
	MINIMUM	AVERAGE	MAXIMUM									
SYSCALLS (N/S)	23.5	2300*	5699									
CPU TIME (H/S)	16	47*	88									
OMVS PROCESS ACTIVITY												
MAXIMUM (TOT)	PROCESSES 1200			USERS 50			PROCESSES PER USER 12					
	MINIMUM	AVERAGE	MAXIMUM	MINIMUM	AVERAGE	MAXIMUM	MINIMUM	AVERAGE	MAXIMUM			
CURRENT (TOT)	99	854	1200	12	13	23	0	1.0*	5.3			
OVERRUNS (N/S)	0	5.5*	333	0	1.5*	4.8	0	1.0*	5.3			
OMVS INTER-PROCESS COMMUNICATION												
MAXIMUM (TOT)	MESSAGE QUEUE IDS 500			SEMAPHORE IDS 500			SHARED MEMORY IDS 500			SHARED MEMORY PAGES 262144		
	MINIMUM	AVERAGE	MAXIMUM	MINIMUM	AVERAGE	MAXIMUM	MINIMUM	AVERAGE	MAXIMUM	MINIMUM	AVERAGE	MAXIMUM
CURRENT (TOT)	100	300	500	100	300	500	100	300	500	100	131072	262144
OVERRUNS (N/S)	0	10*	100	0	10*	100	0	10*	100	0	10*	100
OMVS MEMORY MAP												
MAXIMUM (TOT)	MEMORY MAP STORAGE PAGES 4096			SHARED STORAGE PAGES 131K			MAX SHARED LIBRARY REGION 16M			MAXIMUM QUEUED SIGNALS 100K		
	MINIMUM	AVERAGE	MAXIMUM	MINIMUM	AVERAGE	MAXIMUM	MINIMUM	AVERAGE	MAXIMUM	MINIMUM	AVERAGE	MAXIMUM
CURRENT (TOT)	0.000	0.000	0.000	1542	1541	1542	16M	16M	16M	0.000	50K	99K
OVERRUNS (N/S)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	50K	99K

Units: (TOT) = Total Value, (N/S) = Number per Second, (H/S) = Hundredth of seconds per Second

Figure 207. OMVS Kernel Activity Report

Most values in the report will be reported as MINIMUM, AVERAGE, and MAXIMUM.

All average values derived from accumulated fields are marked with '\*' if the OMVS kernel address space was reinstated during the interval. If the OMVS process limits (MAXIMUM line) have changed, they will be reported as '\*\*\*\*'.



## Field descriptions

Table 176. Fields in the OMVS Kernel Activity Report

Heading	Meaning
OMVS SYSTEM CALL ACTIVITY	
SYSCALLS (N/S)	Number of system calls per second processed by the OMVS kernel address space in this interval.
CPU TIME (H/S)	Time spent to process system calls in hundredths of seconds per second.
OMVS PROCESS ACTIVITY	
MAXIMUM PROCESSES USERS PROCESSES PER USER	Maximum number of processes, users, and processes per user defined by OMVS kernel address space initialization parameters (in Parmlib member BPXPRMxx).  If one of these values has changed (due to an OMVS restart), it will be reported as '****'.
CURRENT PROCESSES	Number of OMVS processes controlled by OMVS during this interval.
CURRENT USERS	Number of OMVS users controlled by OMVS
OVERRUNS PROCESSES	Rate of processes that could not be created by OMVS because the maximum number of processes would have been exceeded.
OVERRUNS USERS	Rate of OMVS users that could not be created by OMVS because the maximum number of users would have been exceeded.
OVERRUNS PROCESSES PER USER	Rate of processes per user that could not be created by OMVS because the maximum number of processes per user would have been exceeded.
OMVS INTER-PROCESS COMMUNICATION	
MAXIMUM MESSAGE QUEUE IDS SEMAPHORE IDS SHARED MEMORY IDS SHARED MEMORY PAGES	Maximum number of message queue IDs, semaphore IDs, shared memory IDs, and shared memory pages defined by OMVS kernel address space initialization parameters (in Parmlib member BPXPRMxx).
CURRENT MESSAGE QUEUE IDS	Number of message queue IDs during this interval.
CURRENT SEMAPHORE IDS	Number of semaphore IDs during this interval.
CURRENT SHARED MEMORY IDS	Number of shared memory IDs during this interval.
CURRENT SHARED MEMORY PAGES	Number of shared memory pages during this interval.
OVERRUNS MESSAGE QUEUE IDS	Rate of message queue IDs that could not be created by OMVS because the maximum number of message queue IDs would have been exceeded.
OVERRUNS SEMAPHORE IDS	Rate of semaphore IDs that could not be created by OMVS because the maximum number of semaphore IDs would have been exceeded.
OVERRUNS SHARED MEMORY IDS	Rate of shared memory IDs that could not be created by OMVS because the maximum number of shared memory IDs would have been exceeded.
OVERRUNS SHARED MEMORY PAGES	Rate of shared memory pages that could not be created by OMVS because the maximum number of shared memory pages would have been exceeded.
OMVS MEMORY MAP	
MAXIMUM MEMORY MAP STORAGE PAGES SHARED STORAGE PAGES	Maximum number of memory map storage pages and shared storage pages defined by OMVS kernel address space initialization parameters (in Parmlib member BPXPRMxx).
CURRENT MEMORY MAP STORAGE PAGES	Number of memory map storage pages during this interval.
CURRENT SHARED STORAGE PAGES	Number of shared storage pages during this interval.
OVERRUNS MEMORY MAP STORAGE PAGES	Rate of memory map storage pages that could not be created by OMVS because the maximum number of memory map storage pages would have been exceeded.

Table 176. Fields in the OMVS Kernel Activity Report (continued)

Heading	Meaning
OVERRUNS SHARED STORAGE PAGES	Rate of shared storage pages that could not be created by OMVS because the maximum number of shared storage pages would have been exceeded.
SHARED LIBRARY REGION	
MAX SHARED LIBRARY REGION	Maximum amount of storage available for shared library region as specified by Parmlib statement SHRLIBRGNSIZE. The values are provided in units of megabytes.
CURRENT SHARED LIBRARY REGION	The current amount of storage in Megabytes available for shared library region.
OVERRUNS SHARED LIBRARY REGION	Rate of attempts to exceed the maximum storage amount for shared library region.
QUEUED SIGNALS	
MAX QUEUED SIGNALS	Maximum amount of queued signals allowed per process as specified by Parmlib statement MAXQUEUEDSIGS
OVERRUNS QUEUED SIGNALS	Rate of attempts to exceed the maximum number of queued signals.

## PAGESP - Page Data Set Activity report

The Page Data Set Activity report provides information about page data set usage for each individual data set. The information about the number of slots used is reported as minimum, maximum, and average values for the interval. Also, the time is provided when the Auxiliary Storage Manager (ASM) considered the data set to be busy, the number of start I/O requests initiated by ASM for the data set, the average page transfer time for each I/O request, and the number of pages transferred to and from the page data set.

The report contains only page data sets that:

- are in use at the end of that RMF measurement interval
- have been deleted during that RMF measurement interval

### How to request this report

Monitor I gathers data for this report automatically. If you want to suppress gathering, you need to specify NOPAGESP.

To produce this report, specify  
REPORTS(PAGESP)

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

### Example URL for the DDS API

<http://ddshost:8803/gpm/rmfpp.xml?reports=PAGESP>

### Using the information given in the report

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

The % IN USE field shows how busy the data set is. If this is above 30% you might see increases in response time. You might then:

- dedicate volumes to page data sets.
- make the sum of all the page space two to four times the number of slots used.
- limit use of VIO=YES

## Contents of the report

PAGE DATA SET ACTIVITY													PAGE 1
z/OS V2R2		SYSTEM ID SYS1		DATE 09/28/2016		INTERVAL 14.59.946							
		RPT VERSION V2R2 RMF		TIME 16.30.00		CYCLE 1.000 SECONDS							
NUMBER OF SAMPLES = 900		PAGE DATA SET AND SCM USAGE											
PAGE SPACE TYPE	VOLUME SERIAL	DEV NUM	DEVICE TYPE	SLOTS ALLOC	---- SLOTS USED ---	BAD SLOTS	% IN USE	PAGE TRANS TIME	NUMBER IO REQ	PAGES XFER'D	V I O	DATA SET NAME	
PLPA	PGT1B5	01B5	33903	180	168 168 168	12	0.00	0.000	0	0		SYS1.PGT1B5.PLPA	
COMMON	PGT1B5	01B5	33903	13500	7145 7148 7146	0	0.00	0.000	16	11		SYS1.PGT1B5.COMMON	
LOCAL	PGT1B5	01B5	33903	540000	40287 41912 41018	0	5.67	0.004	1933	12191	Y	SYS1.PGT1B5.LOCAL	
LOCAL	PGT80A	080A	33903	540000	42338 43947 43012	0	4.67	0.003	2069	13129	Y	SYS1.PGT80A.LOCAL	
LOCAL	PGT80C	080C	33903	540000	41388 43197 42128	0	6.22	0.004	2106	13299	Y	SYS1.PGT80C.LOCAL	
LOCAL	PGT857	0857	33903	540000	39245 40736 39858	0	5.33	0.004	1987	12575	Y	SYS1.PGT857.LOCAL	
LOCAL	PGT859	0859	33903	540000	39809 41273 40438	0	5.89	0.004	1969	12546	Y	SYS1.PGT859.LOCAL	
SCM	N/A	N/A	N/A	131072	58500 58501 58500	0	0.00	0.000	106	106	N/A	N/A	

Figure 208. PAGESP Report

**Data Not Available:** When a page data set comes online during a report interval, an asterisk is placed next to its name and the following message appears instead of measurement data: NOW AVAILABLE FOR SYSTEM USE.

When a page data set has been deleted during a report interval, an asterisk is placed next to its name and the following message appears instead of measurement data: DATA SET DELETED.

**Duration Report:** If you have specified a duration report, certain fields (DEV NUM, VOLUME SERIAL, DEVICE TYPE and SLOTS ALLOC) might be distorted due to a lengthy duration interval. When such a change occurs, it is not reflected in the duration report; these fields are set according to the contents of the first type 75 SMF record encountered.

Table 177. Fields in the Page Data Set Activity report

Field Heading	Meaning
PAGE SPACE TYPE	Page space type, which can be PLPA, COMMON, LOCAL, or SCM (Storage Class Memory).
VOLUME SERIAL	Volume serial number of the volume on which the data set resides. N/A is displayed for page space type SCM.
DEV NUM	Number of the device on which the data set resides. N/A for page space type SCM.
DEVICE TYPE	Device type on which the data set resides. N/A for page space type SCM.
SLOTS ALLOC	The total number of slots each page data set contains. For page space type SCM, the total number of SCM 4K blocks available to ASM is displayed.
SLOTS USED	The number of slots that were being used for paging. For page space type SCM, the total number of SCM 4K blocks in-use by ASM is displayed.
BAD SLOTS	Number of slots that encountered permanent I/O errors. For page space type SCM, the total number of SCM 4K blocks in error is displayed.

## PP - PAGESP

Table 177. Fields in the Page Data Set Activity report (continued)

Field Heading	Meaning
% IN USE	<p>Percentage of time during the reporting interval when the data set was considered busy by the Auxiliary Storage Manager (ASM).</p> <p>At each cycle, RMF tests each data set, and at the end of the interval, the percentage is calculated.</p> $\% \text{ IN USE} = \frac{\# \text{ Busy Samples}}{\# \text{ Samples}} * 100$
PAGE TRANS TIME	<p>Average number of seconds required to complete a page transfer.</p> $\text{PAGE TRANS TIME} = \frac{(\text{USE} * \text{INT}) / \text{N}}{\text{XFER}}$ <p><b>USE</b> Number of samples when the data set was in use  <b>XFER</b> Total number of pages transferred  <b>N</b> Number of samples  <b>INT</b> Interval time (seconds)</p>
NUMBER IO REQ	Total number of I/O requests for the data set made during the interval.
PAGES XFER'D	Number of pages that were transferred to or from the page data set or SCM in units of 4K pages.
VIO	<p>Indication of whether the local paging data set accepts VIO pages. The symbols are:</p> <p><b>Y</b> VIO pages are accepted  <b>N</b> VIO pages are not accepted</p>
DATA SET NAME	<p>Name of the page data set being monitored. A page data set name longer than 35 characters will be truncated to 35 characters in the report. The entire data set name appears in the SMF record.</p> <p><b>Note:</b></p> <ol style="list-style-type: none"> <li>1. If a data set was dynamically introduced during the interval, its data set name is preceded by an asterisk (*).</li> <li>2. When the operating system has detected errors in a data set that prevents its further use, the name of the data set is preceded by two asterisks (**). ASM continues to access the data set in read-only mode, and RMF reports this activity.</li> <li>3. N/A for page space type SCM.</li> </ol>

## Overview reference

Table 178. Overview names in the Page Data Set Activity report

Field Heading or Meaning	Overview Name
SLOTS USED - AVG	PSAVGSL
BAD SLOTS	PSBADS
% IN USE	PSBSY
PAGE TRANS TIME	PSPTT
NUMBER IO REQ	PSART
PAGES XFER'D	PSPT

## PAGING - Paging Activity report

The Paging Activity report provides information about the demands made on the system paging facilities and the use of central storage and external page storage during the interval.

### How to request this report

Monitor I gathers data for this report automatically. If you want to suppress gathering, you need to specify NOPAGING.

To produce this report, specify  
REPORTS(PAGING)

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

### Example URL for the DDS API

`http://ddshost:8803/gpm/rmfpp.xml?reports=PAGING`

## Using the information given in the report

If the non-swap, non-VIO page fault rate (page-ins) is excessively high, it could be the result of over-commitment of central storage.

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

## Contents of the report

The *Paging Activity* report is formatted into the following sections:

- CENTRAL STORAGE PAGING RATES
- CENTRAL STORAGE MOVEMENT AND REQUEST RATES
- FRAME AND SLOT COUNTS
- MEMORY OBJECTS AND HIGH VIRTUAL STORAGE FRAMES

The headers of the sections include the following information:

- The *OPT* field shows the name of the active option member IEAOPTxx. The option member contains parameters that affect system resource manager (SRM) decisions.
- The *LFAREA SIZE* field shows the size of the Large Frame Area in bytes (only available with Enhanced DAT architecture).

**Note:** The FRAME AND SLOT COUNTS section is displayed on one report page together with the CENTRAL STORAGE MOVEMENT AND REQUEST RATES section (Figure 210 on page 423), and therefore does not include the *OPT* or *LFAREA SIZE* fields.

### Central Storage Paging Rates

This section of the *Paging Activity* report monitors paging rates in central storage below the 2 GB bar. The paging rates monitored are organized into two major groups:

- Page-in rates
- Page-out rates.

The page-in and page-out groups are further divided into:

- Swap
- Non-swap (for the page-in group additionally divided into: Block, Non-Block)
- Total (rate and percentage)

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

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

## PP - PAGING

- Address space pages used, broken down into hiperspace data, VIO data and non-VIO data

The rate of page movement within central storage below the 2 GB bar is shown in the bottom left corner of the page.

PAGING ACTIVITY										
z/OS V2R2		SYSTEM ID R71			DATE 09/28/2016		INTERVAL 05.00.000			PAGE 1
		RPT VERSION V2R2 RMF			TIME 10.35.00		CYCLE 1.000 SECONDS			
OPT = IEAOPT00 LFAREA SIZE = 209715200 CENTRAL STORAGE PAGING RATES - IN PAGES PER SECOND										
CATEGORY	SWAP	PAGE IN				PAGE OUT				
		NON SWAP	NON BLOCK	TOTAL	NON SWAP	NON SWAP	TOTAL	NON SWAP		
		BLOCK	BLOCK	RATE	%			RATE	%	
PAGEABLE SYSTEM AREAS (NON-VIO)										
LPA		0.00	0.01	0.01	100					
CSA		0.00	0.00	0.00	0	0.00	0.00	0		
SUM		0.00	0.01	0.01	100	0.00	0.00	0		
ADDRESS SPACES										
HIPERSPACE		0.00		0.00	0	0.00	0.00	0		
VIO		0.00		0.00	0	0.00	0.00	0		
NON-VIO	0.00	0.00	0.00	0.00	0	0.00	0.00	0		
SUM	0.00	0.00	0.00	0.00	0	0.00	0.00	0.00	0	
TOTAL SYSTEM										
HIPERSPACE		0.00		0.00	0	0.00	0.00	0		
VIO		0.00		0.00	0	0.00	0.00	0		
NON-VIO	0.00	0.00	0.01	0.01	100	0.00	0.00	0.00	0	
SUM	0.00	0.00	0.01	0.01	100	0.00	0.00	0.00	0	
SHARED			0.00	0.00		0.00	0.00			
PAGE MOVEMENT WITHIN CENTRAL STORAGE				18.09						
PAGE MOVEMENT TIME %				0.0						
AVERAGE NUMBER OF PAGES PER BLOCK				0.0						
BLOCKS PER SECOND				0.00						
PAGE-IN EVENTS (PAGE FAULT RATE)				0.01						

Figure 209. PAGING Report - Central Storage Paging Rates

Table 179. Fields in the Paging Activity report - Central Storage Paging Rates

Field Heading	Meaning
CATEGORY	The component parts of paging rates identifying these basic components: <ul style="list-style-type: none"> <li>• Pageable system area, non-VIO data</li> <li>• Address space data</li> <li>• Total system data.</li> </ul>
PAGEABLE SYSTEM AREAS (NON-VIO)	The areas of central storage that are not associated with a single address space. This section consists of: <p><b>LPA</b> All values are reported except for swaps</p> <p><b>CSA</b> All values are reported except for swaps</p> <p><b>SUM</b> Sum of LPA and CSA.</p>
ADDRESS SPACES	The areas of central storage that are associated with individual address spaces. This section consists of: <p><b>HIPERSPACE</b></p> <p>All values are reported except for swaps</p> <p><b>VIO</b> All values are reported except for swaps</p> <p><b>NON-VIO</b></p> <p>All values are reported</p> <p><b>SUM</b> Sum of address space hiperspace, VIO and non-VIO.</p>

Table 179. Fields in the Paging Activity report - Central Storage Paging Rates (continued)

Field Heading	Meaning
TOTAL SYSTEM	<p>The sum of system pageable areas and address space values and the following:</p> <p><b>HIPERSPACE</b> Consists of address space hiperspace values</p> <p><b>VIO</b> Consists only of address space VIO values</p> <p><b>NON-VIO</b> Sum of system pageable areas non-VIO and memory non-VIO values</p> <p><b>SUM</b> Sum of system pageable areas sum and address space sum. (The computer system total for paging rates.)</p> <p><b>SHARED</b> Number of shared page group page-ins and page-outs in central storage. The page-in/out rate is included in the SUM values.</p>
PAGE IN	<p>The rate of pages read into central storage.</p> <p><b>SWAP</b> The rate of pages read into central storage as a result of address space swap-ins.  There is no PAGE IN for shared storage due to SWAP.</p> <p><b>NON SWAP/BLOCK</b> The rate of pages read into central storage from auxiliary storage exclusive of address space swap-ins.  Non-VIO paging occurs as a result of a page fault, PGLOAD, or PGFIX. When there are concurrent requests for the same page, only the first generates a page-in because all the requests will be satisfied by the same page.  A hiperspace page-in occurs when referencing a standard hiperspace page residing in auxiliary storage. VIO paging occurs as a result of a page fault or PGLOAD on a VIO window (logical GETs).  VIO pages that are swapped in are not included.  There is no BLOCK for shared storage.</p> <p><b>NON SWAP/NON BLOCK</b> The rate of pages read into central storage from auxiliary storage exclusive of address space swap-ins.  Non-VIO paging occurs as a result of a page fault, PGLOAD, or PGFIX. When there are concurrent requests for the same page, only the first generates a page-in because all the requests will be satisfied by the same page.  A hiperspace page-in occurs when referencing a standard hiperspace page residing in auxiliary storage. VIO paging occurs as a result of a page fault or PGLOAD on a VIO window (logical GETs).  VIO pages that are swapped in are not included.</p> <p><b>TOTAL RATE</b> The rate of the total system pages read into central storage. The rate is the sum of the non-swap page-in rate and the swap page-in rate.</p> <p><b>TOTAL %</b> The percentage of the total page-in rate for each part of the total.</p>

## PP - PAGING

Table 179. Fields in the Paging Activity report - Central Storage Paging Rates (continued)

Field Heading	Meaning
PAGE OUT	<p>The rate of pages written to auxiliary storage.</p> <p><b>SWAP</b> The rate of pages written to auxiliary storage as a result of address space swap outs.</p> <p>There is no PAGE OUT for shared storage due to SWAP.</p> <p><b>NON SWAP</b></p> <p>The rate of pages written to auxiliary storage (forced out) independent of address space swap outs. Non-VIO paging results from a PGOUT (including page stealing and other RSM-generated page-outs). VIO paging results from a PGOUT (including stealing and other RSM-generated page-outs) on a VIO window page (logical PUTs).</p> <p>Included also are the pages trimmed at swap out for logical swap out and from address spaces protected by central storage isolation by means of the IPS minimum working set size specification.</p> <p>A hiperspace page out occurs when a standard hiperspace page is no longer needed in central storage and is written to auxiliary storage.</p> <p>VIO pages transferred as a result of a swap-out are not included.</p> <p><b>TOTAL RATE</b></p> <p>The rate of total system pages written to auxiliary storage. The rate is the sum of the non-swap page-out rate and the swap page-out rate.</p> <p><b>TOTAL %</b></p> <p>The percentage of the total page-out rate for each part of the total.</p>
PAGE MOVEMENT WITHIN CENTRAL STORAGE	The rate of page movement within central storage. This includes each page movement from one frame to another frame independent of the location of the frame.
PAGE MOVEMENT TIME %	The percentage of general purpose processor time, including normalized AAP and IIP times, spent on page movement to obtain or free a particular type of frame for a page to be fixed (that is, a frame below the 16 megabyte line in central processor storage). The calculation is the amount of processor time needed to steal the page (including the time to move the contents of the frames, but not the time to move the new contents into the frames) divided by the length of the interval.
AVERAGE NUMBER OF PAGES PER BLOCK	The average size of address space non-VIO blocks that were paged-in during the interval. It does not include swap or hiperspace pages.
BLOCKS PER SECOND	The rate of page faults for pages that were part of a block.
PAGE-IN EVENTS (PAGE FAULT RATE)	The rate of page faults for all pages in events per second, excluding VIO and Hiperspace. The rate includes pages read from DASD only, not from expanded storage.

**Spreadsheet and Overview reference:** You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following table shows the overview condition names for the Overview report.

Table 180. Overview names in the Paging Activity report - Central Storage Paging Rates

Field Heading or Meaning	Overview Name
PAGE MOVEMENT WITHIN CENTRAL STORAGE	PGMVRT

### Central Storage Movement and Request Rates

The Central Storage Movement and Request Rates section provides paging information about hiperspace and VIO pages and about various types of storage requests.



PAGING ACTIVITY										
z/OS V2R2		SYSTEM ID TRX2		DATE 09/28/2016		INTERVAL 05.00.000				
		RPT VERSION V2R2 RMF		TIME 11.20.00		CYCLE 1.000 SECONDS				
OPT = IEAOPT00 LFAREA SIZE = 209715200 CENTRAL STORAGE MOVEMENT AND REQUEST RATES - IN PAGES PER SECOND										
-----										
SYSTEM UIC: MIN = 65535 MAX = 65535 AVG = 65535										
CENTRAL STORAGE PAGE WRITE PAGE READ ----- FRAME COUNTS -----										
-- RATE -- -- RATE -- -- MIN -- -- MAX -- -- AVG --										
HIPERSPACE 0.00 0.00 2 2 2										
VIO 0.00 0.00 0 0 0										
----- GETMAIN ----- FIXED ----- REF FAULTS -----										
STORAGE REQUESTS REQUESTS FRAMES BACKED REQ < 2GB FRAMES < 2GB 1ST NON-1ST										
RATE 596.37 145.45 2.74 91.51 704.03 0.00										
-----										
FRAME AND SLOT COUNTS										
-----										
(31 SAMPLES)										
CENTRAL STORAGE FRAMES TOTAL AVAILABLE SQA LPA CSA LSQA REGIONS+SWA HV SHARED HV COMMON										
-----										
MIN 1,310,720 39,907 9,632 19,762 9,178 32,416 986,636 176,503 37,173										
MAX 1,310,720 50,500 9,683 19,762 9,190 32,429 997,218 176,503 37,173										
AVG 1,310,720 48,323 9,669 19,762 9,184 32,420 988,825 176,503 37,173										
FIXED FRAMES TOTAL NUCLEUS SQA LPA CSA LSQA REGIONS+SWA <16 MB 16MB-2GB										
-----										
MIN 79,079 3,164 8,791 75 26,458 12,033 28,554 28 15,960										
MAX 89,457 3,164 8,838 75 26,458 12,044 38,886 28 16,074										
AVG 81,823 3,164 8,827 75 26,458 12,036 31,262 28 16,027										
SHARED FRAMES / SLOTS TOTAL CENTRAL STORAGE FIXED TOT FIXED BEL HV 1M HV 4K AUX DASD AUX SCM										
-----										
MIN 1,061,428 179,907 58 0 674 2,088 0 0										
MAX 1,061,529 180,008 58 0 674 2,088 0 0										
AVG 1,061,516 179,995 58 0 674 2,088 0 0										
LOCAL PAGE DATA SET SLOTS TOTAL AVAILABLE BAD NON-VIO VIO										
-----										
MIN 1,802,699 1,802,699 0 0 0										
MAX 1,802,699 1,802,699 0 0 0										
AVG 1,802,699 1,802,699 0 0 0										
SCM PAGING BLOCKS TOTAL AVAILABLE BAD IN-USE										
-----										
MIN 8,388,608 8,372,612 0 15,996										
MAX 8,388,608 8,372,612 0 15,996										
AVG 8,388,608 8,372,612 0 15,996										

Figure 210. PAGING Report - Central Storage Movement and Request Rates / Frame and Slot Counts

Table 181. Fields in the Paging Activity report - Central Storage Movement and Request Rates

Field Heading	Meaning
SYSTEM UIC (MIN, MAX, AVG)	The minimum, maximum and average system high unreferenced interval count. The maximum SYSTEM UIC value can indicate contention for central storage frames. When the SYSTEM UIC is relatively low, contention for central storage is high. Although total paging rates might vary with the type and level of workload, the unreferenced interval count is the best indicator of actual storage contention.
The Central Storage Movement and Request Rates section consists of two categories: CENTRAL STORAGE and STORAGE REQUESTS. The CENTRAL STORAGE category displays the following counts:	
PAGE WRITE RATE	<b>HIPERSPACE</b> Rate of hiperspace pages written to central storage. <b>VIO</b> Rate of VIO pages written to central storage.
PAGE READ RATE	<b>HIPERSPACE</b> Rate of hiperspace pages read from central storage. <b>VIO</b> Rate of VIO pages read from central storage.
FRAME COUNTS	MIN, MAX, and AVG of allocated frame counts. <b>HIPERSPACE</b> Storage frame counts allocated to hiperspace. <b>VIO</b> Storage frame counts allocated to VIO address space.
The STORAGE REQUESTS category displays the following counts:	

## PP - PAGING

Table 181. Fields in the Paging Activity report - Central Storage Movement and Request Rates (continued)

Field Heading	Meaning
GETMAIN	<b>REQUESTS</b> GETMAIN request rate  <b>FRAMES BACKED</b> Rate of pages backed during GETMAIN requests.
FIXED	<b>REQ &lt; 2 GB</b> Rate of fix requests issued for storage below 2 GB (address space only).  <b>FRAMES &lt; 2 GB</b> Rate of pages requested to be fixed for storage below 2 GB (address space only).
REF FAULTS	<b>1ST</b> First page reference faults rate.  <b>NON-1ST</b> Non-first page reference faults rate.

**Spreadsheet and Overview reference:** You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following table shows the overview condition names for the Overview report.

Table 182. Overview names in the Paging Activity report - Central Storage Movement and Request Rates

Field Heading or Meaning	Overview Name
SYSTEM UIC - MAX	MXHUIC
SYSTEM UIC - AVG	AVGHUIC
Overview names in the CENTRAL STORAGE category:	
PAGE WRITE RATE - HIPERSPACE	RSHSPW
PAGE WRITE RATE - VIO	RSVIOW
PAGE READ RATE - HIPERSPACE	RSHSPR
PAGE READ RATE - VIO	RSVIOR
FRAME COUNTS - HIPERSPACE - MIN	RSHSPM
FRAME COUNTS - HIPERSPACE - MAX	RSHSPX
FRAME COUNTS - HIPERSPACE - AVG	RSHSPA
FRAME COUNTS - VIO- MIN	RSVIOM
FRAME COUNTS - VIO- MAX	RSVIOX
FRAME COUNTS - VIO- AVG	RSVIOA

### Frame and Slot Counts

This section of the Paging Activity report (included in Figure 210 on page 423) shows information about the following storage related categories:

- CENTRAL STORAGE FRAMES
- FIXED FRAMES
- SHARED FRAMES
- LOCAL PAGE DATA SET SLOTS
- SCM PAGING BLOCKS

All values are presented as MIN (minimum), MAX (maximum), and AVG (average).

Table 183. Fields in the Paging Activity report - Frame and Slot Counts

Field Heading	Meaning
SAMPLES	<p>The number of valid samples taken in this interval is shown in the upper left corner of this report section in various formats:</p> <ul style="list-style-type: none"> <li>• as (nn SAMPLES) if all samples are valid</li> <li>• as SAMPLES = xx VALID SAMPLES = yy VALID SAMPLES CSA/REGION = zz</li> </ul> <p>if there are invalid samples and therefore the number of valid samples is less than the number of samples. In this case, the number of valid samples, and the number of valid samples for CSA and REGION+SWA values is also displayed to indicate that some of the CENTRAL STORAGE and FIXED FRAMES counts are based on less data.</p>
CENTRAL STORAGE FRAMES	<p><b>TOTAL</b> The total number of central storage frames in the system. <b>AVAILABLE</b> The number of central storage frames that are not in-use by the system. <b>SQA, LPA, CSA, LSQA, REGIONS+SWA, HV SHARED, HV COMMON</b> These columns show the number of central storage frames that are in-use by each of these areas. The REGIONS+SWA value also includes the number of frames used by high virtual private storage.</p> <p>The value of the TOTAL count is not a summation of the AVAILABLE, SQA, LPA, CSA, LSQA, REGIONS+SWA, HV SHARED (high virtual SHARED), and HV COMMON (high virtual COMMON) counts at the end of the interval, but is derived by adding these counts from each valid sample and then reporting the MIN sum, MAX sum, and AVG sum for the complete set of samples.</p> <p>If data is not available for any of the SQA, LPA, CSA, LSQA, or REGIONS+SWA counts, the following text appears across these columns: *** NO COUNTS AVAILABLE ***</p> <p>If there are no valid samples for CSA and REGION+SWA values, 'NO DATA' is displayed in these columns.</p> <p><b>Note:</b> The actual maximum or minimum value of a field might occur at a time when RMF is not sampling.</p>
FIXED FRAMES	<p><b>TOTAL</b> The total number of central storage frames in the system that are in-use by fixed pages. <b>NUCLEUS, SQA, LPA, CSA, LSQA, REGIONS+SWA, &lt;16MB, 16MB-2GB</b></p> <p>These columns show the number of central storage frames that are in-use by fixed pages allocated in each of these areas.</p> <p>The CSA value also includes frames used by fixed pages allocated in High Virtual Common (HV COMMON). The REGIONS+SWA value also includes the number of frames used by high virtual private storage.</p> <p>The SQA value also includes fixed CSA pages. However, pageable CSA pages that have been fixed after allocation are reported in the CSA category.</p> <p>If data is not available for any of the SQA, LPA, CSA, LSQA, or REGIONS+SWA counts, the following text appears across these columns: *** NO COUNTS AVAILABLE ***</p> <p>If there are no valid valid samples for CSA and REGION+SWA values, 'NO DATA' is displayed in these columns.</p>

## PP - PAGING

Table 183. Fields in the Paging Activity report - Frame and Slot Counts (continued)

Field Heading	Meaning
SHARED FRAMES / SLOTS	<p><b>TOTAL</b> The total number of central storage frames and auxiliary slots that are in-use by shared pages.</p> <p><b>CENTRAL STORAGE</b> The total number of central storage frames that are in-use by shared pages.</p> <p><b>FIXED TOT</b> The number of central storage frames that are in-use by shared fixed pages allocated below the 2 GB bar. <b>Note:</b> High virtual shared storage cannot be fixed.</p> <p><b>FIXED BEL</b> The number of central storage frames that are in-use by shared fixed pages allocated below 16 megabytes.</p> <p><b>HV 1M</b> The number of central storage frames that are in-use by shared high virtual 1 MB pages.</p> <p><b>HV 4K</b> The number of central storage frames that are in-use by shared high virtual 4K pages.</p> <p><b>AUX DASD</b> The number of shared pages backed on DASD.</p> <p><b>AUX SCM</b> The number of shared pages backed on Storage Class Memory (SCM). <b>Note:</b> AUX DASD and AUX SCM are also called auxiliary storage slots.</p>
LOCAL PAGE DATA SET SLOTS	<p><b>TOTAL</b> Total number of page data set slots.</p> <p><b>AVAILABLE</b> Number of page data set slots that do not contain any data pages and that are available for use.</p> <p><b>BAD</b> Number of local page data set slots that do not contain any data pages and are unavailable for use because of permanent I/O errors.</p> <p><b>NON-VIO</b> Number of local page data set slots that contain pages belonging to address-space virtual storage.</p> <p><b>VIO</b> Number of local page data set slots that contain pages for VIO data sets.</p>
SCM PAGING BLOCKS	<p><b>TOTAL</b> The total number of 4K SCM paging blocks.</p> <p><b>AVAILABLE</b> The number of SCM blocks that do not contain any data and are available to ASM.</p> <p><b>BAD</b> The number of SCM blocks that do not contain any data and are unavailable for use.</p> <p><b>IN-USE</b> The number of SCM blocks that are in-use by ASM.</p>

**Spreadsheet and Overview reference:** You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following table shows the overview condition names for the Overview report.

Table 184. Overview names in the Paging Activity report - Frame and Slot Counts

Field Heading or Meaning	Overview Name
CENTRAL STORAGE FRAMES - CSA - MAX	MXCSAT
CENTRAL STORAGE FRAMES - CSA - AVG	AVGCSAT
FIXED FRAMES - SQA - MAX	MXSQA
FIXED FRAMES - SQA - AVG	AVGSQA
FIXED FRAMES - CSA - MAX	MXCSAF
FIXED FRAMES - CSA - AVG	AVGCSAF
LOCAL PAGE DATA SET SLOTS - VIO - MAX	MAXVIOF
LOCAL PAGE DATA SET SLOTS - VIO - AVG	AVGVIOF
FIXED FRAMES - 16MB-2GB - MIN	FXBETWM
FIXED FRAMES - 16MB-2GB - MAX	FXBETWX
FIXED FRAMES - 16MB-2GB - AVG	FXBETWA

Table 184. Overview names in the Paging Activity report - Frame and Slot Counts (continued)

Field Heading or Meaning	Overview Name
SHARED FRAMES/SLOTS - TOTAL - AVG	SHRPT
SHARED FRAMES/SLOTS - CENTRAL STORAGE - AVG	SHRPC
SHARED FRAMES/SLOTS - FIXED TOT - AVG	SHRPF
SHARED FRAMES/SLOTS - FIXED BEL - AVG	SHRPB
SHARED FRAMES/SLOTS - HV 1 MB - AVG	SFR1MA
SHARED FRAMES/SLOTS - HV 4K - AVG	SFR4KA
SHARED FRAMES/SLOTS - AUX DASD - AVG	SHRPA
SHARED FRAMES/SLOTS - AUX SCM - AVG	SHRPASCM

## Memory Objects and High Virtual Storage Frames

Figure 211 shows a sample of the MEMORY OBJECTS AND HIGH VIRTUAL STORAGE FRAMES section if Enhanced DAT architecture is available. If Enhanced DAT architecture is not installed, no information about memory objects that can be backed in 1 MB frames is available. Also, no information about 1 MB frames is available and any fields in this section related to such memory objects or frames are not displayed.

All values in this section of the **Paging Activity** report are presented as MIN, MAX, and AVG values.

PAGING ACTIVITY									
z/OS V2R2	SYSTEM ID TRX2	DATE 09/28/2016	INTERVAL 05.00.000	PAGE 3					
	RPT VERSION V2R2 RMF	TIME 11.20.00	CYCLE 1.000 SECONDS						
OPT = IEAOPT00 LFAREA SIZE = 209715200 MEMORY OBJECTS AND HIGH VIRTUAL STORAGE FRAMES									
-----									
MEMORY OBJECTS	COMMON	FIXED 1M	SHARED	SHARED 1M					
-----	-----	-----	-----	-----					
MIN	97	1	4	2					
MAX	97	3	4	2					
AVG	97	2	4	2					
1 MB FRAMES	FIXED		PAGEABLE						
-----	TOTAL	AVAILABLE	IN-USE	TOTAL	AVAILABLE	IN-USE			
-----	-----	-----	-----	-----	-----	-----			
MIN	200	0	160	560	0	560			
MAX	200	40	200	560	0	560			
AVG	200	30	170	560	0	560			
HIGH SHARED FRAMES	TOTAL	CENTRAL STORAGE	BACKED 1M			AUX DASD	AUX SCM		
-----	-----	-----	-----			-----	-----		
MIN	136902.1M	176,503	674			0	0		
MAX	136902.1M	176,503	674			0	0		
AVG	136902.1M	176,503	674			0	0		
HIGH COMMON FRAMES	TOTAL	CENTRAL STORAGE	BACKED 1M	FIXED	FIXED 1M	AUX DASD	AUX SCM		
-----	-----	-----	-----	-----	-----	-----	-----		
MIN	17301504	37,173	70	12,542	30	0	0		
MAX	17301504	37,173	70	12,542	30	0	0		
AVG	17301504	37,173	70	12,542	30	0	0		

Figure 211. PAGING Report - Memory Objects and High Virtual Storage Frames

## PP - PAGING

Table 185. Fields in the Paging Activity report - Memory Objects and High Virtual Storage Frames

Field Heading	Meaning
MEMORY OBJECTS	<p><b>COMMON</b> Number of memory objects allocated in the high virtual common storage of the system.</p> <p><b>FIXED 1M</b> Number of fixed memory objects that are allocated in the system and can be backed in 1 MB frames.</p> <p><b>SHARED</b> Number of memory objects allocated in the high virtual shared storage of the system.</p> <p><b>SHARED 1M</b> Number of shared memory objects that are allocated in the system and can be backed in 1 MB frames.</p>
1 MB FRAMES - FIXED	<p><b>TOTAL</b> Total number of 1 MB frames, that can be used by fixed memory objects. This value is equal to the size of the Large Frame Area in megabytes.</p> <p><b>AVAILABLE</b> Number of 1 MB frames in the Large Frame Area that are not in-use.</p> <p><b>IN-USE</b> Number of 1 MB frames in the Large Frame Area that are in-use by fixed memory objects.</p>
1 MB FRAMES - PAGEABLE	<p><b>TOTAL</b> Total number of 1 MB frames that can be used by pageable and DREF memory objects..</p> <p><b>AVAILABLE</b> Number of 1 MB frames that are not in-use by pageable and DREF memory objects.</p> <p><b>IN-USE</b> Number of 1 MB frames that are in-use by pageable and DREF memory objects regardless of whether the frames are actually used for 1 MB pages or used to satisfy 4 KB space requests on a constrained system.</p>
HIGH SHARED FRAMES	<p><b>TOTAL</b> Size of high virtual shared area in units of 4 KB pages.</p> <p><b>CENTRAL STORAGE</b> Number of pages from high virtual shared storage that are backed in central storage (in units of 4 KB).</p> <p><b>BACKED 1M</b> Number of high virtual shared memory 1 MB pages that are backed in central storage</p> <p><b>AUX DASD</b> Number of auxiliary storage slots used for high virtual shared pages that are backed on DASD.</p> <p><b>AUX SCM</b> Number of auxiliary storage slots used for high virtual shared pages that are backed on SCM storage.</p>

Table 185. Fields in the Paging Activity report - Memory Objects and High Virtual Storage Frames (continued)

Field Heading	Meaning
HIGH COMMON FRAMES	<p><b>TOTAL</b> Size of high virtual common area in units of 4 KB pages.</p> <p><b>CENTRAL STORAGE</b> Number of pages from high virtual common storage that are backed in central storage (in units of 4 KB).</p> <p><b>BACKED 1M</b> Number of high virtual common memory 1 MB pages that are backed in central storage.</p> <p><b>FIXED</b> Number of pages from high virtual common storage that are fixed in central storage (in units of 4 KB).</p> <p><b>FIXED 1M</b> Number of high virtual common memory 1 MB pages that are fixed in central storage.</p> <p><b>AUX DASD</b> Number of auxiliary storage slots used for high virtual common pages that are backed on DASD.</p> <p><b>AUX SCM</b> Number of auxiliary storage slots used for high virtual common pages that are backed on SCM storage.</p>

**Spreadsheet and Overview reference:** You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following table shows the overview condition names for the Overview report.

Table 186. Overview names in the Paging Activity report - Memory Objects and High Virtual Storage Frames

Field Heading or Meaning	Overview Name
MEMORY OBJECTS - COMMON AVG	CMOA
MEMORY OBJECTS - FIXED 1M AVG	LMOA
MEMORY OBJECTS - SHARED AVG	SMOA
MEMORY OBJECTS - SHARED 1M AVG	SMO1MA
FRAMES - 1 MB AVG	LFRA
1 MB FRAMES - FIXED TOTAL AVG	LFRTA
1 MB FRAMES - FIXED IN-USE AVG	LFRRUA
1 MB FRAMES - PAGEABLE TOTAL AVG	LPFRTA
1 MB FRAMES - PAGEABLE AVAILABLE AVG	LPFRAA
1 MB FRAMES - PAGEABLE IN-USE AVG	LPFRUA
HIGH SHARED FRAMES - TOTAL AVG	SFRTA
HIGH SHARED FRAMES - CENTRAL STORAGE AVG	SFRA
HIGH SHARED FRAMES - BACKED 1M AVG	SFR1MA
HIGH SHARED FRAMES - AUX DASD AVG	SAUXSA
HIGH SHARED FRAMES - AUX SCM AVG	SAUXSSA
HIGH COMMON FRAMES - TOTAL AVG	CFRTA
HIGH COMMON FRAMES - CENTRAL STORAGE AVG	CFRA
HIGH COMMON FRAMES - BACKED 1M AVG	CFR1MA
HIGH COMMON FRAMES - FIXED AVG	CFFRA
HIGH COMMON FRAMES - FIXED 1M AVG	CFFR1MA
HIGH COMMON FRAMES - AUX DASD AVG	CAUXSA

Table 186. Overview names in the Paging Activity report - Memory Objects and High Virtual Storage Frames (continued)

Field Heading or Meaning	Overview Name
HIGH COMMON FRAMES - AUX SCM AVG	CAUXSSA

## PCIE - PCIE Activity Report

The PCIE Activity Report provides statistics and performance measurements on PCI Express based functions (PCIE functions) allocated by at least one z/OS address space for a period of time within the reporting interval. A PCIE function is captured by the report if one of the following feature activities has been detected:

- RDMA (Remote Direct Memory Access) over Converged Enhanced Ethernet
- zEnterprise Data Compression (zEDC) capability using zEDC Express
- SMC-Direct over Internal Shared Memory (ISM) virtual PCIE function.

### How to request this report

If the currently active SMFPRMxx parameter settings indicate that SMF record type 74 subtype 9 is to be collected, then RMF Monitor III gathers the data required for the PCIE Activity Report into this SMF record.

To produce this report, specify  
REPORTS(PCIE)

This single-system report is only available in XML output format. Therefore, you need to specify the XPRPTS ddname in your Postprocessor job. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

### Example URL for the DDS API

<http://ddshost:8803/gpm/rmfpp.xml?reports=PCIE>

### Contents of the report

The *PCIE Activity Report* is divided into three sections:

- **General PCIE Activity**
- **Hardware Accelerator Activity**
- **Hardware Accelerator Compression Activity**

The **General PCIE Activity** section shows measurements for all PCIE functions independent from the type of the exploited hardware feature. The measurements reflect the activity of the z/OS system on which RMF data collection took place. They comprise data rates about the communication of z/OS programs with PCIE functions by means of PCI operations that are transferring data blocks from z/OS to the PCIE function (PCI LOAD, PCI STORE, PCI STORE BLOCK, and REFRESH PCI TRANSLATIONS) as well as Read/Write Transfer data rates.

The **Hardware Accelerator Activity** section and the **Hardware Accelerator Compression Activity** section have single system scope and are leveraging the measurements displayed in the **General PCIE Activity** section. They are only displayed if the hardware feature zEnterprise Data Compression (zEDC) is used for compression acceleration. In this case, they display:

- common accelerator metrics, like for example, total request execution time, or the amount of transferred data



- compression specific metrics, like for example the amount of compressed data and the number and throuput of compression requests
- device driver buffer statistics.

**RMF Postprocessor Interval Report [System TA2] : PCIE Activity Report**

RMF Version : :IOS V2R2 SMF Data : :IOS V2R2  
Start : 12/18/2014-22:20:00 End : 12/18/2014-22:25:00 Interval : 05:00:000 minutes

**General PCIE Activity**

Function ID	Function CHID	Function Name	Function Type	Function Status	Owner Job Name	Owner Address Space ID	Function Allocation Time	PCI Load Operations Rate	PCI Store Operations Rate	PCI Store Block Operations Rate	Refresh PCI Translations Operations Rate	DMA Address Space Count	Read Transfer Rate	Write Transfer Rate	Packets Received Rate	Packets Transmitted Rate	Work Units Processed Rate	Adapter Utilization
006C	0204	Hardware Accelerator	1014044B	Allocated	FPGHWAM	0013	300	0	102	0	14.0	1					61701	<.001
007C	025C	Hardware Accelerator	1014044B	Allocated	FPGHWAM	0013	300	0	102	0	14.0	1					66363	<.001
00A2	013C	10GbE RoCE Express	15B31004	Allocated	TAM390	002C	300	0.113	5999	0	1.69	1	0.267	288	10315	144648		

**Hardware Accelerator Activity**

Function ID	Time Busy %	Request Execution Time	Std Dev for Request Execution Time	Request Queue Time	Std Dev for Request Queue Time	Request Size	Transfer Rate Total
006C	0.286	28.0	8.07	65.7	140	47.6	4.87
007C	0.304	29.7	19.4	80.0	138	47.9	4.90

**Hardware Accelerator Compression Activity**

Function ID	Compression Request Rate	Compression Throughput	Compression Ratio	Decompression Request Rate	Decompression Throughput	Decompression Ratio	Buffer Pool Size	Buffer Pool Utilization
006C	102	2.91	1.47	0.437	0.009	0.455	16	0
007C	0.437	0.033	2.92	102	1.98	0.877	16	0

Figure 212. Postprocessor PCIE Activity Report

Table 187. Fields in the PCIE Activity Report

Field Heading	Meaning
<b>General PCIE Activity</b>	
Function ID	Identifier of the monitored PCIE function.  This field also identifies applicable PCIE functions in the <b>Hardware Accelerator Activity</b> and <b>Hardware Accelerator Compression Activity</b> sections.
Function CHID	Physical or virtual channel identifier for the PCIE function.
Function Name	Device name for the PCIE function.
Function Status	The PCIE function status can be one of the following:  <b>Allocated</b> The function is allocated and in use at the end of the reporting interval.  <b>Re-Allocated</b> The function was de-allocated during the interval but has been re-allocated again. It is in use at the end of the reporting interval.  <b>De-Allocated</b> The function was de-allocated during the interval and is unused at the end of the reporting interval.  <b>De-Allocate-Pending</b> The function is in the process of de-allocation.  <b>Error</b> The function is in permanent error.  <b>Unknown</b> The function status is unknown.
Owner Job Name	Job name of the owner who allocated the PCIE function.
Owner Address Space ID	Address space ID of the owner who allocated the PCIE function.
Function Allocation Time	Time in seconds for which the PCIE function was allocated or de-allocate-pending during this interval.
PCI Load Operations Rate	Rate of PCI Load operations executed during the reporting interval.
PCI Store Operations Rate	Rate of PCI Store operations executed during the reporting interval.
PCI Store Block Operations Rate	Rate of PCI Store Block operations executed during the reporting interval.
Refresh PCI Translations Operations Rate	Rate of Refresh PCI Translations operations executed during the reporting interval.

## PP - PCIE

Table 187. Fields in the PCIE Activity Report (continued)

Field Heading	Meaning
DMA Address Space Count	Number of defined DMA address spaces.
Read Transfer Rate	The number of megabytes per second that a RoCE device received on the external Ethernet interface. On zEC12 or zBC12, this field designates the number of megabytes per second that were transferred by DMA reads from all defined DMA address spaces to the PCIE function.
Write Transfer Rate	The number of megabytes per second transmitted on a RoCE or SMC-D device. On zEC12 or zBC12, this field designates the number of megabytes per second that were transferred by DMA writes from the PCIE function to all defined DMA address spaces.
Packets Received Rate	Number of packets per second that were received on the external Ethernet interface of the RoCE device. This value is not reported on zEC12 and zBC12 hardware.
Packets Transmitted Rate	Number of packets per second that were transmitted on the external Ethernet interface of the RoCE device. This value is not reported on zEC12 and zBC12 hardware.
Work Units Processed Rate	Number of work units per second that were processed by the zEDC device. This value is not reported on zEC12 and zBC12 hardware.
Adapter Utilization	Utilization of the zEDC device. This value is not reported on zEC12 and zBC12 hardware.
Physical Network Id Port 1 and 2	Physical-network identifier (PNET ID) that identifies the first or second port of the RoCE device or virtual PCIE function.
<b>Hardware Accelerator Activity</b>	
Time Busy %	The percentage of time that this partition kept the hardware accelerator busy.
Request Execution Time	The average time in microseconds the hardware accelerator used to process a request.
Std Dev for Request Execution Time	The standard deviation of the request execution time.
Request Queue Time	The average queue time in microseconds that was spent for a request. This value has single system scope but is affected by activity from other partitions sharing the hardware accelerator.
Std Dev for Request Queue Time	The standard deviation of the request queue time.
Request Size	The average number of kilobytes transferred per request.
Transfer Rate Total	The number of megabytes per second transferred by DMA operations.
<b>Hardware Accelerator Compression Activity</b>	
Compression Request Rate	The number of compression requests per second.
Compression Throughput	The number of megabytes compressed per second.
Compression Ratio	The ratio between input and output bytes compressed within this interval.
Decompression Request Rate	The number of decompression requests per second.
Decompression Throughput	The number of megabytes decompressed per second.
Decompression Ratio	The ratio between input and output bytes decompressed within this interval.
Buffer Pool Size	The total size of memory in megabytes that is allocated to the buffer pool.
Buffer Pool Utilization	The average utilization of the buffer pool that z/OS kept for in-use buffers.

## Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following table shows the overview condition names for the Overview report.

Table 188. Overview conditions in the PCIE Activity Report

Field Heading or Meaning	Overview Name
PCI Load Operations Rate	PCILOAD
PCI Store Operations Rate	PCISTOR
PCI Store Block Operations Rate	PCISTBL

Table 188. Overview conditions in the PCIE Activity Report (continued)

Field Heading or Meaning	Overview Name
Refresh PCI Translations Operations Rate	PCIRPTR
DMA Read Rate (on zEC12 or zBC12 hardware only)	PCIDMAR
DMA Write Rate (on zEC12 or zBC12 hardware only)	PCIDMAW
Number of megabytes received per second (RoCE on z13 only and Virtual PCIE functions)	PCIBYTR
Number of megabytes transmitted per second (RoCE on z13 only)	PCIBYTT
Number of packets received per second (RoCE on z13 only)	PCIPAKR
Number of packets transmitted per second (RoCE on z13 only)	PCIPAKT
Number of work units processed per second (zEDC on z13 only)	PCIWUP
PCI Function Utilization (zEDC on z13 only)	PCIUTIL
<b>Hardware Accelerator Activity</b>	
Time Busy %	FPGBUSY
Request Execution Time	FPGRTIM
Request Queue Time	FPGQTIM
Request Size	FPGBYTR
Transfer Rate Total	FPGBYTS
<b>Hardware Accelerator Compression Activity</b>	
Compression Request Rate	FPGCORS
Compression Throughput	FPGCOBS
Compression Ratio	FPGCORT
Decompression Request Rate	FPGDCRS
Decompression Throughput	FPGDCBS
Decompression Ratio	FPGDCRT
Buffer Pool Size	FPGBPSZ
Buffer Pool Utilization	FPGBPRT

## SCM - SCM Activity Report

The SCM Activity Report provides statistics and performance measurements on Storage Class Memory (SCM) activity.

### How to request this report

If the currently active SMFPRMxx parameter settings indicate that SMF record type 74 subtype 10 is to be collected, then RMF Monitor III gathers the data required for the SCM Activity Report into this SMF record.

To produce this report, specify  
REPORTS(SCM)

This single-system report is only available in XML output format. Therefore, you need to specify the XPRPTS ddname in your Postprocessor job. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

### Example URL for the DDS API

http://ddshost:8803/gpm/rmfpp.xml?reports=SCM

## Contents of the report

The *Storage Class Memory (SCM) Activity Report* consists of two segments:

1. EADM device/subchannel level information. The EADM (extended asynchronous data mover) device summary segment provides the rate of start subchannel (SSCH) instructions for all EADM devices together with response time statistics consisting of pending, IOP queue and initial command response time.
2. Flash Express card level information. For each Flash Express card, the report provides measurements at both the LPAR and CPC level. The total number of requests, the rate at which requests are processed by the adapter, the rate at which data units were read and written, the average response and IOP queue time is displayed.

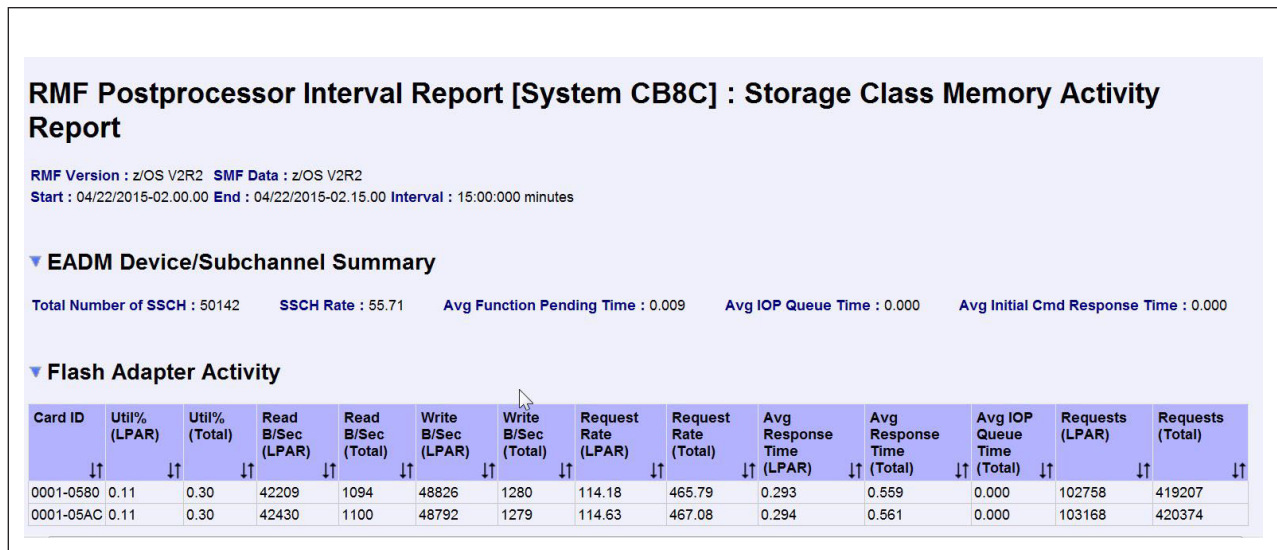


Figure 213. Storage Class Memory (SCM) Activity Report

Table 189. Fields in the SCM Activity Report

Field Heading	Meaning
EADM Device / Subchannel Summary	
This section provides summary information about the extended asynchronous data mover (EADM) devices or subchannels. EADM subchannels are similar to I/O subchannels in a way that I/O instructions can be issued. However, they do not have channel paths or device numbers assigned, and they are not defined in the I/O configuration. They are created automatically during IPL.	
Total Number of SSCH	The total number of SSCH instructions to all EADM devices in the report interval.
SSCH Rate	The number of SSCH instructions to all EADM devices per second.
Avg Function Pending Time	The average function pending time across all EADM devices in milliseconds. This is similar to function pending time for traditional I/O devices, which is the amount of time between when the SSCH is issued and the first command in the channel program is accepted.  $AVG = \frac{\text{Sum( Function Pending Time )}}{\text{Measurement Event Count}}$

Table 189. Fields in the SCM Activity Report (continued)

Field Heading	Meaning
Avg IOP Queue Time	The average IOP queue time across all EADM devices in milliseconds. This is unique to EADM devices. It represents the amount of time the request is not accepted by the adapter because it would exceed its maximum capacity. For a particular I/O request, this may occur multiple times. $\text{AVG} = \frac{\text{Sum( IOP Queue Time )}}{\text{Measurement Event Count}}$
Avg Initial Cmd Response Time	The average initial command response time across all EADM devices in milliseconds. This is the time from when the first command does not immediately proceed to execute until the successful start of execution at the SCM resource part. $\text{AVG} = \frac{\text{Sum( Initial Command Response Time )}}{\text{Measurement Event Count}}$
Flash Adapter Activity	
Card ID	The identifier of the flash adapter card.
Following fields are displayed at a system-wide level (Total) and for the current LPAR whereby IOP Queue Time is only available at the total level.	
Util(%)	The average utilization of the flash card during the interval as reported by the SCM measurement facility.
Read(B/Sec)	Bytes read per second.
Write(B/Sec)	Bytes written per second.
Request Rate	The requests processed per second.
Requests	The total number of requests.
Avg Response Time	The average response time per request in milliseconds.
Avg IOP Queue Time	The average IOP queue time per request in milliseconds.

## SDELAY - Serialization Delay report

In large systems, it may be difficult to detect and debug performance problems due to resource contention. System dumps or traditional performance reports may not be adequate tools to identify the address space that is causing a contention.

For this purpose, RMF provides global resource serialization (GRS) enqueue and latch performance statistics, as well as system suspend lock contention information to help users in analyzing serialization-related performance problems.

### How to request this report

RMF Monitor III gathers the data required for the Serialization Delay report by default in SMF record type 72 subtype 5.

To produce this report, specify  
REPORTS(SDELAY)

**Note:** The SDELAY report is only available as an interval report, not as a duration report.

If you do not want to use this report, you should suppress the associated SMF data collection for record type 72-5. Methods how to achieve this are listed in the *z/OS RMF User's Guide* in section *Defining SMF record writing*.

This single-system report is only available in XML output format. Therefore, you need to specify the XPRPTS ddname in your Postprocessor job. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

### Example URL for the DDS API

<http://ddshost:8803/gpm/rmfpp.xml?reports=SDELAY&sysid=SYSF>

## Contents of the report

The *Serialization Delay* report provides contention information on system and address space level for different types of suspend locks, GRS latches, and GRS ENQs. Reported suspend lock types (with their abbreviations used in the report in parentheses) are: CMS lock (CMS), CMS Enqueue/Dequeue lock (CMSEQDQ), CMS Latch lock (CMSLatch), CMS SMF lock (CMSSMF), LOCAL lock (Local), and CML lock (CML).

The **Serialization Delay Report** consists of two sections:

- the *Serialization Delay Summary* (see “Serialization Delay Summary” )
- the *Serialization Delay Details* (see “Serialization Delay Details” on page 438 )

### Serialization Delay Summary

The *Serialization Delay Summary* section contains system-wide summary data for all address spaces and is divided into three subsections:

- The *System Locks* subsection displays summary data for system suspend locks.
- The *GRS Latch Set Creator* subsection displays summary data about GRS latches.
- The *GRS Enqueue* subsection displays summary data about GRS enqueue requests.

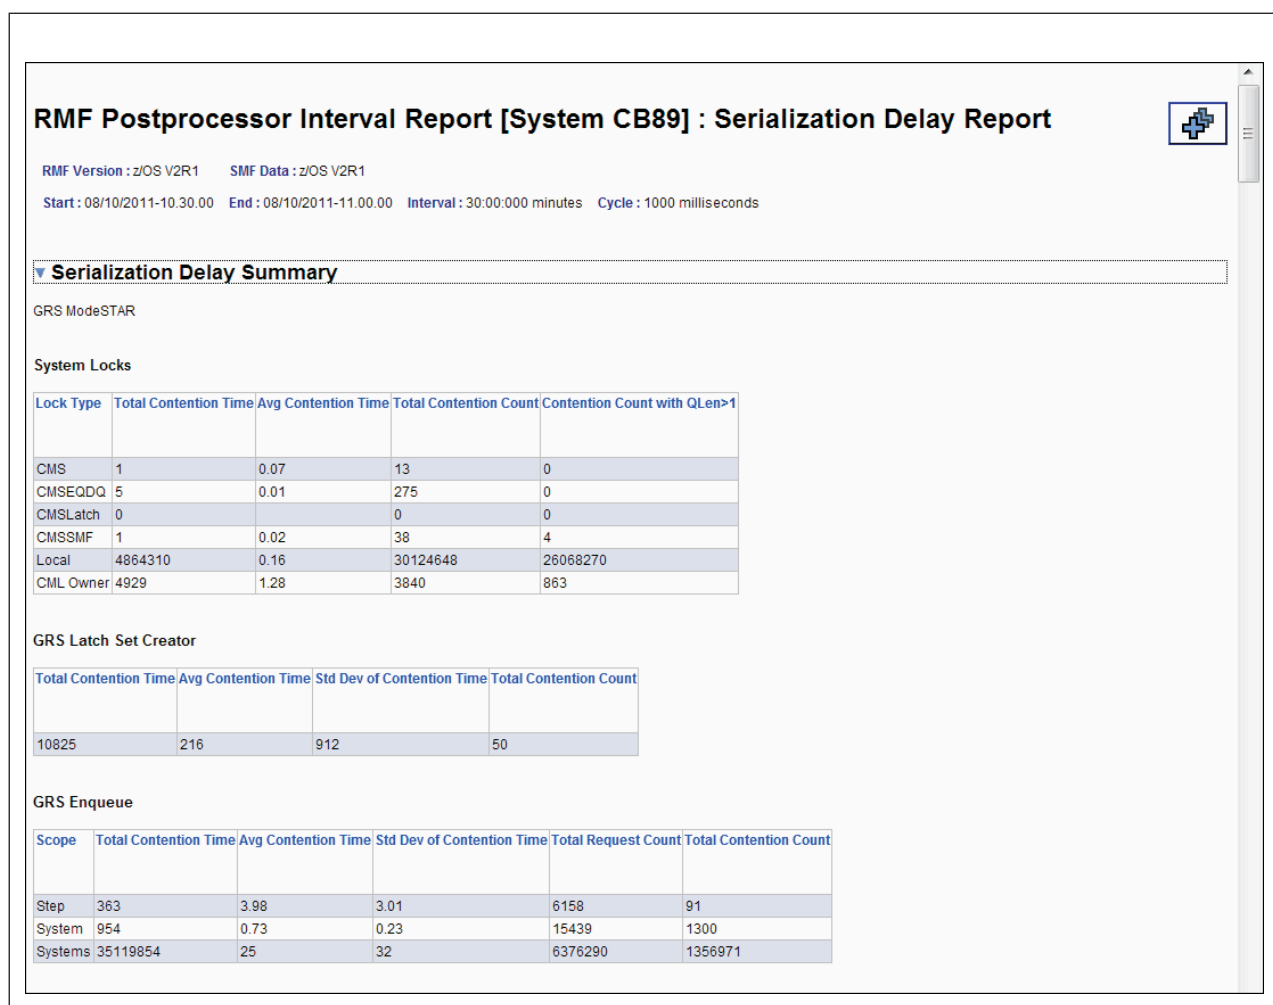


Figure 214. SDELAY Report - Serialization Delay Summary

Table 190. Fields in the Serialization Delay Summary section

Field Heading	Meaning
GRS Mode	The operation mode of GRS: <ul style="list-style-type: none"> <li>• NONE</li> <li>• RING</li> <li>• STAR</li> </ul>
<b>System Locks</b> – contains system-wide summary data on system suspend locks for all address spaces.	
Lock Type	Displays the system suspend lock type: <b>CMS</b> CMS lock <b>CMSEQDQ</b> CMS Enqueue/Dequeue lock <b>CMSLatch</b> CMS Latch lock <b>CMSSMF</b> CMS SMF lock <b>Local</b> LOCAL lock <b>CML Owner</b> CML lock owner
Total Contention Time	The total amount of time in milliseconds that a unit of work was suspended by a lock of the indicated type.
Avg Contention Time	The average amount of time in milliseconds that a unit of work was suspended by a lock of the indicated type.

## PP - SDELAY

Table 190. Fields in the *Serialization Delay Summary* section (continued)

Field Heading	Meaning
Total Contention Count	The total number of times that a unit of work was suspended by a lock of the indicated type.
Contention Count with QLen>1	The total number of times that a unit of work was suspended by a lock of the indicated type when there was already at least one other unit of work suspended for the lock (that is, queue length > 1).
<b>GRS Latch Set Creator</b> – contains summary data about GRS latches for all address spaces.	
Total Contention Time	The total amount of time in milliseconds that latch obtain requests were suspended.
Avg Contention Time	The average amount of time in milliseconds that latch obtain requests were suspended.
Std Dev of Contention Time	The standard deviation of the total contention time in milliseconds.
Total Contention Count	The total number of suspended latch obtain requests.
<b>GRS Enqueue</b> – contains summary data about GRS enqueue requests for all address spaces.	
Scope	The scope of an GRS enqueue request: <ul style="list-style-type: none"> <li>• STEP</li> <li>• SYSTEM</li> <li>• SYSTEMS</li> </ul> <p>One line is displayed for requests of a certain scope.</p>
Total Contention Time	The total amount of time in milliseconds that the GRS ENQ requests with the specified <i>Scope</i> were suspended.
Avg Contention Time	The average amount of time in milliseconds that the GRS ENQ requests with the specified <i>Scope</i> were suspended.
Std Dev of Contention Time	The standard deviation of the <i>Total Contention Time</i> in milliseconds.
Total Request Count	The total number of GRS ENQ requests with the specified <i>Scope</i> .
Total Contention Count	The total number of GRS ENQ requests with the specified <i>Scope</i> that were suspended.

### Serialization Delay Details

The *Serialization Delay Details* section provides the following information in four subsections:

- The *CMS Lock Details* subsection contains detail data about CMS/CMSEQDQ/CMSLatch/CMSSMF locks per address space (see Figure 215 on page 439).
- The *CML and Local Lock Details* subsection contains detail data about CML and LOCAL locks per address space (see Figure 216 on page 440).
- The *GRS Latch Details* subsection contains detail data about GRS latches (see Figure 217 on page 441).
- The *GRS Enqueue Details* subsection contains detail data about GRS enqueue requests (see Figure 218 on page 442).

**Note:** For each lock type, a maximum of the top twenty address spaces with the longest contention times are reported.



▼ **Serialization Delay Details**

**CMS Lock Details**

Address Space ID	Job Name	Service Class Name	Service Class Period	CMS Total Contention Time	CMS Avg Contention Time	CMS Total Contention Count	CMS Contention Count with QLen>1	CMSEQDQ Total Contention Time	CMSEQDQ Avg Contention Time	CMSEQDQ Total Contention Count	CMSEQDQ Contention Count with QLen>1	CMSLatch Total Contention Time	CMSLatch Avg Contention Time
001F	APPC	STCLOW	1	1	0.25	4	0	0	0.00	3	0		
00C5	T016023	TSOLOW	1	0	0.00	1	0						
0018	IXGLOGR	SYSTEM	1	0	0.00	1	0	0	0.00	1	0		
0009	SMSPDSE1	SYSTEM	1	0	0.00	1	0	0	0.00	35	0		
0006	XCFAS	SYSTEM	1	0	0.00	3	0						
0049	PFA	STCLOW	1	0	0.00	1	0						
0092	RMFGAT	SYSSTC	1	0	0.00	1	0	0	0.00	2	0		
000B	CONSOLE	SYSTEM	1	0	0.00	1	0						
0007	GRS	SYSTEM	1					4	0.01	205	0		
0008	SMSPDSE	SYSTEM	1					0	0.00	25	0		
0015	SMS	SYSSTC	1					0	0.00	2	0		
0050	CATALOG	SYSTEM	1					0	0.00	1	0		
002A	NETVIEW	SYSSTC	1					0	0.00	1	0		
009A	CICS2A33	STCHI	1										
00C3	T016090	TSOLOW	1										

Figure 215. SDELAY Report - Serialization Delay Details - CMS Lock Details

Table 191. Fields in the Serialization Delay Details section - CMS Lock Details

Field Heading	Meaning
<i>CMS Lock Details</i>	– contains detail data about CMS/CMSEQDQ/CMSLatch/CMSSMF locks per address space.
Address Space ID	The hexadecimal address space identifier (ASID) of the job for which lock data was collected.
Jobname	The name of the job.
Service Class Name	The name of the service class that the job has been running in.
Service Class Period	The service class period that the job has been running in.
CMS - Total Contention Time CMSEQDQ - Total Contention Time CMSLatch - Total Contention Time CMSSMF - Total Contention Time	The total amount of time in milliseconds that a unit of work of the indicated address space was suspended on the respective lock type.
CMS - Avg Contention Time CMSEQDQ - Avg Contention Time CMSLatch - Avg Contention Time CMSSMF - Avg Contention Time	The average amount of time in milliseconds that a unit of work of the indicated address space was suspended on the respective lock type.
CMS - Total Contention Count CMSEQDQ - Total Contention Count CMSLatch - Total Contention Count CMSSMF - Total Contention Count	The number of times that a unit of work of the indicated address space was suspended on the respective lock type.
CMS - Contention Count with QLen>1 CMSEQDQ - Contention Count with QLen>1 CMSLatch - Contention Count with QLen>1 CMSSMF - Contention Count with QLen>1	The number of times that a unit of work of the indicated address space was suspended on the respective lock type when there was already at least one other unit of work suspended for the lock.

PP - SDELAY

CML and Local Lock Details														
Address Space ID	Job Name	Service Class Name	Service Class Period	CML Lock Owner Total Contention Time	CML Lock Owner Avg Contention Time	CML Lock Owner Total Contention Count	CML Lock Owner Contention Count with QLen>1	Local Lock Total Contention Time	Local Lock Avg Contention Time	Local Lock Total Contention Count	Local Lock Contention Count with QLen>1	CML Lock Requestor Total Contention Time	CML Lock Requestor Avg Contention Time	CML Lock Requestor Total Contention Count
0007	GRS	SYSTEM	1	1281	0.92	1392	536	4857691	0.16	30083088	26063879			
0006	XCFAS	SYSTEM	1	2065	2.96	696	236	5069	2.75	1840	449			
0016	IOSAS	SYSTEM	1					510	0.15	3226	871	14	0.63	22
0018	IXGLOGR	SYSTEM	1					251	0.08	3082	85	359	7.97	45
003E	RMF	SYSSTC	1					208	0.02	9096	1969	14	0.60	23
0055	BPXAS	SYSSTC	1					78	0.06	1256	254			
000A	SMSVSAM	SYSTEM	1	0	0.00	3	0	72	0.02	2846	210	350	8.33	42
000C	WLM	SYSTEM	1	1	0.16	6	4	71	0.16	425	42	32	0.84	38
000B	CONSOLE	SYSTEM	1	84	0.10	825	58	53	0.05	903	100			
001F	APPC	STCLOW	1					43	0.03	1158	94			
00AB	SOAKER	STCLOW	1					29	0.00	3858	0			
0055	BPXAS	OMVSLow	2					24	0.08	288	49			
00D8	SOAKER	STCLOW	1					23	0.00	3826	0			
002A	NETVIEW	SYSSTC	1					20	0.04	466	20	54	0.12	424
0041	SOAKER	STCLOW	1					20	0.00	4021	0			

Figure 216. SDELAY Report - Serialization Delay Details - CML and Local Lock Details

Table 192. Fields in the Serialization Delay Details section - CML and Local Lock Details

Field Heading	Meaning
<i>CML and Local Lock Details</i> – contains detail data about CML and LOCAL locks per address space.	
Address Space ID	The hexadecimal address space identifier (ASID) of the job for which lock data was collected.
Jobname	The name of the job.
Service Class Name	The name of the service class that the job has been running in.
Service Class Period	The service class period that the job has been running in.
CML Lock Owner - Total Contention Time	The total amount of time in milliseconds that a unit of work from another address space was suspended when requesting the LOCAL lock of the indicated address space.
CML Lock Owner - Avg Contention Time	The average amount of time in milliseconds that a unit of work from another address space was suspended when requesting the LOCAL lock of the indicated address space.
CML Lock Owner - Total Contention Count	The number of times that a unit of work from another address space was suspended when requesting the LOCAL lock of the indicated address space.
CML Lock Owner - Contention Count with QLen>1	The number of times that a unit of work from another address space was suspended when requesting the LOCAL lock of the indicated address space and there was already at least one other unit of work waiting for this lock.
Local Lock - Total Contention Time	The total amount of time in milliseconds that a unit of work of the indicated address space was suspended on a LOCAL lock.
Local Lock - Avg Contention Time	The average amount of time in milliseconds that a unit of work of the indicated address space was suspended on a LOCAL lock.
Local Lock - Total Contention Count	The number of times that a unit of work of the indicated address space was suspended on a LOCAL lock.
Local Lock - Contention Count with QLen>1	The number of times that a unit of work of the indicated address space was suspended on a LOCAL lock when there was already at least one other unit of work suspended.
CML Lock Requestor - Total Contention Time	The total amount of time in milliseconds that a unit of work of the indicated address space was suspended when requesting the LOCAL lock of another address space.

Table 192. Fields in the Serialization Delay Details section - CML and Local Lock Details (continued)

Field Heading	Meaning
CML Lock Requestor - Contention Time	The average amount of time in milliseconds that a unit of work of the indicated address space was suspended when requesting the LOCAL lock of another address space.
CML Lock Requestor - Total Contention Count	The number of times that a unit of work from this address space was suspended when requesting the LOCAL lock of another address space.
CML Lock Requestor - Contention Count with QLen>1	The number of times that a unit of work from this address space was suspended when requesting the LOCAL lock of another address space and there was already at least one other unit of work waiting for that lock.

Address Space ID	Job Name	Service Class Name	Service Class Period	Latch Set Creator Total Contention Time	Latch Set Creator Avg Contention Time	Latch Set Creator Std Dev Contention Time	Latch Set Creator Total Contention Count	Latch Requestor Total Contention Time	Latch Requestor Avg Contention Time	Latch Requestor Std Dev Contention Time	Latch Requestor Total Contention Count
0010	OMVS	SYSTEM	1	4927	821	2011	6	4926	4926		1
0049	PFA	STCLOW	1	4332	4332		1	4332	4332		1
000C	WLM	SYSTEM	1	1258	33	8.59	38	1258	33	8.59	38
0016	IOSAS	SYSTEM	1	222	222		1	222	222		1
003F	RRS	STCHI	1	72	36	47	2	72	36	47	2
0018	IXGLOGR	SYSTEM	1	11	5.50	7.45	2	11	5.50	7.45	2
0055	BPXAS	SYSSTC	1					1	0.25	0.00	4
0055	BPXAS	OMVSLOW	2					0	0.00		1

Figure 217. SDELAY Report - Serialization Delay Details - GRS Latch Details

Table 193. Fields in the Serialization Delay Details section - GRS Latch Details

Field Heading	Meaning
<i>GRS Latch Details</i> – contains detail data about GRS latches.	
Address Space ID	The hexadecimal address space identifier (ASID) of the job for which lock data was collected.
Jobname	The name of the job.
Service Class Name	The name of the service class that the job has been running in.
Service Class Period	The service class period that the job has been running in.
In the following field descriptions, the term <i>Latch Set Creator</i> denotes statistics for latch obtain requests <i>against latch sets created by this address space</i> and <i>Latch Requestor</i> denotes statistics for latch obtain requests <i>issued from this address space</i> :	
Latch Set Creator - Total Contention Time Latch Requestor - Total Contention Time	The amount of contention time in milliseconds that was caused by latch obtain requests.
Latch Set Creator - Avg Contention Time Latch Requestor - Avg Contention Time	The average amount of contention time in milliseconds that was caused by latch obtain requests.
Latch Set Creator - Std Dev of Contention Time Latch Requestor - Std Dev of Contention Time	The standard deviation of the total contention time.
Latch Set Creator - Total Contention Count Latch Requestor - Total Contention Count	The number of times a latch obtain request was suspended.

## PP - SDELAY

GRS Enqueue Details															
Address Space ID	Job Name	Service Class Name	Service Class Period	ENQ STEP Total Contention Time	ENQ STEP Avg Contention Time	ENQ STEP Std Dev Contention Time	ENQ STEP Request Count	ENQ STEP Contention Count	ENQ SYSTEM Total Contention Time	ENQ SYSTEM Avg Contention Time	ENQ SYSTEM Std Dev Contention Time	ENQ SYSTEM Request Count	ENQ SYSTEM Contention Count	ENQ SYSTEMS Total Contention Time	ENQ SYSTEMS Avg Contention Time
0055	BPXAS	SYSSTC	1	270	5.00	2.80	67	54	0			12	0		
0055	BPXAS	OMVSLow	2	75	4.16	2.31	21	18							
002A	NETVIEW	SYSSTC	1	18	0.94	2.01	1717	19	0			139	0		
0029	HSAMPROC	SYSSTC	1	0			28	0							
0018	IXGLOGR	SYSTEM	1	0			390	0	0			642	0	128	128
0010	OMVS	SYSTEM	1	0			18	0	0			58	0	0	
000C	WLM	SYSTEM	1	0			579	0						0	
000B	CONSOLE	SYSTEM	1	0			240	0	0	0.00	0.00	256	2		
0006	XCFAS	SYSTEM	1	0			16	0	0			8	0		

Figure 218. SDELAY Report - Serialization Delay Details - GRS Enqueue Details

Table 194. Fields in the Serialization Delay Details section - GRS Enqueue Details

Field Heading	Meaning
<i>GRS Enqueue Details</i> – contains detail data about GRS enqueue requests.	
Address Space ID	The hexadecimal address space identifier (ASID) of the job for which lock data was collected.
Jobname	The name of the job.
Service Class Name	The name of the service class that the job has been running in.
Service Class Period	The service class period that the job has been running in.
ENQ STEP - Total Contention Time ENQ SYSTEM - Total Contention Time ENQ SYSTEMS - Total Contention Time	The total amount of contention time in milliseconds that was caused by GRS ENQ requests of the indicated scope for this address space.
ENQ STEP - Avg Contention Time ENQ SYSTEM - Avg Contention Time ENQ SYSTEMS - Avg Contention Time	The average amount of contention time in milliseconds that was caused by GRS ENQ requests of the indicated scope for this address space.
ENQ STEP - Std Dev of Contention Time ENQ SYSTEM - Std Dev of Contention Time ENQ SYSTEMS - Std Dev of Contention Time	The standard deviation of the total contention time in milliseconds for GRS ENQ requests of the indicated scope for this address space.
ENQ STEP - Request Count ENQ SYSTEM - Request Count ENQ SYSTEMS - Request Count	The total number of GRS ENQ requests of the indicated scope for this address space.
ENQ STEP - Contention Count ENQ SYSTEM - Contention Count ENQ SYSTEMS - Contention Count	The total number of GRS ENQ requests of the indicated scope that were suspended for this address space.

## SDEVICE - Shared Device Activity report

This section describes the Shared Device report. There are two types:

- The Shared Direct Access Device Activity Report
- The Shared Magnetic Tape Device Report

The report gives you an overall performance picture of DASD and TAPE devices that are shared between z/OS systems in a sysplex.

For each shared DASD or tape device the report contains one line for each system that has access to it. The additional system line shows the device activity contributed by all systems in the sysplex.

## How to request this report

Monitor I gathers data for the DASD Activity report automatically with the default option `DEVICE(DASD)`. If you want to suppress gathering, you need to specify `DEVICE(NODASD)`.

To gather data for the TAPE Activity report, specify `DEVICE(TAPE)`.

To produce this report, specify  
`SYSRPTS(SDEVICE(options))`

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

**Note:** The report requires matching device numbers (the physical device must have the same device number on all systems), or self-defining devices to give meaningful results.

### Example URL for the DDS API

`http://ddshost:8803/gpm/rmfpp.xml?reports=SDEVICE(NMBR(2000:2200))`

## Using the information given in the report

The summary line allows you to identify a bottleneck caused by device delay in the sysplex. Furthermore, it allows you to see each systems share in the bottleneck.

The summary device activity rate and the device utilization show the total load on the device. The single-system values show the share of each system.

## Shared Direct Access Device Activity report

The following example reports about a sysplex consisting of two systems (SYSD and SYSE). Only two devices are shown.

Both devices have the same device number on both systems.

**PP - SDEVICE**

SHARED DIRECT ACCESS DEVICE ACTIVITY																	PAGE 1		
z/OS V2R2		SYSPLX SYSDPLEX		DATE 09/28/2016		INTERVAL 15.00.000													
		RPT VERSION V2R2 RMF		TIME 12.30.00		CYCLE 1.000 SECONDS													
TOTAL SAMPLES(AVG) = 900.0 (MAX) = 900.0 (MIN) = 900.0																			
DEV NUM	DEVICE TYPE	VOLUME SERIAL	PAV	SMF SYS ID	IODF SUFF	LCU	DEVICE ACTIVITY RATE	AVG RESP TIME	AVG IOSQ TIME	AVG CMR DLY	AVG DB DLY	AVG INT DLY	AVG PEND TIME	AVG DISC TIME	AVG CONN TIME	% DEV CONN	% DEV UTIL	% DEV RESV	AVG NUMBER ALLOC
2180	33909	SYST10		*ALL			0.036	0.248	0.000	0.000	0.000	0.005	0.120	0.000	0.128	0.00	0.00	0.0	0.0
			1	SYSD	00	0062	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.00	0.00	0.0	0.0
			1	SYSE	00	0062	0.036	0.248	0.000	0.000	0.000	0.005	0.120	0.000	0.128	0.00	0.00	0.0	0.0
2181	33909	SYST11		*ALL			0.071	0.264	0.000	0.006	0.000	0.000	0.130	0.000	0.134	0.00	0.00	0.0	0.0
			1	SYSD	00	0062	0.036	0.280	0.000	0.012	0.000	0.000	0.136	0.000	0.144	0.00	0.00	0.0	0.0
			1	SYSE	00	0062	0.036	0.248	0.000	0.000	0.000	0.000	0.124	0.000	0.124	0.00	0.00	0.0	0.0

Figure 219. Shared DASD Activity Report

For the field descriptions, see Table 195.

**Shared Magnetic Tape Device Activity report**

The following example reports about a sysplex consisting of four systems (P40, P50, P60, and P70). Only some devices are shown.

SHARED MAGNETIC TAPE DEVICE ACTIVITY																	PAGE 1		
z/OS V2R2		SYSPLX SYSDPLEX		DATE 09/28/2016		INTERVAL 30.00.001													
		RPT VERSION V2R2 RMF		TIME 00.15.00		CYCLE 1.000 SECONDS													
TOTAL SAMPLES(AVG) = 1800.0 (MAX) = 1800.0 (MIN) = 1800.0																			
DEV NUM	DEVICE TYPE	VOLUME SERIAL	SMF SYS ID	IODF SUFF	LCU	DEVICE ACTIVITY RATE	AVG RESP TIME	AVG IOSQ TIME	AVG CMR DLY	AVG DB DLY	AVG PEND TIME	AVG DISC TIME	AVG CONN TIME	% DEV CONN	% DEV UTIL	% DEV RESV	NUMBER OF MOUNTS	AVG MOUNT TIME	TIME DEVICE ALLOC
05D0	3490	TAP508	*ALL			9.511	42.97	1.134	11.71	2.442	27.51	7.312	7.014	6.64	13.6	0.0	4	19	57:37
		TAP508	P40	A1	0006	0.012	451.8	227.4	39.02	131.8	217.7	0.156	6.621	0.01	0.01	0.0	0	0	13:37
		TAP508	P50	A1	0006	9.474	40.92	0.157	11.52	1.034	26.35	7.379	7.023	6.61	13.6	0.0	4	19	12:01
		TAP508	P60	A1	0006	0.012	152.9	45.24	33.03	29.94	101.5	0.342	5.845	0.01	0.01	0.0	0	0	15:39
		TAP508	P70	A1	0006	0.012	1130	409.3	78.81	560.8	714.7	0.145	6.014	0.01	0.01	0.0	0	0	16:18
05D5	3490	TAP606	*ALL			9.829	54.13	3.174	14.18	4.034	34.53	9.913	6.514	6.39	16.2	0.0	8	28	1:04:04
		TAP606	P40	A1	0006	0.012	1347	636.4	31.82	640.8	702.6	0.158	8.036	0.01	0.01	0.0	0	0	16:27
		TAP606	P50	A1	0006	1.109	37.01	3.235	8.147	7.942	25.94	1.969	5.868	0.65	0.86	0.0	1	1:05	17:17
		TAP606	P60	A1	0006	8.696	50.71	0.000	14.72	1.248	33.01	11.02	6.689	5.73	15.3	0.0	7	22	16:11
		TAP606	P70	A1	0006	0.012	2145	955.8	60.86	1061	1182	0.147	7.625	0.01	0.01	0.0	0	0	14:07

Figure 220. Shared Magnetic Tape Device Activity Report

**Field descriptions**

The table Table 195 contains all report fields and the description how the values shown are calculated.

Table 195. Fields in the Shared Device Activity Reports

Field Heading	Meaning
SYSPLX	Sysplex name
DATE	This is the earliest date found in all records used to process this report

Table 195. Fields in the Shared Device Activity Reports (continued)

Field Heading	Meaning
INTERVAL	<p>This is the longest interval which can be built by using integer multiple interval lengths of all SMF record interval lengths.</p> <p><b>Note:</b> This value is referenced below as the "Common Interval Length" (abbreviated as CIL).</p> <p>The interval length of each system is referenced as the "System Interval Length" (abbreviated as SIL).</p> <p>If the SMF or RMF interval options and/or SYNCH option do not match it might not be possible to find such a matching interval. In this case an interval of 1 hour (the maximum possible interval length ) will be taken.</p>
CYCLE	<p>This is the cycle value found in the first SMF record used to process this report.</p> <p><b>Note:</b> If the CYCLE values of all SMF records used to build the report do not match, no report is generated.</p>
TOTAL SAMPLES	<p>This field shows the total number of samples used to build this report.</p> <p><b>Note:</b> This "Single System Total Samples" (abbreviated as STS) might differ for each system. Therefore, the MIN/MAX/AVG value of all system total sample values is shown. If these value differ strongly, the system summary line values might not be very meaningful.</p>
DEV NUM	<p>The four-digit hexadecimal device number of a physical I/O device.</p> <p><b>Note:</b> The device number is printed only on the summary line if it is equal for all systems. Otherwise, the device number is printed on each system line. The summary line contains the same device number as the first system line in this case.</p>
DEVICE TYPE	<p>The device type of the volume.</p>
VOLUME SERIAL	<p>The volume serial number of the volume mounted on the device at the end of the reporting interval.</p> <p><b>Note:</b> In the Shared Direct Access Device report, this field is printed only on the summary line. The system line field is blank.</p> <p>In the Shared Magnetic Tape Device report, this field is printed on the summary line and on that system line which had this volume mounted at end of interval. If no volume was mounted in any system on that device at end of interval, this field is blank on the summary line and on the system lines.</p>
PAV	<p>The number of parallel access volumes (base and alias) which were active at the end of the reporting interval.</p> <p>If the number has changed during the reporting interval, it is followed by an '*'.</p> <p>If the device is a HyperPAV base device, the number is followed by an 'H', for example, 5.4H . The value is the average number of HyperPAV volumes (base and alias) in that interval.</p> $\text{Average \# of HPAV devices} = \frac{\text{Accumulated \# of HPAV devices}}{\text{Number of Samples}}$
SMF SYS ID	<p>The SMF System IDs of the systems found sharing this device. The summary line contains the text string *ALL. The system ID is followed by an asterisk, if for any reason either no data, or only partial data could be presented.</p> <p><b>Note:</b> Data is also considered to be partial, if the interval length of that system differs from the interval shown in the report header.</p>
IODF SUFF	<p>The IODF suffix in effect for this system</p> <p><b>Note:</b> In the single-system Device Activity report, the header contains the complete IODF name and suffix. Here, only the suffix is presented.</p>
LCU	<p>The number of the logical control unit to which the device belongs.</p> <p>On the system summary line, this field is blank.</p> <p><b>Note:</b> The LCU number for the same device can differ between the different systems even if the device is being shared.</p>

## PP - SDEVICE

Table 195. Fields in the Shared Device Activity Reports (continued)

Field Heading	Meaning
DEVICE ACTIVITY RATE	<p>The rate at which start subchannel (SSCH) instructions to the device completed successfully.</p> <p>The calculation for the single system line is:</p> $\text{DEVICE ACTIVITY RATE} = \frac{\text{SSCH Count}}{\text{CIL}}$ <p>In the summary line, this field contains the sum of the rates for each single system.  <b>Note:</b> For multi-exposure devices the field reflects the value of the entire device. This is true also in all following fields.</p>
AVG RESP TIME	<p>The average number of milliseconds the device required to complete an I/O request.</p> <p>The average response time consist of two parts, the average service time and the average IOS queue time.</p> <p>For the single system line this is:</p> $\text{AVG RESP TIME} = \text{Avg IOSQT} + \text{AVG PEND TIME} + \text{AVG DISC TIME} + \text{AVG CONN TIME}$ <p>In the summary line, the same formula is used using the corresponding summary line fields.</p>
AVG IOSQ TIME	<p>The average number of milliseconds an I/O request must wait on an IOS queue before a SSCH instruction can be issued.</p> <p>Using the abbreviation for "Device Activity Rate (<b>DAR</b>)", the calculation for the single system line is:</p> $\text{AVG IOSQ TIME} = \frac{\text{IOS Queue Count}}{\text{STS} * \text{DAR}}$ <p>In the summary line, this field contains the weighted average IOS queue times of the single systems. The weighting factor used is the Measurement Event Count (<b>MEC</b>). The weighting factor for System SYSi is therefore, assuming we have n systems:</p> $W_i = \frac{\text{MEC}(\text{SYS}_i)}{\text{MEC}(\text{SYS}_1) + \text{MEC}(\text{SYS}_2) + \dots + \text{MEC}(\text{SYS}_n)}$ <p>The weighted AVG IOS queue time is therefore calculated as:</p> $\text{WEIGHTED AVG IOS QT} = (\text{AVG IOS QT}(\text{SYS}_1)) * W_1 + (\text{AVG IOS QT}(\text{SYS}_2)) * W_2 + \dots + (\text{AVG IOS QT}(\text{SYS}_n)) * W_n$ <p><b>Note:</b> This weighting algorithm is different to the LCU summary line algorithm used in the single-system Device Activity report.</p> <p>The measurement event count (MEC) is the same as the number of start subchannel instructions (SSCH), unless there has been a timer overflow error in the channel.</p>
AVG CMR DLY	<p>The average number of milliseconds that a successfully initiated start or resume function needs until the first command is indicated as accepted by the device.</p> $\text{AVG CMR DLY} = \frac{\text{Initial Command Response Time}}{\text{MEC}}$



Table 195. Fields in the Shared Device Activity Reports (continued)

Field Heading	Meaning
AVG DB DLY	<p>The average number of milliseconds of delay that I/O requests to this device encountered because the device was busy.</p> <p>The calculation for the single system line is:</p> $\text{AVG DB DLY} = \frac{\text{Total DB DLY Time}}{\text{MEC}}$ <p>In the summary line, this field contains the weighted average of the individual system AVG DB DELAY times. The weighting algorithm used is the same as described in AVG IOSQ TIME.</p>
AVG INT DLY	<p>The average interrupt delay time in units of milliseconds encountered for I/O requests to this device. For each I/O request, the time is measured from when the I/O operation is complete to when the operating system begins to process the status.</p> $\text{AVG INT DLY} = \frac{\text{Device Interrupt Delay Time}}{\text{MEC}}$
AVG PEND TIME	<p>The average number of milliseconds an I/O request must wait in the hardware.</p> <p>The calculation for the single system line is:</p> $\text{AVG PEND TIME} = \frac{\text{Total Pending Time}}{\text{MEC}}$ <p>In the summary line, this field contains the weighted average of the individual systems AVG PEND times. The weighting algorithm used is the same as described in AVG IOSQ TIME.</p>
AVG DISC TIME	<p>The average number of milliseconds the device was disconnected while processing an SSCH instruction.</p> <p>The calculation for the single system line is:</p> $\text{AVG DISC TIME} = \frac{\text{Total Disconnect Time}}{\text{MEC}}$ <p>In the summary line, this field contains the weighted average of the individual systems AVG DISC TIME times. The weighting algorithm used is the same as described in AVG IOSQ TIME.</p>
AVG CONN TIME	<p>The average number of milliseconds the device was connected to a channel path and actually transferring data between the device and central storage.</p> <p>The calculation for the single system line is:</p> $\text{AVG CONN TIME} = \frac{\text{Total Connect Time}}{\text{MEC}}$ <p>In the summary line, this field contains the weighted average of the individual systems AVG CONN TIME times. The weighting algorithm used is the same as described in AVG IOSQ TIME.</p>
% DEV CONN	<p>The percentage of time during the interval when the device was connected to a channel path.</p> <p>The calculation for the single system line is:</p> $\% \text{ DEV CONN} = 100 * \frac{\text{Device Connect Time}}{\text{CIL}}$ <p>In the summary line, this field contains the sum of the single system % DEV CONN values.</p>

## PP - SDEVICE

Table 195. Fields in the Shared Device Activity Reports (continued)

Field Heading	Meaning
% DEV UTIL	<p>The percentage of time during the interval when the device was in use. This percentage includes both the time when the device was involved in I/O operations (connect and disconnect time) and the time when it was reserved but not involved in an I/O operation.</p> <p>Using the abbreviations for "reserved but not involved in an I/O operation (UTL)", "single system total samples (STS)", the single system value is calculated as:</p> $\% \text{ DEV UTIL} = 100 * \frac{(\text{Device Connect Time} + \text{Disconnect Time})}{\text{CIL}} + 100 * \frac{\text{UTL} * \text{SIL}}{\text{STS} * \text{CIL}}$ <p>In the summary line, this field contains the sum of the single system % DEV UTIL values.</p>
% DEV RESV	<p>The percentage of time during the interval when a shared device was reserved by the system on which RMF was started.</p> <p>Using the abbreviations for "number reserved samples (DRP)", and "system total samples (STS)", the calculation used for the single system line is:</p> $\% \text{ DEV RESV} = 100 * \frac{\text{DRP} * \text{SIL}}{\text{STS} * \text{CIL}}$ <p>This is the percent of the common interval length time which this system holds an reserve on this device.</p> <p>In the summary line, this field contains the sum of the single system % DEV RESV values.</p>
AVG NUMBER ALLOC	<p>The average number of data control blocks (DCBs) and access method control blocks (ACBs) concurrently allocated for each volume.</p> <p>This field is reported only for the shared direct access storage devices.</p> <p>The calculation used for the single system line is:</p> $\text{AVG NUMBER ALLOC} = \frac{\text{Num Alloc} * \text{System Interval Length}}{\text{STS} * \text{Common Interval Length}}$ <p>In the summary line, this field contains the sum of the single system AVG NUMBER ALLOC values.</p>
NUMBER OF MOUNTS	<p>The number of tape mounts, shown as an integer value, detected by RMF.</p> <p>This field is reported only for magnetic tape devices.</p> <p>If the tape mount was pending at the <b>first</b> cycle of the interval, an asterisk is placed <b>before</b> the numerical value of the tape mount. If the tape mount was pending at the <b>last</b> cycle of the interval, an asterisk is placed immediately <b>following</b> the numerical value of the tape mount.</p> <p>If a mount-pending condition is detected at the first cycle of the interval, the mount count for the interval increments by one.</p> <p>At the single system line the value is displayed as a integer value allowing a range between 0 and 9999.</p> <p>In the summary line, the mount count for the tape device is shown as the sum of the system line values. No indications, however, are displayed to show the mount pending status at the begin or end of the reporting interval.</p>

Table 195. Fields in the Shared Device Activity Reports (continued)

Field Heading	Meaning
AVG MOUNT TIME	<p>The average mount time pending for every device, expressed in the form of HH:MM:SS.</p> <p>Using the abbreviations “Mount Pending Samples (MTP)”, and “Mount Total Count (MTC)”, the calculation used for the single system line is:</p> $\text{AVG MOUNT TIME} = \frac{\text{MTP} * \text{SIL}}{\text{STS} * \text{MTC}}$ <p><b>Note:</b> The mount time value has units of full seconds, while the interval length usually is shown in units of milliseconds.</p> <p>If the mount count or the sample count is zero, the result is zero.</p> <p>In the summary line, the average mount time is calculated as:</p> $\text{AVG MOUNT TIME} = \frac{\text{AVG Mount Time}(\text{SYS1}) * \text{MTC}(\text{SYS1}) + \dots}{\sum \text{All System MTC Counts}}$ <p>This field is reported only for magnetic tape devices.</p>
TIME DEVICE ALLOC	<p>The total time the device was allocated during the interval, expressed in the form of HH:MM:SS.</p> <p>This field is reported only for magnetic tape devices.</p> $\text{TIME DEVICE ALLOC} = \frac{\text{ALC} * \text{SIL}}{\text{STS}}$ <p>If the system total sample count is zero, the result is zero.</p> <p>In the summary line, the field contains the sum of the system line values.</p> <p><b>Note:</b> The mount time value has units of full seconds, while the interval length in the report header is shown in units of milliseconds.</p>

## Spreadsheet and Overview reference

You can make this report available through Overview records in a spreadsheet, using the Spreadsheet Reporter.

The following table shows all criteria and the corresponding Overview criterion names for creating Overview records. For details, see the *z/OS RMF User's Guide*.

Table 196. Overview names in the Shared DASD Activity Report

Field Heading or Meaning	Overview Name
Percent reserved	DR
Percent mount pending	DMTPEND
Percent device utilization	DVUTL
Device activity rate	DART
Average connect time	DCTAVG
Average disconnect time	DDTAVG
Average pending time	DPTAVG
Average IOS queue time	DQTAVG
Average response time	DRTAVG
Average device busy delay time	DBDL
Average initial command response delay time	CMRDL
Average interrupt delay time	INTDL

---

## TRACE - Trace Activity report

The Trace Activity report provides information from various trace variables.

The report shows snapshots of each of the specified variables along with timing information. RMF trace treats values collected as unsigned binary integers. See *z/OS RMF User's Guide* for a description of how to specify the TRACE options. A description of the trace variables is in Table 198 on page 451.

**Note:** Monitor I gathers and reports all trace variables the way they are provided by the system. Monitor I cannot influence the format, and does not perform any calculation.

### How to request this report

To gather data for this report, specify as a Monitor I gatherer option  
TRACE(variable [,options list])

To produce this report, specify  
REPORTS(TRACE)

**Note:** The TRACE report is only available as an interval report, not as a duration report.

### Using the information given in the report

You can use the Trace report to monitor the SRM multiprogramming level (MPL) adjustment or monitor the contention detected and handled by the system. You can see how the system handles contention by tracing the following variables: RCVUICA, RCVCPUA and RCVPTR.

### Contents of the report

RMF reports all trace variables that contain invalid data on a separate report page.

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

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

TRACE ACTIVITY												
z/OS V2R2		SYSTEM ID SYS1			DATE 09/28/2016		INTERVAL 14.59.946				PAGE 1	
		RPT VERSION V2R2 RMF			TIME 16.30.00		CYCLE 1.000 SECONDS					
SECONDS/SET= 60.00		CYCLES/SAMPLE= 1		NUMBER OF SAMPLES= 900		SAMPLES/SET= 60		NUMBER OF SETS= 15				
TIME	* RCEESINU	* MCVMGAGE	* CCVUTILP	* RCVDPAGRT	* RCVDPTR	* SMCABFLS	* MCVFRCNT	* MCVSTCRI	* RCVMFXA	* RCEDFRS		
MM.SS.TT	* MAXIMUM	* AVERAGE	* AVERAGE	* END	* MAXIMUM	* MAXIMUM	* END	* AVERAGE	* MAXIMUM	* END		
31.00.00	* 727108	* 16003.00	* 89.60	* 23	* 52	* 0	* 300	* 225.68	* 14	* 132773		
32.00.00	* 726553	* 16003.00	* 97.90	* 14	* 15	* 0	* 300	* 220.00	* 15	* 132781		
33.00.00	* 721806	* 16003.00	* 93.85	* 364	* 8	* 0	* 300	* 233.90	* 16	* 132788		
34.00.00	* 714439	* 16003.00	* 98.53	* 191	* 24	* 0	* 300	* 233.60	* 13	* 132807		
35.00.00	* 713054	* 16003.00	* 83.60	* 0	* 2	* 0	* 300	* 223.58	* 16	* 132814		
36.00.00	* 704624	* 16003.00	* 83.11	* 0	* 5	* 0	* 300	* 221.70	* 15	* 132822		
37.00.00	* 708467	* 16003.00	* 92.88	* 0	* 25	* 0	* 300	* 229.28	* 14	* 132836		
38.00.00	* 706935	* 16003.00	* 90.08	* 0	* 0	* 0	* 300	* 226.15	* 15	* 132837		
39.00.00	* 709305	* 16003.00	* 98.31	* 0	* 3	* 0	* 300	* 214.08	* 14	* 132854		
40.00.00	* 716450	* 16003.00	* 96.86	* 357	* 62	* 0	* 300	* 217.25	* 14	* 132944		
41.00.00	* 725154	* 16003.00	* 100.46	* 52	* 0	* 0	* 300	* 234.45	* 13	* 132948		
42.00.00	* 704379	* 16003.00	* 102.36	* 83	* 29	* 0	* 300	* 249.68	* 14	* 132956		
43.00.00	* 697476	* 16003.00	* 98.73	* 0	* 18	* 0	* 300	* 225.03	* 13	* 132968		
44.00.00	* 692562	* 16003.00	* 99.91	* 89	* 69	* 0	* 300	* 238.70	* 13	* 132973		
45.00.00	* 699938	* 16003.00	* 102.36	* 8	* 71	* 0	* 300	* 225.41	* 13	* 132975		
MAXIMUM*	727108	* 16003.00	* 102.36	* 364	* 71	* 0	* 300	* 249.68	* 16	* 132975		
MINIMUM*	692562	* 16003.00	* 83.11	* 0	* 0	* 0	* 300	* 214.08	* 13	* 132773		
AVERAGE*	711216.66	* 16003.00	* 95.24	* 78.73	* 25.53	* 0.00	* 300.00	* 227.90	* 14.13	* 132871.73		

Figure 221. TRACE Report

Table 197. Fields in the Trace Activity Report

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

Table 198. Variables in the Trace Activity Report

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 logical processors (threads)
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

## PP - TRACE

Table 198. Variables in the Trace Activity Report (continued)

Variable	Value
MCVFCNT	Number of pages needed to be stolen by force steal routine
MCVMGAGE	Expanded storage migration age
MCVSBTF	Long term percentage of eligible storage that is actually fixed
MCVSIPR	Common page-in rate
MCVSTCRI	Highest system UIC
MCVTWSS	Common target working set size
OMDGAMRE	Maximum number of messages on the action message retention facility (AMRF) queue. If a large number of action messages are retained on the AMRF queue for a particular period, it may mean more operators are needed for that period.
OMDGCMDI	Number of commands issued per second.
OMDGOREB	Maximum number of operator reply entries (OREs) on the system reply queue. To eliminate thrashing, use this number to monitor and adjust the ORE buffer limit set at IPL time. To dynamically adjust this limit, use the CONTROL M command.
OMDGWQEB	Maximum number of WTO queue elements (WQEs) on the system output queue. To eliminate thrashing (excessive data movement which confines system to doing little useful work), use this number to monitor and adjust the WTO buffer time limit set at IPL time. To dynamically adjust this limit, use the CONTROL M command.
OMDGWTLI	Number of write-to-logs (WTLs) issued per second indicating the number of records going to SYSLOG within a time period. To control the number of data sets produced during the day, vary the number of records per SYSLOG data set.
OMDGWTOI	Total number of lines of messages, write-to-operators (WTOs) issued per second. Use it to determine the peak message rate period and the average message rate.
RAXESCT	Number of common storage pages on expanded storage
RAXFMCT	Number of frames allocated to common storage
RCEAEC	Total number of expanded storage E frames currently on the ESTE queue
RCEAECLO	Available expanded storage low threshold
RCEAECOK	Available expanded storage satisfactory threshold
RCEAFC	Total number of frames currently on all available frame queues
RCEAFCLO	Available central storage low threshold
RCEAFCOK	Available central storage satisfactory threshold
RCEBELFX	Total number of fixed pages below the 16M line in central storage, which is the sum of page-fixed LSQA, SQA (excluding reserved SQA), and V=R allocated pages
RCECOMPI	Number of common area pages paged-in
RCECOMPO	Number of common area pages paged-out
RCEDFRS	Number of times a deferred frame allocation has been satisfied
RCEESINU	Number of in-use expanded storage frames
RCEESREA	Number of non-VIO pages read from expanded storage
RCEESWRT	Number of pages written to expanded storage frames
RCEHSPEM	Total number of hiperspace pages migrated from expanded storage to auxiliary storage
RCEHSPER	Total number of hiperspace pages in the system read from expanded storage to central storage
RCEHSPEW	Total number of hiperspace pages written from central storage to expanded storage
RCEHSPPI	Total number of hiperspace pages paged in from auxiliary storage
RCEHSPPO	Total number of hiperspace pages paged out to auxiliary storage
RCELPAPI	Number of PLPA and PLPA directory pages paged-in
RCEMBEL	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

Table 198. Variables in the Trace Activity Report (continued)

Variable	Value
RCEPOOL	Number of frames currently available to the system. Frames that are backing permanent storage (nucleus frames, hardware storage area frames, FLPA frames or fixed BLDL frames), bad frames and offline frames are excluded.
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
RMCASOISC	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
RMCAXSXC	Auxiliary storage shortage swap count
RMCTTRPC	Number of pages used for transaction elements
SMCABFLS	Number of records lost because of a shortage of buffers

Table 198. Variables in the Trace Activity Report (continued)

Variable	Value
SMCABFWT	Number of buffers written
SMCACNBF	Current number of buffers
SMCADSCT	Number of records lost because of a full data set
SMCANMFL	Current number of full buffers
SMCARCWT	Number of records written

### Standard deviation output

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

### Example

1.123E +01 is the same as  $1.123 \times 10.1$  or 11.23

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

---

## VSTOR - Virtual Storage Activity report

The Virtual Storage Activity report provides information about the use of virtual storage that can help your installation manage its use of virtual storage.

The report is formatted into the following four sections, each with a separate subheading:

- COMMON STORAGE SUMMARY (see “Common Storage Summary section” on page 456 )
- COMMON STORAGE DETAIL (see “Common Storage Detail section” on page 458 )
- PRIVATE AREA SUMMARY (see “Private Area Summary section” on page 459 )
- PRIVATE AREA DETAIL (see “Private Area Detail section” on page 462 )

Most of the information you need to begin managing virtual storage appears in the common storage summary report. When this report indicates a problem, you can request one of the three more comprehensive reports for additional information. The structure and fields of the different reports are described in “Contents of the report” on page 455.

### Free and allocated storage

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

### Using the information given in the report

Information on virtual storage use is particularly helpful in the process of long-term measurements. It helps you, for example, understand your current use of virtual storage, see the relationship between increased use of your system and



increased demands on virtual storage, and predict future constraints before they occur. This ability to predict a future constraint is useful for the virtual storage resource because actions that can relieve a virtual storage constraint generally require significant time to plan and implement. The report can also help you determine the effect of any actions, such as moving a large application above the 16-megabyte line or installing products that take advantage of expanded addressing.

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

## How to request this report

Monitor I gathers data for this report automatically with the default option VSTOR(S). See the *z/OS RMF User's Guide* for details. If you want to suppress gathering, you need to specify NOVSTOR.

To produce this report, specify

```
REPORTS(VSTOR(S))
REPORTS(VSTOR(D))
REPORTS(VSTOR(D[,jobname1,jobname2,...]))
REPORTS(VSTOR(jobname1[,jobname2,...]))
```

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

### Example URLs for the DDS API

```
http://ddshost:8803/gpm/rmfpp.xml?reports=VSTOR
http://ddshost:8803/gpm/rmfpp.xml?reports=VSTOR(D)
http://ddshost:8803/gpm/rmfpp.xml?reports=VSTOR(D,rmf)
```

## Contents of the report

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

## Data gathering considerations

To minimize overhead, RMF does not sample virtual storage data at every cycle. RMF takes one sample of virtual storage data for every ten RMF cycles. For example, if the RMF cycle is one second, RMF samples virtual storage data every ten seconds. In this case, RMF provides time stamps (accurate to within a ten-second range) for each minimum and maximum value on the report. The time stamp shows the time when RMF first observed the minimum or maximum value in the sample.

## Changing the sample to cycles ratio

### Programming Interface Information

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

Increasing the 1:10 ratio - for example, to 1:15 - decreases the accuracy of the virtual storage data RMF collects, in the sense that RMF is less likely to capture such data as a peak value when it samples less frequently. Increasing the ratio, however, does decrease RMF overhead. It is recommended to increase the ratio when measuring virtual storage activity of system address spaces like CATALOG, VTAM, DB2, IMS or other. This reduces the impact to the measured address space. Due to this impact, virtual storage activity of system address spaces should only be measured for a short period of time when diagnosing a special situation.

#### Example:

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

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

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

For more information on the use of AMASPZAP, see *z/OS MVS Diagnosis: Tools and Service Aids*.

### End Programming Interface Information

## Common Storage Summary section

The common storage summary section enables you to measure the use of virtual storage with minimal overhead. It contains the information you need to understand your current use of virtual storage. If you archive the data, you can use differences over time to predict a problem or constraint before it becomes critical. It also helps you to verify the size values set for CSA and SQA at IPL time and determine if you are using common storage effectively. Because RMF does not sample virtual storage data at every cycle, the value reported for NUMBER OF SAMPLES is less than the number of cycles.

VIRTUAL STORAGE ACTIVITY									
z/OS V2R2			SYSTEM ID TRX1		DATE 09/28/2016		INTERVAL 05.00.000		
			RPT VERSION V2R2 RMF		TIME 10.35.00		CYCLE 1.000 SECONDS		
COMMON STORAGE SUMMARY									
NUMBER OF SAMPLES 30									
STATIC STORAGE MAP			ALLOCATED CSA/SQA						
AREA	ADDRESS	SIZE	----- BELOW 16M -----		----- EXTENDED (ABOVE 16M) -----				
			MIN	MAX	AVG	MIN	MAX	AVG	
EPVT	21100000	1519M	SQA	284K 11.20.25	284K 11.20.25	284K	33.7M 11.24.05	33.9M 11.20.25	33.8M
ECSA	E4BC000	300M	CSA	404K 11.20.25	404K 11.20.25	404K	44.9M 11.20.45	44.9M 11.20.35	44.9M
EMLPA	0	0K							
EFLPA	E4B9000	12K							
EPLPA	A7F6000	60.8M	ALLOCATED CSA BY KEY						
ESQA	1C2B000	140M	0	148K 11.20.25	148K 11.20.25	148K	31.5M 11.20.25	31.5M 11.20.25	31.5M
ENUC	1000000	12.2M	1	128K 11.20.25	128K 11.20.25	128K	1844K 11.20.25	1844K 11.20.25	1844K
----- 16 MEG BOUNDARY -----			2	40K 11.20.25	40K 11.20.25	40K	4K 11.20.25	4K 11.20.25	4K
NUCLEUS	FD3000	180K	3	0K 11.20.25	0K	0K	0K 11.20.25	0K	0K
SQA	E52000	1540K	4	0K 11.20.25	0K	0K	4K 11.20.25	4K 11.20.25	4K
PLPA	C71000	1924K	5	4K 11.20.25	4K 11.20.25	4K	2828K 11.20.25	2828K 11.20.25	2828K
FLPA	C70000	4K	6	84K 11.20.25	84K 11.20.25	84K	9008K 11.20.45	9056K 11.20.35	9033K
MLPA	0	0K	7	0K 11.20.25	0K	0K	0K 11.20.25	0K	0K
CSA	900000	3520K	8-F	0K 11.20.25	0K	0K	0K 11.20.25	0K	0K
PRIVATE	2000	9208K	SQA EXPANSION INTO CSA						
PSA	0	8K		0K 11.20.25	0K	0K	0K 11.20.25	0K	0K
PLPA INTERMODULE SPACE - 2K IN PLPA AND 276K IN EPLPA									
PLPA SPACE REDUNDANT WITH MLPA/FLPA - 0K IN PLPA AND 11K IN EPLPA									
			----- BELOW 16M -----		----- ABOVE 16M -----				
			MIN	MAX	AVG	MIN	MAX	AVG	
CSA									
FREE PAGES (BYTES)	3116K 11.20.25	3116K 11.20.25	3116K	255M 11.20.35	255M 11.20.45	255M			
LARGEST FREE BLOCK	3116K 11.20.25	3116K 11.20.25	3116K	255M 11.20.35	255M 11.20.45	255M			
ALLOCATED AREA SIZE	404K 11.20.25	404K 11.20.25	404K	44.9M 11.20.45	44.9M 11.20.35	44.9M			
SQA									
FREE PAGES (BYTES)	1256K 11.20.25	1256K 11.20.25	1256K	106M 11.20.25	106M 11.24.05	106M			
LARGEST FREE BLOCK	896K 11.20.25	896K 11.20.25	896K	106M 11.20.25	106M 11.20.25	106M			
ALLOCATED AREA SIZE	644K 11.20.25	644K 11.20.25	644K	140M 11.20.25	140M 11.20.25	140M			
MAXIMUM POSSIBLE USER REGION - 9208K BELOW AND 1510M ABOVE			MAXIMUM POSSIBLE USER REGION - 9208K BELOW AND 1514M ABOVE						

Figure 222. VSTOR report - Common Storage Summary

Table 199. Fields in the Virtual Storage Activity Report - Common Storage Summary

Field Heading	Meaning
STATIC STORAGE MAP	The major storage areas above and below the 16-megabyte line. It includes the name of each area, the address of its lower boundary, and its size, reported in bytes.
ALLOCATED CSA/SQA	The MIN, MAX, and AVG values for allocated CSA and SQA, both below and above the 16-megabyte line. RMF calculates each size by adding the number of bytes assigned to each SQA or CSA subpool. The report also breaks down allocated CSA by key.
SQA EXPANSION INTO CSA	The MIN, MAX, and AVG size of any expansion of SQA into CSA.
PLPA INTERMODULE SPACE	The amount of unused space between the modules in both the PLPA and the EPLPA (the expanded PLPA). If your installation uses a pack list (in the IEAPAK00 Parmlib member), the values reported can help you determine the cost of your packing algorithm in relation to its benefit, a reduction in LPA paging rates, as shown in the paging report.
PLPA SPACE REDUNDANT WITH MLPA/FLPA	The amount of space for PLPA occupied by modules that also exist in (E)MLPA and/or (E)FLPA. For EPLPA, reports the amount of space occupied by modules that also exist in (E)MLPA or (E)FLPA.
FREE PAGES (BYTES)	The MIN, MAX, and AVG values, in bytes, for the amount of free storage.
LARGEST FREE BLOCK	The MIN, MAX, and AVG values, in bytes, for the size of the largest free block. The size of the largest free block, when compared to the total amount of free storage, is a measure of fragmentation within the common storage area. For example, when the size of the largest free block is close to the size of free storage, there is little fragmentation. The size of the largest free block is also the size of the largest GETMAIN that the system can currently satisfy within CSA or SQA.

Table 199. Fields in the Virtual Storage Activity Report - Common Storage Summary (continued)

Field Heading	Meaning
ALLOCATED AREA SIZE	The MIN, MAX, and AVG values, in bytes, for the size of the allocated area. RMF calculates this value as the difference between the highest and lowest address occupied by allocated storage. This includes all free blocks that lie between allocated blocks. Because free blocks between allocated blocks cause an increase in the virtual address range needed to hold the allocated blocks, consider this value when determining the size of CSA (and ECSA) and SQA (and ESQA). Significant fragmentation causes this number to be much larger than the amount of storage actually used. <b>Note:</b> Because the system allocates storage in the ESQA area for both ends of the address range, the allocated area size is always the same as the total size.
MAXIMUM POSSIBLE USER REGION	The largest size specified on the REGION= JCL parameter that this system can satisfy (assuming a minimal number of DD statements). RMF determines this value, reported for below and above the 16-megabyte line, by examining its own private area; it calculates the size by finding the difference between the bottom of its allocated area at RMF initialization and the start of the user region. There is no guarantee that a particular job can obtain a region of the reported size. For a job that runs under an initiator (as opposed to a job that runs as a started task), the storage that the initiator obtains, reduces the size of the region that the job can obtain. The number of DD statements in the JCL can also reduce the size of the region. The reported region size, however, can indicate whether a job with a known region requirement is likely to obtain the region it requires under the system conditions reflected in the report.

## Overview reference

Table 200. Overview names in the Virtual Storage Activity Report

Field Heading or Meaning	Overview Name
SQA EXPANSION INTO CSA	SQAE
LARGEST FREE BLOCK - MIN, CSA	CSAFB
LARGEST FREE BLOCK - MIN, SQA	SQAFB

## Common Storage Detail section

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

VIRTUAL STORAGE ACTIVITY																
z/OS V2R2		SYSTEM ID TRX1		DATE 09/28/2016		INTERVAL 05.00.000		PAGE 2								
		RPT VERSION V2R2 RMF		TIME 10.35.00		CYCLE 1.000 SECONDS										
COMMON STORAGE DETAIL																
ALLOCATED CSA BY SUBPOOL BY KEY (BELOW 16 MEG)				ALLOCATED SQA BY SUBPOOL (BELOW 16M)												
SUBPOOL 227		SUBPOOL 228		SUBPOOL 231		SUBPOOL 241		SUBPOOL	MIN	MAX	AVG					
----- MINIMUM -----																
0	28K	11.20.25	4K	11.20.25	4K	11.20.25	112K	11.20.25	226	24K	11.20.25	24K	11.20.25	24K	11.20.25	24K
1			4K	11.20.25			124K	11.20.25	239	36K	11.20.25	36K	11.20.25	36K	11.20.25	36K
2						40K	11.20.25		245	224K	11.20.25	224K	11.20.25	224K	11.20.25	224K
3																
4																
5							4K	11.20.25								
6	12K	11.20.25	52K	11.20.25	12K	11.20.25	8K	11.20.25								
7																
8-F																
ALL	40K	11.20.25	60K	11.20.25	16K	11.20.25	288K	11.20.25								
----- MAXIMUM -----																
0	28K	11.20.25	4K	11.20.25	4K	11.20.25	112K	11.20.25								
1			4K	11.20.25			124K	11.20.25								
2						40K	11.20.25									
3																
4																
5							4K	11.20.25								
6	12K	11.20.25	52K	11.20.25	12K	11.20.25	8K	11.20.25								
7																
8-F																
ALL	40K	11.20.25	60K	11.20.25	16K	11.20.25	288K	11.20.25								
----- AVERAGE -----																
0	28K		4K		4K		112K									
1			4K				124K									
2							40K									
3																
4																
5							4K									
6	12K		52K		12K		8K									
7																
8-F																
ALL	40K		60K		16K		288K									

Figure 223. VSTOR report - Common Storage Detail

Table 201. Fields in the Virtual Storage Activity Report - Common Storage Detail Section

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

## Private Area Summary section

The optional private area summary section presents information about how a specific address space is using its private virtual storage. RMF uses the job name you specify when you request the report to identify the address space. In choosing a job, note that gathering data for a specific address space requires additional RMF overhead. Note that RMF can gather private area data only when a job is active at the beginning of the interval, and various conditions can limit RMF's ability to report complete private area data. These conditions, and the actions RMF takes, are described later in this section under "Partial private area data" on page 461.

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

**Note:** Measuring virtual storage activity for a specific job may have significant impact on the performance of the job. When requesting a VSTOR report, system address spaces like CATALOG, VTAM, DB2, IMS or other, should be specified as *jobname* only for a short period of time when diagnosing a special performance

situation.

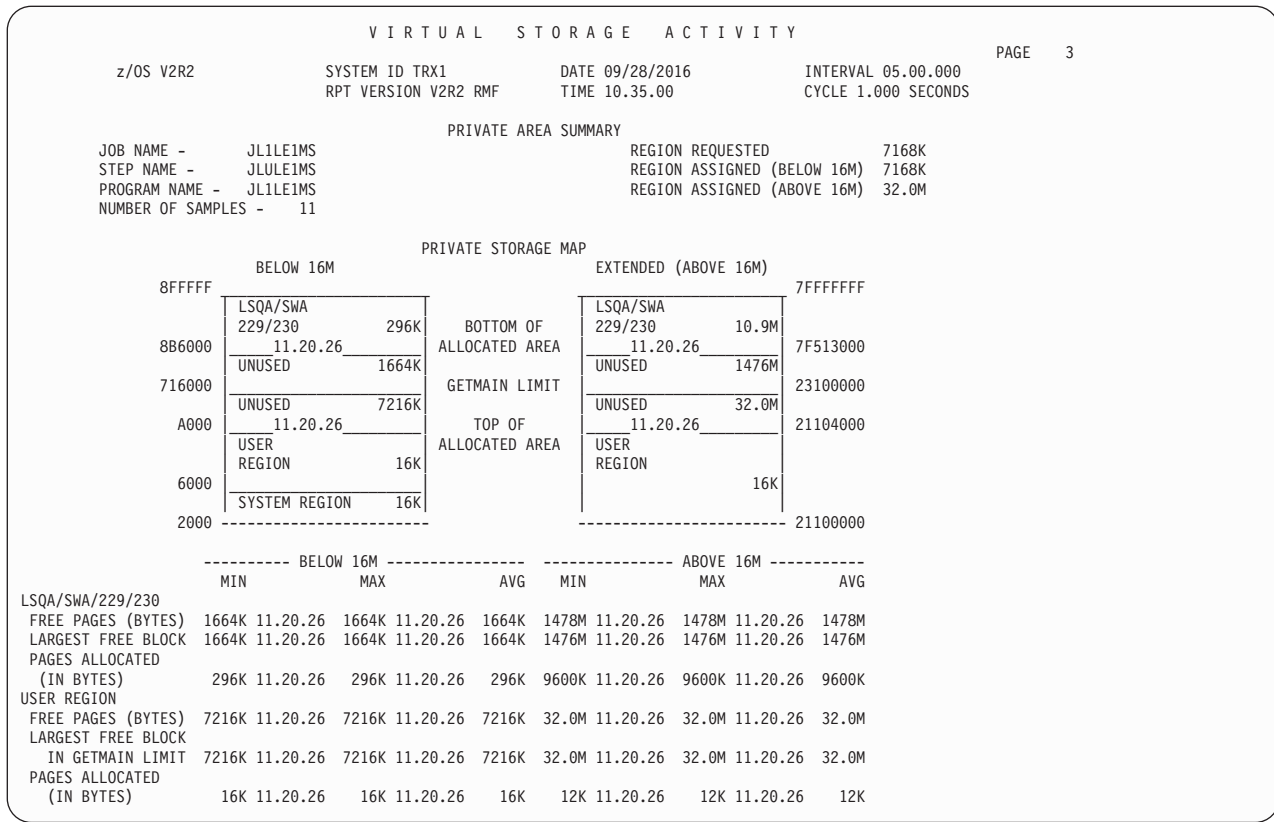


Figure 224. VSTOR report - Private Area Summary

Table 202. Fields in the Virtual Storage Activity Report - Private Area Summary

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

### Private Storage Map

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

Table 203. Fields in the Virtual Storage Activity Report - Private Storage Map

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

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

### Bottom half of Private Storage Map

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

Table 204. Fields in the Virtual Storage Activity Report - Bottom Half

Field Heading	Meaning
FREE PAGES (BYTES)	The MIN, MAX, and AVG values, in bytes, for the amount of free storage.
LARGEST FREE BLOCK	The MIN, MAX, and AVG values, in bytes, for the size of the largest free block in the available amount of free storage.
PAGES ALLOCATED (IN BYTES)	The MIN, MAX, and AVG values for the amount of allocated storage. <b>Note:</b> The MAX value for PAGES ALLOCATED has the same meaning as fields in SMF record type 30. For LSQA, these fields are SMF30ARB and SMF30EAR. For the user region, these fields are SMF30URB and SMF30EUR. However, RMF might report a smaller number than SMF does. RMF reports the highest value that it sampled while SMF reports the highest value that occurred. If the highest value occurred when RMF was not taking a sample, RMF misses the actual peak value.

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

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

### Partial private area data

Private area reporting works best for jobs that are running at least one interval. You can, of course, monitor other jobs, but there are some conditions that mean RMF can collect little or no data. These conditions, which are related to the way virtual storage reporting works, are:

1. RMF searches for any requested jobs at the beginning of each interval. If it does not find a job, it does not monitor the job during the interval. In this case, RMF



issues a message to the operator and produces a report. The report, however, contains no data; instead, the following message appears:

JOB WAS NOT ACTIVE AT THE BEGINNING OF THIS INTERVAL

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

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

2. If a job that RMF is monitoring terminates and is then restarted, the report for the interval in which it terminated includes data only up to the point when the job terminated. RMF resumes its monitoring of the restarted job at the beginning of the interval following the interval during which the job was restarted.
3. If a job RMF is monitoring is swapped out at the time RMF takes a sample of virtual storage data, RMF does not cause a swap-in; it skips the sample for that job. Thus, the number of samples for a swappable job may be less than expected. If a job is swapped out every time RMF tries to take a sample during an interval, RMF reports no data for that interval.

### Private Area Detail section

The optional *Private Area Detail* section provides information about the number of bytes of allocated blocks by area below the 16-megabyte line. In the header, the job name and the memory limit in bytes for this address space is displayed.

VIRTUAL STORAGE ACTIVITY					PAGE	4
z/OS V2R2	SYSTEM ID TRX1	DATE 09/28/2016	INTERVAL 05.00.000			
	RPT VERSION V2R2 RMF	TIME 10.35.00	CYCLE 1.000 SECONDS			
PRIVATE AREA DETAIL						
JOB NAME - RMFGAT	MEMORY LIMIT - 16384P					
NUMBER OF BYTES OF	ALLOCATED BLOCKS	BY AREA (BELOW 16 MEG)				
SUBPOOL (AREA)	MIN	MAX	AVG			
230	76K 15.55.00	76K 15.55.00	76K			
236 (SWA)	84K 15.55.00	84K 15.55.00	84K			
237 (SWA)	24K 15.55.00	24K 15.55.00	24K			
255 (LSQA)	20K 15.55.00	20K 15.55.00	20K			
USER REGION						
0	112K 15.55.00	112K 15.55.00	112K			
1	12K 15.55.00	12K 15.55.00	12K			
4	4K 15.55.00	4K 15.55.00	4K			
9	8K 15.55.00	8K 15.55.00	8K			
10	8K 15.55.00	8K 15.55.00	8K			
252 (REENTRANT)	728K 15.55.00	728K 15.55.00	728K			
HIGH VIRTUAL MEMORY USAGE (ABOVE 2GB)						
BYTES	MIN	MAX	AVG	PEAK		
PRIVATE	7.000M 15.55.00	7.000M 15.55.00	7.000M	7.000M		
SHARED	0 15.55.00	0	0	0		
COMMON	0 15.55.00	0	0	0		
MEMORY OBJECTS						
PRIVATE	6 15.55.00	6 15.55.00	6			
SHARED	0 15.55.00	0	0			
COMMON	0 15.55.00	0	0			
FIXED 1 MB	0 15.55.00	0	0			
SHARED 1 MB	0 15.55.00	0	0			
1 MB FRAMES						
FIXED	0 15.55.00	0	0			
PAGEABLE	0 15.55.00	0	0			

Figure 225. VSTOR report - Private Area Detail and High Virtual Memory Usage (above 2GB)

Table 205. Fields in the Virtual Storage Activity report - Private Area Detail section

Field Heading	Meaning
JOB NAME	The job analyzed in the report.



Table 205. Fields in the Virtual Storage Activity report - Private Area Detail section (continued)

Field Heading	Meaning
MEMORY LIMIT	The memory limit in bytes for this address space.
NUMBER OF BYTES OF ALLOCATED BLOCKS BY AREA (BELOW 16 MEG)	
SUBPOOL (AREA) / USER REGION	The MIN, MAX, and AVG values for the number of bytes of allocated blocks during the report interval, broken down by subpool and by area (LSQA, SWA).

## High Virtual Memory Usage section

The optional *High Virtual Memory Usage* section provides information about the high virtual memory usage above the 2-gigabyte line.

See Figure 225 on page 462 for an example of the High Virtual Memory Usage section.

Table 206. Fields in the Virtual Storage Activity report - High Virtual Memory Usage section

Field Heading	Meaning
BYTES	The MIN, MAX, and AVG values for the number of bytes in <b>PRIVATE</b> , <b>SHARED</b> , and <b>COMMON</b> memory objects allocated with the indicated job as the owner. In addition, the PEAK useable storage since the start of the job are shown. Unlike MIN, MAX and AVG, the PEAK values report usable storage only. Hidden storage, like guard areas, are not included so that the reported PEAK value may be smaller than the MIN, MAX or AVG values.
MEMORY OBJECTS	The MIN, MAX, and AVG values for the number of <b>PRIVATE</b> , <b>SHARED</b> , and <b>COMMON</b> memory objects allocated with the indicated job. Fixed memory objects and shared memory objects that are backed in 1 MB frames are also reported if the Enhanced DAT Architecture is supported.
1 MB FRAMES	The MIN, MAX, and AVG values for the number of 1 MB frames that are used by <b>FIXED</b> and <b>PAGEABLE</b> memory objects with the indicated job as the owner (only available with Enhanced DAT Architecture). The PAGEABLE value also includes 1 MB frames that are used by DREF memory objects.  Frames that are either used by shared 1 MB pages or to satisfy 4 KB space requests on a constrained system are not included.

## WLMGL - Workload Activity report

The Workload Activity report (WLMGL) can be used to request a variety of reports, as shown in “The WLMGL option list” on page 464.

### How to request this report

Monitor I gathers data for this report automatically. If you want to suppress gathering, you need to specify NOWKLD.

To produce this report, specify  
SYSRPTS(WLMGL(options))

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

### Example URL for the DDS API

[http://ddshost:8803/gpm/rmfpp.xml?reports=WLMGL\(SCLASS\(BATCH\)\)](http://ddshost:8803/gpm/rmfpp.xml?reports=WLMGL(SCLASS(BATCH)))

For more information, see the *z/OS RMF User's Guide*.

## The WLMGL option list

You can select among various types of WLMGL reports by specifying the SYSRPTS WLMGL options:

### SCPER

All service class periods found for a service class.

This report contains detailed information about:

- Subsystem delays
- Response time goals vs actuals
- General execution delays
- Response time distribution.

### SCLASS

Summary of data for all service class periods defined for a service class.

### WGROU

Summary of data for all service classes defined in a workload definition.

### POLICY

Summary of data for all workloads defined in the active service policy.

### WGPER

All service classes, including one line for each service class period, defined in a workload definition.

### RCLASS

All report classes defined in a service policy.

### RCPER

All periods found for a report class. The report has the same structure and information as the Service Class Period report.

**RTD** Response Time Distribution is displayed in WLMGL Service/Report Class Period reports (RTD is default, NORTD otherwise). This suboption can only be specified together with suboptions SCPER or RCPER; otherwise, it is ignored.

In addition, you have the **SYSNAM** option to select systems to be included in the report. All systems must belong to the same sysplex.

## Contents of the report

The following sections describe the various WLMGL report types as introduced in "The WLMGL option list." "Field descriptions for all reports" on page 471 describes the report headers and explains all of the fields displayed by these reports.

Note that for all report types, the reporting interval is finished in case of a policy change or a policy refresh. A policy refresh is initiated either by the operator or by the Workload Manager component itself when a processor speed change or a IEAOPTxx parameter change occurs.

## Service Class Period report

To request a Service Class Period report, specify:

```
SYSRPTS(WLMGL(SCPER(service_class)))
```

For example, to produce a report for all service class periods of service class STCLOW, specify:

```
SYSRPTS(WLMGL(SCPER(STCLOW)))
```

Figure 226 shows a report for service class STCLOW where the first service class period (PERIOD=1) is defined with an execution velocity goal. For service class periods with an execution velocity goal, the WLMGL report displays one tabular response time distribution for each system in the sysplex. The example depicted contains two response time distributions for systems TRX1 and TRX2.

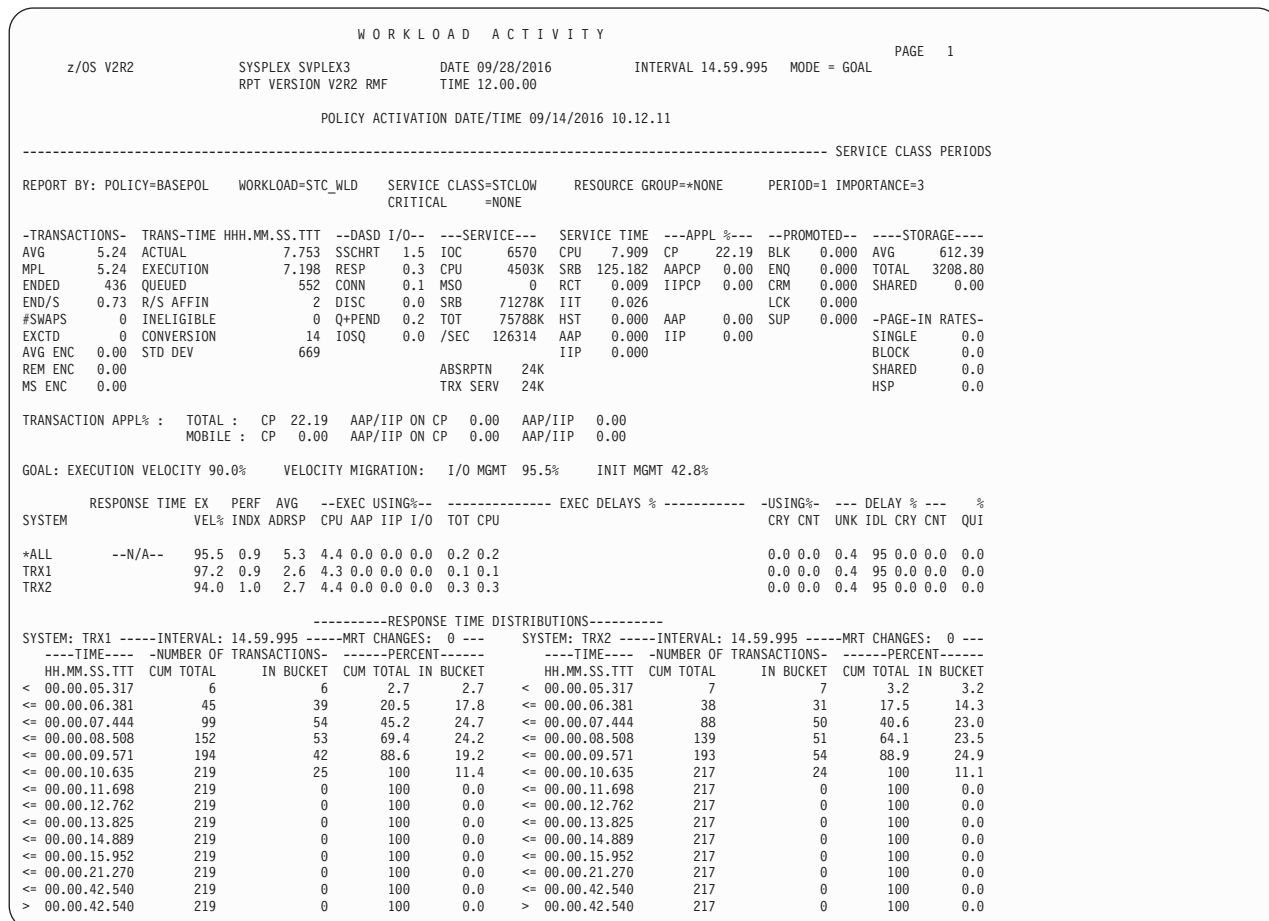


Figure 226. WLMGL - Service Class Period report - with execution velocity goal for Period 1 and response time distributions



Figure 228 shows a report for service class CICSLOW with subsystem delay data included, where the first period is defined with a percentile response time goal.

**Note:** For transaction service classes of subsystem work managers, like CICS or IMS, no service consumption and execution delays are reported. This data is reported with the service classes for the regions.

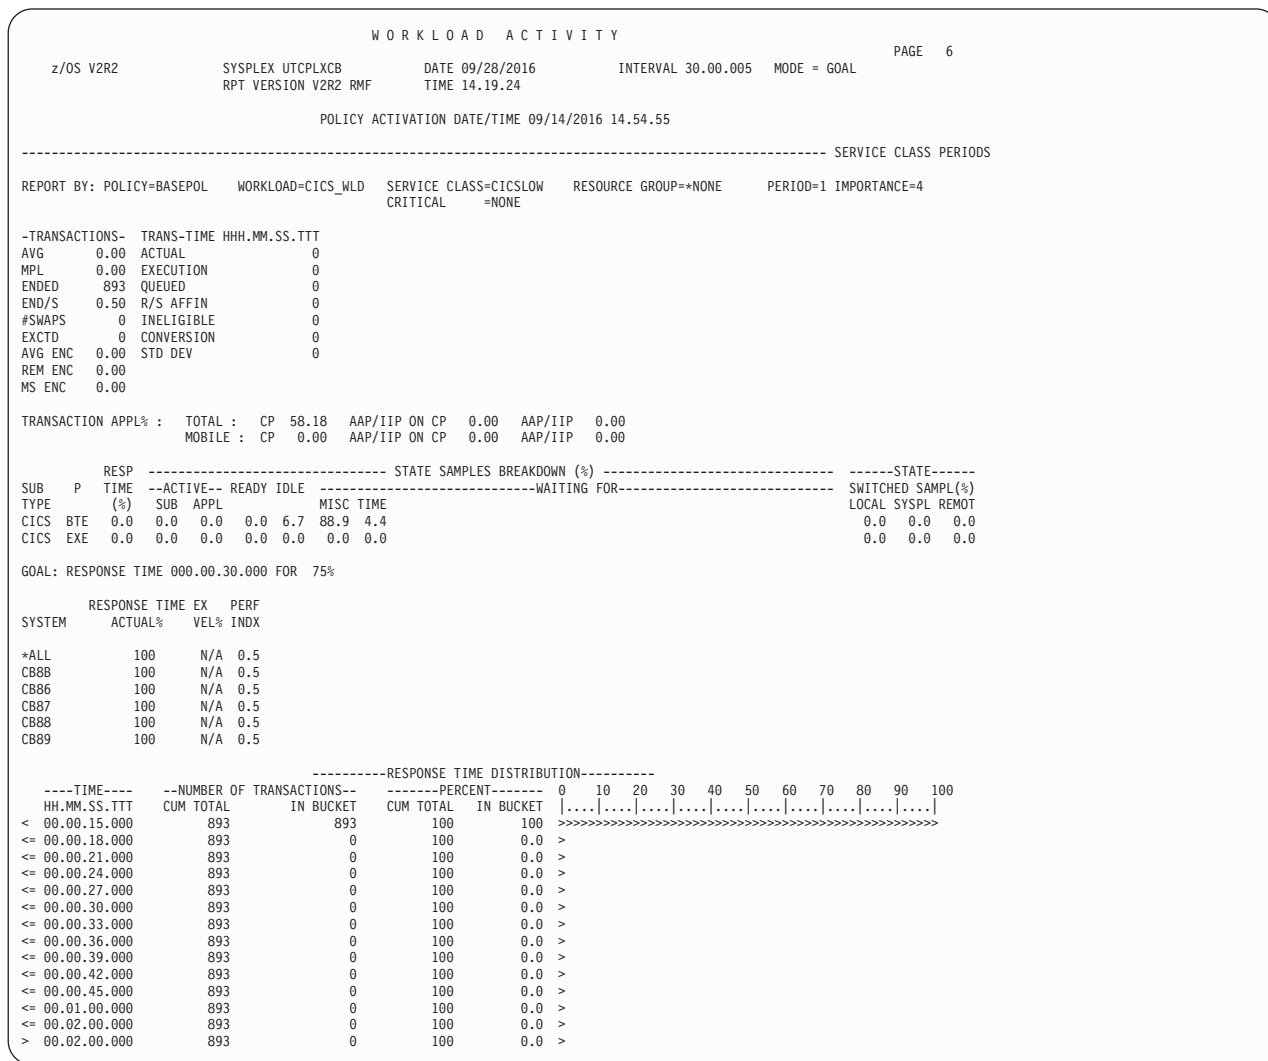


Figure 228. WLMGL - Service Class Period report - with subsystem data and response time distribution for response time goal

## Service Class report

To request this report, specify:  
 SYSRPTS(WLMGL(SCLASS(service\_class)))

For example, to produce the report shown in Figure 229 on page 468, specify:  
 SYSRPTS(WLMGL(SCLASS(STCLOW)))

The report can contain an additional part SERVICE CLASSES BEING SERVED if address spaces in the service class in this report are doing work for transactions that were classified to a another service class.

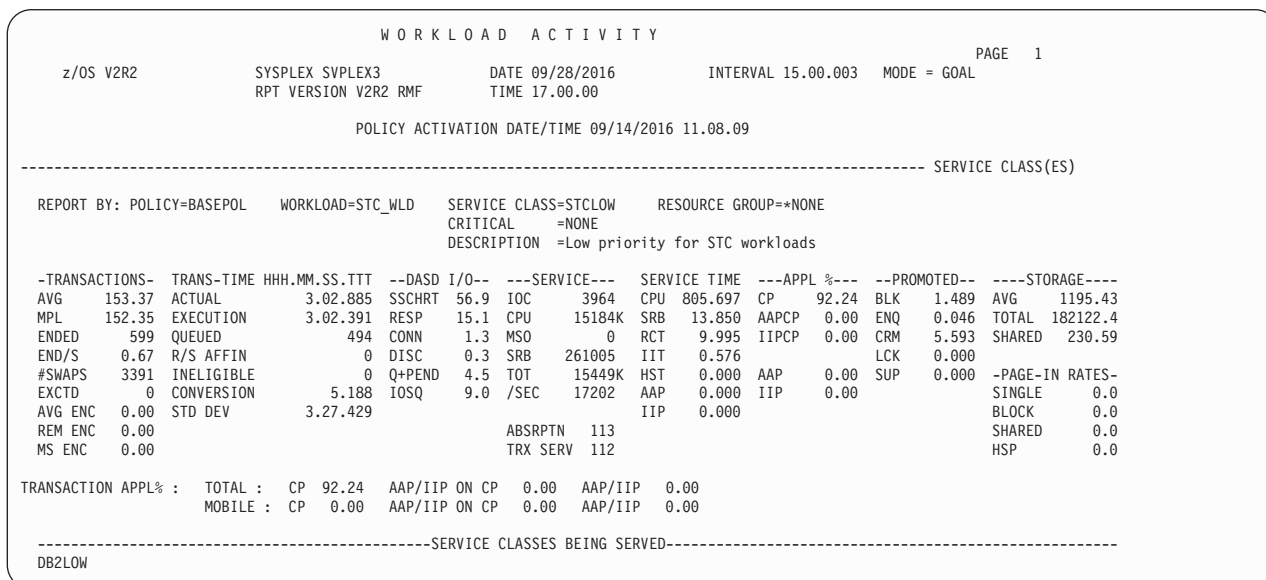


Figure 229. WLMGL - Service Class Report

## Workload Group report

To request this report, specify:

SYSRPTS(WLMGL(WGROUP(workload\_group)))

For example, to produce the report for a workload group called OMVS\_WLD, specify:

SYSRPTS(WLMGL(WGROUP(OMVS\_WLD)))

The report has the same layout as a Service Class report, but all service classes associated with the workload OMVS\_WLD are combined in a workload summary.

## Workload Group and Service Class Period report

To request this report, specify:

SYSRPTS(WLMGL(WGPER(workload\_group)))

Each service class associated with the specified workload group is listed with its service class periods and the defined and achieved GOALS for the periods.

For example, to produce the report for a workload group called STC\_WLD with its associated service classes STCHIGH and STCLOW, specify:

SYSRPTS(WLMGL(WGPER(STC\_WLD)))

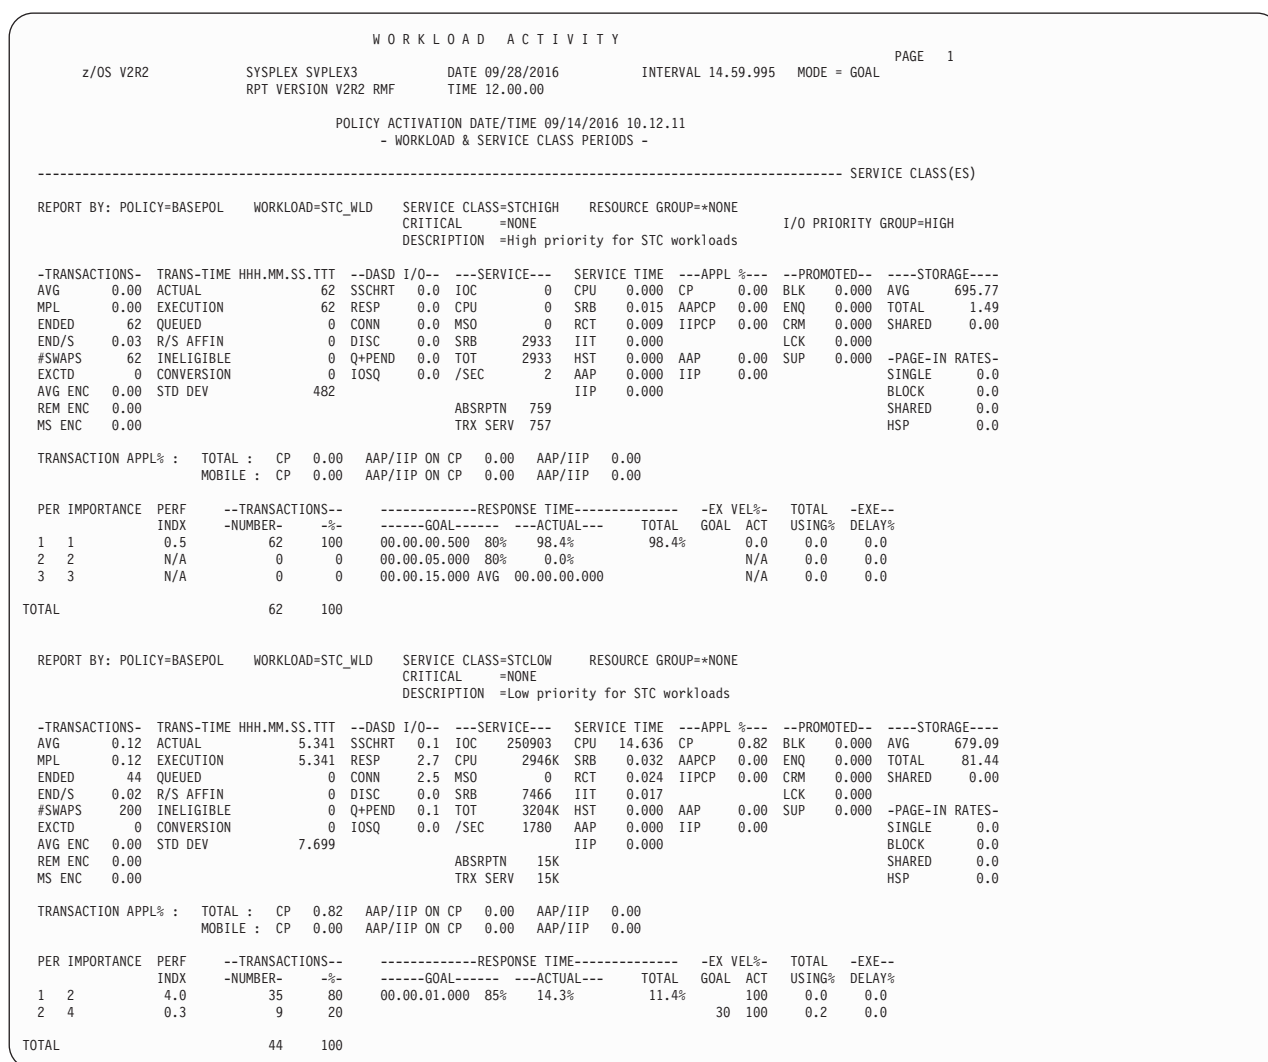


Figure 230. WLMGL - Workload Group with associated service classes

## Policy Summary report

To request this report, specify:

SYSRPTS(WLMGL(POLICY(policy)))

For example, to produce a report for a policy named HOLIDAY, specify:

SYSRPTS(WLMGL(POLICY(HOLIDAY)))

Except for the separation line, a Policy Summary report has the same layout as the Workload Group report, and combines all workload groups associated to the policy.

## Report Class report

To request this report, specify:

SYSRPTS(WLMGL(RCLASS(report\_class)))

For example, to produce a report for a report class called BATCH, specify:

SYSRPTS(WLMGL(RCLASS(BATCH)))

Optionally, classification rules can assign incoming work to a report class. Report classes are for additional reporting data, across service classes, or for monitoring special work.

The report has nearly the same contents as the Service Class report, just the workload group name cannot be associated to a report class and is therefore not reported.

### Report Class Period report

To request this report, specify:

```
SYSRPTS(WLMGL(RCPER(report_class)))
```

The report has nearly the same contents as the Service Class Period report, but reporting of subsystem delays and response time distribution data is possible only for homogeneous report class periods.

All goal-related data in the report is based on the corresponding service class period.

A report class period is called **homogeneous** if all its transactions are being assigned to the same service class period.

**Example:** You classify all TSO users to run in service class TSOPROD and distinguish the departments for reporting purposes in report classes TSOEPTA, TSOEPTB, and TSOEPTC. This definition done in the WLM application creates homogeneous report classes.

All other report class periods are called **heterogeneous**.

### Service Policy page

The Service Policy page is created automatically for each interval. This page contains the following information:

- Information about the installation of the service definition
- Service definition coefficients and normalization factors for special purpose processors
- Information about each system, including the system ID, parmlib member suffix, service units/second, effective capacity percentage, and the interval start time and length
- Resource group definitions
- Resource group actuals for each service class

This page is always included at the end of an interval.



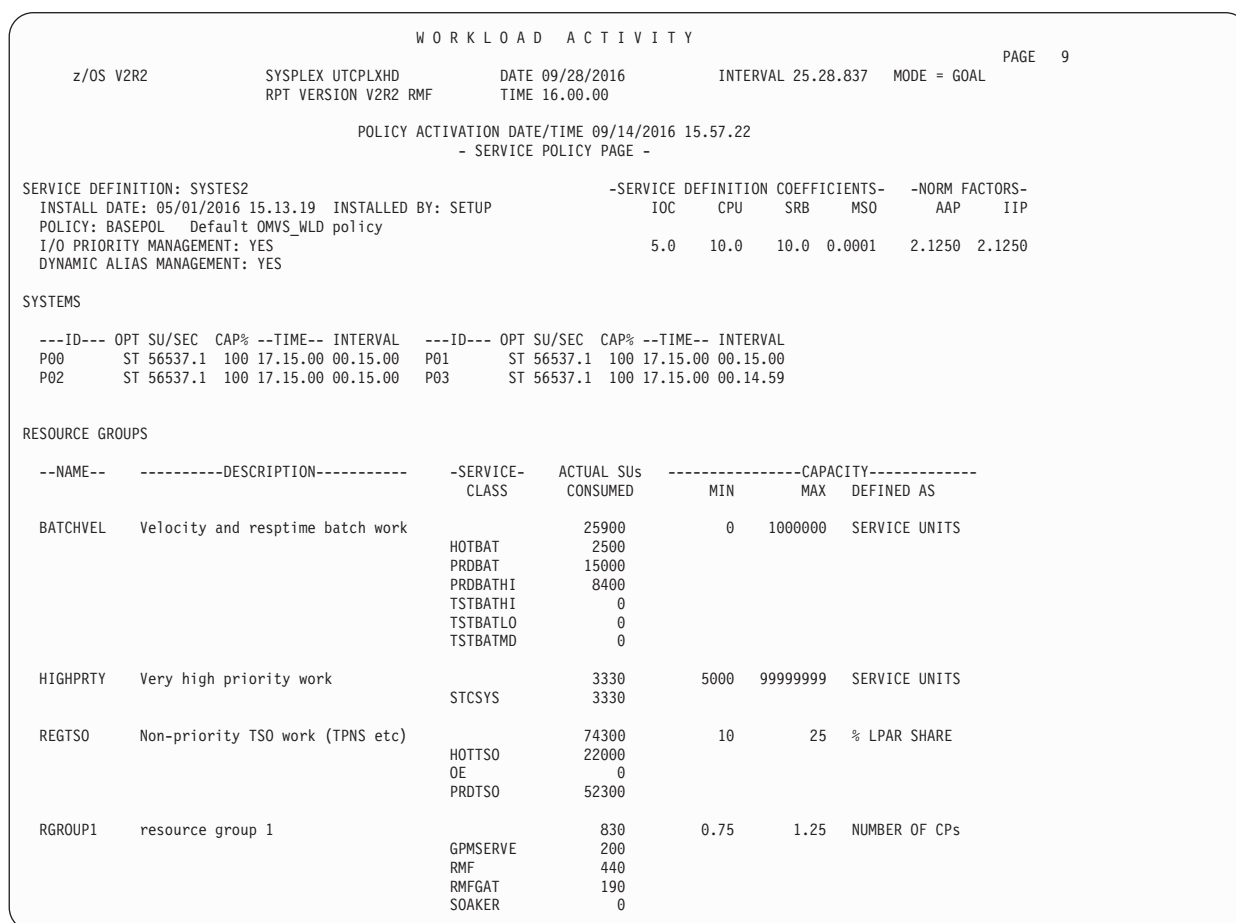


Figure 231. WLMGL Report - Service Policy Page

## Field descriptions for all reports

Table 207. Fields in the Workload Activity Report

Field Heading	Meaning
CRITICAL	This indication reports on the specification in the service policy, not whether SRM is using or ignoring it. This information can be found in Monitor II (ARD report) and in Monitor III (for example, DELAY report). <b>CPU</b> CPU critical <b>STORAGE</b> Storage critical <b>STORAGE + CPU</b> Both storage and CPU critical <b>NONE</b> Neither storage nor CPU critical
I/O PRIORITY GROUP=HIGH	This indication is reported for service classes assigned to I/O priority group HIGH in the active service policy.

Table 207. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
TRANSACTIONS	<p>Number of transactions.</p> <p><b>AVG</b> The average number of active transactions during the interval including independent enclaves. On a period level, this is the transaction active time for the single period, divided by the RMF interval time. Otherwise, it is the sum of the transaction active time for all summarized periods, divided by the RMF interval time.</p> <p>To get the average number of transactions related to active address spaces, subtract <b>AVG ENC</b> from <b>AVG</b>.</p> <p><b>MPL</b> The average number of transactions resident in central storage during the interval. On a period level, it is the transaction residency time for a single period, divided by the RMF interval time. Otherwise, it is the sum of the transaction residency time for all summarized periods, divided by the RMF interval time.</p> <p>To get the average number of transactions related to resident address spaces, subtract <b>AVG ENC</b> from <b>MPL</b>.</p> <p><b>ENDED</b> The number of transactions that ended during the interval. On a period level, this is the number of transactions that ended during that period. Otherwise, it is the total number of transactions that ended for all the summarized periods.</p> <p><b>END/S</b> The number of transactions that ended per second.</p> <p><b>#SWAPS</b> The total number of swaps. On a period level it is the number of swaps occurred during the single period. Otherwise, it is the sum of the number of SWAPS that occurred during all summarized periods.</p> <p><b>EXCTD</b> Count of times a subsystem work manager reported that an execution phase has completed. A single transaction could have zero or more execution phases.</p>
AVG ENC	<p>The average number of independent enclaves during the interval. From a sysplex scope, this is the sum of active time for enclaves that originated on the respective system either for the single period or for all summarized periods divided by the RMF interval time.</p>
REM ENC	<p>The average number of foreign enclaves during the interval. From a sysplex scope, this is the sum of active time for enclaves that originated on a remote system in the sysplex, but are executing on the respective system either for the single period or for all summarized periods divided by the RMF interval time.</p>
MS ENC	<p>Average number of multi-system enclaves during the interval. From a sysplex scope, this is the sum of active time for enclaves that originated on the respective system and are executing on one or more remote systems in the sysplex in parallel either for the single period or for all summarized periods divided by the RMF interval time.</p>

Table 207. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
TRANS.-TIME	<p>Transaction time in HHH.MM.SS.TTT units. All times are reported in the period the transaction ended.</p> <p>The time a job was delayed due to TYPRUN=HOLD or TYPRUN=JCLHOLD is NOT included in any of the transaction times.</p> <p><b>ACTUAL</b> The actual amount of time required to complete the work submitted under the service class. This is the total response time including EXECUTION, QUEUED, R/S AFFIN, and INELIGIBLE.</p> <p><b>EXECUTION</b> The average execution time of ended transactions.</p> <p><b>QUEUED</b> Average time a job was delayed for reasons other than the ones mentioned below. This field therefore basically includes the time a job was delayed for initiation.</p> <p>For TSO users, this can be a portion of LOGON processing. For APPC this is the time the transaction spent on an APPC queue.</p> <p><b>R/S AFFIN</b> Average time the job was delayed due to resource or system affinity scheduling delay. This means that resource(s) required for the job to run were not available at some point while the job was queued to JES2.</p> <p><b>INELIGIBLE</b> Average time the job was delayed due to operational delays or JES scheduling delays, examples are:</p> <ul style="list-style-type: none"> <li>• Job held by operator</li> <li>• Job class or job queue held</li> <li>• Duplicate jobname serialization</li> <li>• Job class execution limits</li> </ul> <p><b>CONVERSION</b> Average time the job was delayed due to JCL conversion.</p> <p>Jobs held during conversion (due to affinity, HSM recall, or enqueue contention) contribute only to conversion time, not to ineligible or R/S affinity times.</p> <p>CONVERSION time is not included in the total response time (ACTUAL).</p> <p><b>STD DEV</b> Standard deviation of ACTUAL.</p> <p>Standard deviation is a measure of variability of the data in the sample. The higher the standard deviation, the more spread out it looks on a graph.</p>
DASD I/O	<p>Information about DASD I/O activities.</p> <p><b>SSCHRT</b> Number of start subchannels SSCH per second in the reported interval.</p> <p><b>RESP</b> Average DASD response time (in milliseconds) of the transactions in this group. This is the sum of the average connect time (CONN), the average disconnect time (DISC), the average wait time (Q+PEND), and the IOS queue time (IOSQ).</p> <p><b>CONN</b> Average DASD connection time of the transactions in this group, as reported by the channel measurement subsystem.</p> <p><b>DISC</b> Average DASD disconnect time of the transactions in this group, as reported by the channel measurement subsystem.</p> <p><b>Q+PEND</b> Average DASD wait time (queue time + pending time) of the transactions in this group. This does not include IOSQ time, as reported by the channel measurement subsystem.</p> <p><b>IOSQ</b> Average time the transactions in this group spent on the IOS queue, based on sampled delays.</p>

Table 207. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
SERVICE	<p>The service is calculated by multiplying the received service units with the appropriate service definition coefficient.</p> <p><b>IOC</b> Total amount of input/output service received.</p> <p><b>CPU</b> Total amount of task and preemptible-class SRB processor service received.</p> <p><b>MSO</b> Total amount of main storage occupancy service received.</p> <p><b>SRB</b> Total amount of non-preemptible SRB service received.</p> <p><b>TOT</b> Sum of CPU, SRB, IOC, and MSO service.</p> <p><b>/SEC</b> Rate at which service is provided in service units per second.</p> <p><b>ABSRPTN</b> Absorption rate at which service is used while transactions are resident in main storage. This is the total service divided by the transaction residency time.</p> <p><b>TRX SERV</b> Rate at which service is used by transactions that are active, but not necessarily in storage. This is the total service divided by the transaction active time.</p>
SERVICE TIME	<p>This category is made up of the following:</p> <p><b>CPU</b> Task and preemptible-class SRB (enclave) time in seconds consumed on general purpose and special purpose processors.</p> <p><b>SRB</b> Service request block time in seconds.</p> <p><b>RCT</b> Region control task time in seconds.</p> <p><b>IIT</b> I/O interrupt time in seconds.</p> <p><b>HST</b> Hiperspace service time in seconds.</p> <p><b>AAP</b> zAAP service time in seconds.</p> <p><b>IIP</b> zIIP service time in seconds.</p> <p><b>Note:</b></p> <ol style="list-style-type: none"> <li>1. If special purpose processors are running faster than general purpose processors, AAP and IIP times are not normalized.</li> <li>2. Normalized AAP and IIP times are included in CPU time.</li> </ol>

Table 207. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
APPL%	<p>Percentage of the processor time used by transactions running on the different processor types. The calculation is:</p> $\text{APPL\%} = \frac{\text{Processor time used}}{\text{Interval length} * \text{Multithreading maximum capacity factor}} * 100$ <p><b>CP</b> Percentage of the processor time used by transactions running on general purpose processors in the service or report class period. The calculation of the processor time is based on the time values displayed under field heading SERVICE TIME.</p> <p>Processor time used = CPU + SRB + RCT + IIT + HST – (AAPNF * AAP) – (IIPNF * IIP)</p> <p>The AAP and IIP times may be normalized to general purpose processor time from a faster zAAP or zIIP where AAPNF and IIPNF are the zAAP and zIIP normalization factors. They can be found under field heading NORM FACTORS on the POLICY page of the Workload Activity report.</p> <p><b>AAPCP</b> Percentage of the processor time used by zAAP eligible transactions running on general purpose processors. This is a subset of APPL% CP.</p> <p><b>IIPCP</b> Percentage of the processor time used by zIIP eligible transactions running on general purpose processors. This is a subset of APPL% CP.</p> <p><b>AAP</b> Percentage of the processor time used by transactions executed on zAAPs in the service or report class period.</p> <p><b>IIP</b> Percentage of the processor time used by transactions executed on zIIPs in the service or report class period.</p> <p><b>Notes:</b></p> <ol style="list-style-type: none"> <li>1. APPL% shows processor utilization based on uniprocessor capacity. This means that the values can exceed 100% in systems with more than one processor.</li> <li>2. If the multithreading mode is set to 1, a multithreading maximum capacity factor of 1 is used for the APPL% calculation.</li> <li>3. The interval length in a sysplex is the common interval length.</li> <li>4. In a sysplex, the values for seconds and CPU time percentages are meaningful only if all processors have the same speed and the multithreading mode is the same on all systems. You can use the SYSRPTS WLMGL SYSNAM option to select only a subset of the systems to be included in the report.</li> <li>5. AAPCP or IIPCP may report values greater than zero even if no special purpose processors are configured or if they are varied offline, because the PROJECTCPU option is specified in the active IEAOPT Parmlib member. This information can be used to understand the benefit of adding special purpose processors to your system.</li> </ol>
PROMOTED	<p>CPU time in seconds that transactions in this group were running at a promoted dispatching priority, separated by the reason for the promotion:</p> <p><b>BLK</b> CPU time in seconds consumed while the dispatching priority of work with low importance was temporarily raised to help blocked workloads</p> <p><b>ENQ</b> CPU time in seconds consumed while the dispatching priority was temporarily raised by enqueue management because the work held a resource that other work needed.</p> <p><b>CRM</b> CPU time in seconds consumed while the dispatching priority was temporarily raised by chronic resource contention management because the work held a resource that other work needed</p> <p><b>LCK</b> In HiperDispatch mode, the CPU time in seconds consumed while the dispatching priority was temporarily raised to shorten the lock hold time of a local suspend lock held by the work unit.</p> <p><b>SUP</b> CPU time in seconds consumed while the dispatching priority for a work unit was temporarily raised by the z/OS supervisor to a higher dispatching priority than assigned by WLM.</p>

Table 207. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
STORAGE	<p>Amount of storage frames.</p> <p><b>AVG</b> Weighted average number of central and expanded storage frames allocated to active ASIDs. This value is the sum of the number of central and expanded frames weighted by the transaction residency time for each active ASID, divided by the total transaction residency time.  <b>Note:</b> Enclave transaction residency or active time is not included in the calculation of this value.</p> <p><b>TOTAL</b> Total number of central and expanded storage frames allocated to resident ASIDs. This value is the sum of the total number of central and expanded frames weighted by the transaction residency time for each active ASID, divided by the RMF interval time.</p> <p><b>SHARED</b>                      Total number of shared storage pages allocated to resident ASIDs.</p>
PAGE-IN RATES	<p><b>SINGLE</b>                      The average rate at which pages are read into central storage while transactions are resident in central storage. On a single period level this is the total number of page-ins during the period, divided by transaction residency time. For all other levels it is the sum of the total number of page-ins for all periods summarized, divided by the sum of the transaction residency time for all periods being summarized.</p> <p><b>BLOCK</b> Rate of demand page-ins from DASD for blocked pages, expressed in pages per seconds.</p> <p><b>SHARED</b>                      Rate of shared storage page-ins</p> <p><b>HSP</b> Rate of standard hiperspace pages read into central storage from auxiliary storage.  <b>Note:</b> Enclave transaction residency time is not included in the calculation of these values because there is no paging on behalf of enclaves.</p>
<i>Transaction Application Time %</i>	
<p>When transaction processor usage is reported to WLM through IWM4RPT or IWM4MNTF services, the consumed service units are accounted to the transaction service or report classes, and deducted from the region's service and report classes. If the number of transactions is very small and a single transaction reports high processor times, it can occur that processor times become negative. In such a case RMF, displays asterisk (*).</p>	
TOTAL	<p>Total percentage of the processor time used by transactions running on the different processor types.</p> <p><b>CP</b> Total percentage of general purpose processor time used by transactions.</p> <p><b>AAP/IIP ON CP</b>                      Total percentage of general purpose processor time used by transactions eligible to run on specialty processors.</p> <p><b>AAP/IIP</b>                      Total percentage of specialty processor time used by transactions.</p>
MOBILE	<p>Percentage of the processor time used by transactions classified with reporting attribute MOBILE running on the different processor types. MOBILE is a subset of TOTAL.</p> <p><b>CP</b> Percentage of general purpose processor time used by transactions classified with reporting attribute MOBILE.</p> <p><b>AAP/IIP ON CP</b>                      Percentage of general purpose processor time used by transactions classified with reporting attribute MOBILE, eligible to run on specialty processors.</p> <p><b>AAP/IIP</b>                      Percentage of specialty processor time used by transactions classified with reporting attribute MOBILE.</p>
<i>Service Classes being Served</i>	
SERVICE CLASSES BEING SERVED	<p>This section is only available if address spaces are doing work for transactions that were classified to another service class. The name of each service class being served by the reported service class (see name in the separation line) is displayed.</p>
<i>Work Manager/Resource Manager</i>	

Table 207. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
SUB TYPE	The name (for example CICS or IMS) represents the subsystem type (4 characters) as used in the classification rules in the WLM administration application. The subsystem's documentation should explain the meaning that product attributes to the specific states.
P	The phase identified as <b>BTE</b> indicates the representation of the states incurred in the begin-to-end phase of a transaction <b>EXE</b> indicates the representation of the states incurred in the execution phase of a transaction.
RESP TIME (%)	The transaction response time percentage in either the <i>BEGIN-TO-END</i> phase, or the <i>EXECUTION</i> phase.
STATE SAMPLES BREAKDOWN (%)	Identifies the percentages of samples that a transaction has been detected in the reported states.  <b>ACTIVE SUB</b> The active subsystem state sample percentage. Active indicates that there is a program executing on behalf of the work request from the perspective of the work manager. This does not mean that the program is active from the BCP's perspective.  <b>ACTIVE APPL</b> The active application state sample percentage in contrast to the active subsystem state sample percentage. This allows a subsystem to differentiate between work requests processed by the subsystem itself ( <b>ACTIVE SUB</b> ) and work requests processed by an application invoked by the subsystem.  <b>READY</b> The ready state sample percentage. Ready indicates that there is a program ready to execute on behalf of the work request described by the monitoring environment, but the work manager has given priority to another work request.  <b>IDLE</b> The idle state sample percentage. Idle indicates that no work request (or transaction) is allowed to run.

Table 207. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
STATE SAMPLES BREAKDOWN (%) - WAITING FOR	<p>STATE SAMPLES BREAKDOWN (%) - continued.</p> <p>This category presents up to fifteen named delay reason states having the highest non-zero values. These are sorted by the sum of the BTE and EXE rows in each column. The values of the remaining less important states will be accumulated, if applicable, and presented as delay reason OTHR.</p> <p>Here is a list explaining the delay reasons.</p> <p><b>LOCK</b> waiting for lock.</p> <p><b>I/O</b> waiting for I/O indicates that the work manager is waiting on an activity related to an I/O request. This may either be an actual I/O operation or some function associated with an I/O request.</p> <p><b>CONV</b> waiting for conversation could have been used in conjunction with IWMMSWCH to identify where the target is located.</p> <p><b>DIST</b> waiting for distributed request indicates at a high level that some function or data must be routed prior to resumption of the work request. This is to be contrasted with 'waiting on conversation', which is a low level view of the precise resource that is needed. A distributed request could involve 'waiting on conversation' as part of its processing.</p> <p><b>LOCL</b> waiting for a session to be established locally, for example, on the current MVS image.</p> <p><b>SYSP</b> waiting for a session to be established somewhere in the sysplex.</p> <p><b>REMT</b> waiting for a session to be established somewhere in the network.</p> <p><b>TIME</b> waiting for timer.</p> <p><b>LTCH</b> waiting for a latch.</p> <p><b>PROD</b> waiting for another product.</p> <p><b>MISC</b> waiting for unidentified resource, possibly among another specific category, but which may not be readily determined.</p> <p><b>SSLT</b> waiting for an SSL thread.</p> <p><b>REGT</b> waiting for a regular thread.</p> <p><b>WORK</b> waiting for registration to a work table.</p> <p><b>BPMI</b> waiting for I/O resulting from a DB2 buffer pool miss.</p> <p><b>TYPn/TYnn</b> The generic delay state defined by the subsystem (possible values: TYP1 - TYP9 and TY10 - TY15). If the subsystem uses the WLM service IWM4MGDD (Define Descriptions for Generic Delay States) to provide a description for a generic delay state, RMF displays a legend with the delay state description. If the subsystem did not use the IWM4MGDD service, the legend is omitted. For further explanation of the generic delay state types please refer to the subsystem documentation.</p>
STATE SWITCHED SAMPL(%)	<p>Subsystem state samples - continued</p> <p><b>LOCAL</b> State representing transactions for which there are logical continuations on this MVS image. Subsystems might set this state when they function ship a transaction to another component within the same MVS image.</p> <p><b>SYSP</b> State representing transactions for which there are logical continuations on another MVS image in the sysplex. Subsystems might set this state when they function ship a transaction to another component on another image in the sysplex.</p> <p><b>REMOT</b> State representing transactions for which there are logical continuations somewhere within the network. Subsystems might set this state when they function ship a transaction to another component within the network.</p>
<p><i>Service or Report Class period: goal and actual values</i></p>	
<p>If measurement data for systems in a sysplex is available, this section starts with an *ALL line showing the average or cumulated values for the sysplex. The *ALL line is followed by one line for each system.</p>	



Table 207. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
GOAL	<p>This line shows the goal specified in the WLM service policy for a service class period. For information about available goal types refer to <i>z/OS MVS Planning: Workload Management</i>.</p> <p>In a report for a homogeneous report class period, the goal of the corresponding service class period is printed. For heterogeneous report class periods, N/A is printed.</p>
VELOCITY MIGRATION	<p>The following two values are only provided for periods with an execution velocity goal:</p> <p><b>I/O MGMT</b> I/O Priority Management</p> <p>Value of achieved execution velocity including I/O using and delay samples.</p> <p>If WLM I/O delay management is enabled in the service definition, this value matches EX VEL%. Otherwise, this is the value that would be observed if WLM I/O management were enabled and no other changes that would affect the execution velocity calculation were made.</p> <p>You see your current definition in the Service Policy page (see Figure 231 on page 471).</p> <p><b>INIT MGMT</b> Initiator Management</p> <p>Value of achieved execution velocity including batch initiator delay samples.</p> <p>If WLM batch initiator management is enabled in the service definition, this value matches EX VEL%. Otherwise, this is the value that would be observed if WLM batch initiator management were enabled and no other changes that would affect the execution velocity calculation were made.</p> <p>You find a description of these delays in this table for the field TRANS.-TIME.</p>
RESPONSE TIME	<p>This column either shows:</p> <ul style="list-style-type: none"> <li>• for an AVG response time goal: the measured average response times</li> <li>• for a percentile response time goal: the percentages of the transactions that met the response time goal</li> <li>• for an execution velocity goal, a system or a discretionary goal: N/A</li> <li>• for heterogeneous report class periods: N/A</li> </ul>
EX VEL %	<p>The execution velocity measures the portion of the acceptable processor and storage delays relative to the total execution time. For details about the execution velocity, see “Common Monitor III report measurements” on page 12.</p>
PERF INDX	<p>The performance index for a period represents how close a period came to reaching the goal (PI is 1.0 if goal is reached), and how much this period suffered versus its goal. See Table 82 on page 191 for more details about the performance index.</p>
AVG ADRSP	<p>Average number of address spaces and enclaves that contributed delay and using samples to this class.</p>
EXEC USING%	<p>The following using samples are measured as percentages of the total samples:</p> <p><b>CPU</b> Standard CP using samples. This value includes using samples of zAAP and zIIP work executing on general purpose processors (standard CPs).</p> <p><b>AAP</b> zAAP using samples.</p> <p><b>IIP</b> zIIP using samples.</p> <p><b>I/O</b> I/O using samples.</p> <p><b>Note:</b></p> <ol style="list-style-type: none"> <li>1. Use the APCUSGP (AAP on CP Using%) overview condition to retrieve the using samples of zAAP work executing on general purpose processors (standard CPs).</li> <li>2. Use the IPCUSGP (IIP on CP Using%) overview condition to retrieve the using samples of zIIP work executing on standard CPs.</li> </ol>

Table 207. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
EXEC DELAYS %	<p>General execution delays included in TOT (total). Each dispatchable unit sampled can increase one of the CPU or paging delay samples. Besides the TOT value, only the seven highest values contributing to TOT will be shown. The remaining less important values will be accumulated and presented as OTH.</p> <p><b>TOT</b> Total delay used by SRM in its execution velocity calculation.</p> <p><b>CPU</b> CPU delay. A TCB or SRB is waiting to be dispatched (other than the first in-line behind sampler), or a TCB is waiting for a LOCAL lock.</p> <p><b>AAP</b> zAAP-eligible work is delayed because it is waiting for a processor that can run zAAP work.</p> <p><b>IIP</b> zIIP-eligible work is delayed because it is waiting for a processor that can run zIIP work.</p> <p><b>I/O</b> I/O delay. A TCB or SRB has initiated an I/O request that is delayed obtaining a path to the device. This includes IOSQ and Q+PEND components (see "CONN" on page 473 for a description).  <b>Note:</b> It depends on the definition in WLM whether this value is part of the TOTAL value or not, by default it is <b>not</b> contained in TOTAL.</p> <p><b>CAP</b> CPU capping delay. A TCB or SRB is marked non-dispatchable because</p> <ul style="list-style-type: none"> <li>• a resource group maximum is being enforced</li> <li>• or because of discretionary goal management. That is, if certain types of work are overachieving their goals, that work may be capped so that the resources may be diverted to run discretionary work (see also section 'Using Discretionary Goals' in <i>z/OS MVS Planning: Workload Management</i>).</li> </ul> <p>This value is NOT part of the CPU delay.</p> <p><b>SIN</b> Swap-In delay. Swap-In has started but not completed.</p> <p><b>MPL</b> MPL delay. Ready but swap-in has not started.</p> <p><b>Q MPL</b> Queue MPL - work is waiting for a server address space or batch initiator.</p> <p><b>SRV PRV</b> Private area paging delay for a server address space.</p> <p><b>SRV VIO</b> VIO paging delay for a server address space.</p> <p><b>SRV SHS</b> Hiperspace paging delay for a server address space.</p> <p><b>SRV SIN</b> Swap-in delay for a server address space.</p> <p><b>SRV MPL</b> MPL delay for a server address space.</p> <p><b>AUX PRV</b> Auxiliary paging from private.</p> <p><b>AUX COM</b> Auxiliary paging from common.</p> <p><b>AUX XME</b> Auxiliary paging from cross memory.</p> <p><b>AUX VIO</b> Auxiliary paging from VIO.</p> <p><b>AUX SHS</b> Auxiliary paging from standard hiperspaces.</p> <p><b>AUX EHS</b> Auxiliary paging from ESO hiperspaces (a page being read was not in the ESO hiperspace, it has to be read from DASD by the program managing the hiperspace).</p>
USING%	<p>Percentage of using states:</p> <p><b>CRY</b> Crypto using state — a TCB or SRB was found to be using a cryptographic asynchronous message processor (CAP) or an adjunct processor (AP).</p> <p><b>CNT</b> Contention using state - work is holding resources.</p>

Table 207. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
DELAY %	<p>The following states are NOT included in the TOTAL EXECUTION DELAYS.</p> <p><b>UNK</b> State is unknown. The address space or enclave was not found to be using or delayed for any sampled resource, but z/OS has not been notified that it is idle.</p> <p><b>IDL</b> Idle state. Work is in STIMER wait, TSO terminal wait, APPC wait, OMVS input or output wait, or an initiator is waiting for work.</p> <p><b>CRY</b> Crypto delay state — a TCB or SRB was found to be waiting for a CAP, an AP or a processor feature queue.</p> <p><b>CNT</b> Contention delay state - work is waiting for resources.</p>
% QUI	<p>Quiesce state. Some work in this period has been RESET with the QUIESCE keyword. This is the percentage of address spaces and enclaves quiesced during the reporting interval.</p>
RESPONSE TIME DISTRIBUTION (for service/report class periods with a response time goal only)	<p>WLM maintains counts of how many transactions were completed within a particular time.</p> <p>The <i>response time goal</i> defined for each service class period is split into 14 response time buckets where:</p> <ul style="list-style-type: none"> <li>• bucket 1 covers the gap from 0 to half the goal</li> <li>• buckets 2 to 11 cover the gap between half the goal to 1.5 times the goal evenly divided</li> <li>• bucket 12 covers two times the goal</li> <li>• bucket 13 covers four times the goal</li> <li>• bucket 14 covers the gap from four times the goal to infinity</li> </ul> <p>The chart presents the sysplex-wide view on the</p> <ul style="list-style-type: none"> <li>• number of total (ended) transactions,</li> <li>• response time,</li> <li>• number of total (ended) transactions in percent,</li> <li>• and a graphical illustration of the percentage.</li> </ul> <p><b>TIME</b> Response time associated to this bucket.</p> <p><b>NUMBER OF TRANSACTIONS</b>  Number of transactions that completed for this period.</p> <p><b>CUM TOTAL</b>  Cumulative number of transactions so far</p> <p><b>IN BUCKET</b>  Number of transactions in this bucket</p> <p><b>PERCENT</b>  Percentage</p> <p><b>CUM TOTAL</b>  Cumulative percentage of transactions so far</p> <p><b>IN BUCKET</b>  Percentage of transactions associated to the bucket</p> <p><b>Percent Scale</b>  Graphical presentation of each bucket</p>

**PP - WLMGL**

Table 207. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
RESPONSE TIME DISTRIBUTIONS (for service/report class periods with an execution velocity goal only)	<p>WLM maintains counts of how many transactions were completed within a particular time.</p> <p>The <i>average of all response times (midpoint)</i> measured for each service/report class period is split into 14 response time buckets where:</p> <ul style="list-style-type: none"> <li>• bucket 1 covers the gap from 0 to half the midpoint</li> <li>• buckets 2 to 11 cover the gap between half to 1.5 times the midpoint evenly divided</li> <li>• bucket 12 covers two times the midpoint</li> <li>• bucket 13 covers four times the midpoint</li> <li>• bucket 14 covers the gap from four times the midpoint to infinity</li> </ul> <p>For each system in the sysplex, the report presents a tabular representation of:</p> <ul style="list-style-type: none"> <li>• the number of total (ended) transactions</li> <li>• the response time</li> <li>• the number of total (ended) transactions in percent.</li> </ul> <p><b>SYSTEM</b> System Name</p> <p><b>INTERVAL</b> Measurement interval (Time since last midpoint change)</p> <p><b>MRT CHANGES</b> Number of midpoint changes during the SMF interval</p> <p><b>TIME</b> Response time associated to this bucket</p> <p><b>NUMBER OF TRANSACTIONS</b> Number of transactions that completed for this period</p> <p><b>CUM TOTAL</b> Cumulative number of transactions so far</p> <p><b>IN BUCKET</b> Number of transactions in this bucket</p> <p><b>PERCENT</b> Percentage</p> <p><b>CUM TOTAL</b> Cumulative percentage of transactions so far</p> <p><b>IN BUCKET</b> Percentage of transactions associated to the bucket</p>

Table 208. Fields in the WLMGL Report - POLICY

Field Heading	Meaning
<i>Service Policy Page</i>	
SERVICE DEFINITION	Service definition name and description. The service definition includes defined goals for each of the service classes in the workload.
INSTALL DATE	Service definition installation date and time.
INSTALLED BY	Userid and system name that last installed this service definition.
SERVICE DEFINITION COEFFICIENTS	Service definitions coefficients as defined in the service policy to determine the: <b>IOC</b> Number of countable EXCP instructions. <b>CPU</b> Task processor time. <b>SRB</b> SRB processor time. <b>MSO</b> Approximate storage use for each service class period.
NORM FACTORS	Normalization factors for special purpose processors: <b>AAP</b> Normalization factor for zAAP. Multiply zAAP service times or service units with this value to calculate the CP equivalent value. <b>IIP</b> Normalization factor for zIIP. Multiply zIIP service times or service units with this value to calculate the CP equivalent value.
POLICY	Policy name and description.

Table 208. Fields in the WLMGL Report - POLICY (continued)

Field Heading	Meaning
I/O PRIORITY MANAGEMENT	<p><b>YES</b> Workload management dynamically manages I/O priorities based on service class goals and importance. Workload management also includes I/O delays in the calculation of execution velocity.</p> <p><b>NO</b> I/O priorities are the same as dispatching priorities.</p>
DYNAMIC ALIAS MANAGEMENT	YES if WLM dynamic management of parallel access volumes is active, otherwise <b>NO</b> .
SYSTEMS	<p>The system names contributing to the sysplex report.</p> <p><b>ID</b> System name.</p> <p><b>OPT</b> IEAOPTxx Parmlib member suffix (parameter description that control resource and workload management algorithms in the system resources manager).</p> <p><b>SU/SEC</b> Nominal capacity rating in service units per second per online CPU.</p> <p><b>CAP%</b> Percentage of effective capacity available to the CPU.</p> <p>The value is 100, if the machine is working at its full nominal capacity. If the machine is working in power-save mode or cycle-steering mode, the value is less than 100. If the nominal or effective processor capacity cannot be determined, N/A is reported.</p> <p><b>TIME</b> Begin time of the interval for this system.</p> <p><b>INTERVAL</b> Interval length for this system in HH.MM.SS.</p>
RESOURCE GROUPS	<p>A resource group is an amount of processor capacity across one or more MVS systems. The report includes resource group name and description as:</p> <p><b>NAME</b> Resource group name.</p> <p><b>DESCRIPTION</b> Resource group name description.</p> <p><b>SERVICE CLASS</b> Service class name associated with this RESOURCE GROUP.</p> <p><b>ACTUAL SUs CONSUMED</b> Unweighted CPU and SRB service units consumed per second by all service classes in the resource group.</p> <p>Unweighted CPU and SRB service units consumed per second by each service class. <b>Note:</b> The reported service units do not include service units consumed on zAAPs or zIIPs.</p> <p><b>CAPACITY</b> Resource group capacity limits.</p> <p><b>MIN</b> Minimum amount of service that the resource group should receive if demand exists.</p> <p><b>MAX</b> Maximum amount of service that the resource group should be allowed to consume.</p> <p><b>DEFINED AS</b> The method how the resource group's capacity is defined:</p> <ul style="list-style-type: none"> <li>• SERVICE UNITS: in unweighted CPU and SRB service units across the sysplex</li> <li>• % LPAR SHARE: as percentage of the LPAR share (scope of the resource group is system-wide)</li> <li>• NUMBER OF CPs: as number of CPs (scope of the resource group is system-wide).</li> </ul>

## Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*. The following table shows the exception and overview condition names for the Overview report.

Table 209. Exception and Overview names in the Workload Activity Report

Field Heading or Meaning	Overview Name
TRANSACTIONS -	
- AVG	TRANSAVG
- MPL	TRANSMPL
- ENDED	TRANSTOT
- END/S	TRANS
- #SWAPS	SPERTRA
- AVG ENC	ENCAVG
- REM ENC	ENCREM
- MS ENC	ENCMS
TRANSACTION TIME -	
- ACTUAL	RTIMETOT
- EXECUTION	RTIME
- QUEUED	RTIMEQUE
- R/S AFFIN	TRANSADT
- INELIGIBLE	TRANSIQT
- CONVERSION	TRANSCVT
DASD I/O -	
- SSCHRT	SSCHRT
- RESP	RESP
- CONN	CONN
- DISC	DISC
- Q+PEND	QPEND
- IOSQ	IOSQ
SERVICE -	
<b>Note:</b> The Overview output is a rate, not a count. For details see <i>z/OS RMF User's Guide</i> .	
- IOC	IOSRV
- CPU	CPUSRV
- MSO	MSOSRV
- SRB	SRBSRV
- TOT	TOTSRV
- ABSRPTN	ABSRPTN
- TRX SERV	TRXSERV
SERVICE TIME -	
- CPU	TCBSEC/TCBPER
- SRB	SRBSEC/SRBPER
- RCT	RCTSEC
- IIT	IITSEC
- HST	HSTSEC
- AAP	AAPSEC/AAPNSEC
- IIP	IIPSEC/IIPNSEC
PROMOTED -	
- BLK	PROMSEC/PROMPER
- ENQ	EPROMSEC/EPROMPER
- CRM	CPROMSEC/CPROMPER
- LCK	LPROMSEC/LPROMPER

Table 209. Exception and Overview names in the Workload Activity Report (continued)

Field Heading or Meaning	Overview Name
- SUP	SPROMSEC/SPROMPER
APPL % -	
- CP	APPLSEC/APPLPER
- AAPCP	APPLIFCP/APPLAPCP
- IIPCP	APPLIPCP
- AAP	APPLAAP
- IIP	APPLIIP
- AAP time on general purpose processors	AAPCPSEC
- Total number of EXCPs	EXCP
- EXCP rate	EXCPRT
STORAGE -	
- TOTAL	STOTOT
- CENTRAL	STOCEN
- SHARED	STOSHR
PAGE-IN RATES -	
- SINGLE	SINGLE
- BLOCK	BLOCK
- SHARED	SHARED
- HSP	HSP
EX VEL %	EXVEL
PERF INDEX	PI
USING% -	
- CPU	CPUUSGP
- AAP	AAPUSGP
- IIP	IIPUSGP
- I/O	IOUSGP
- AAP on CP	APCUSGP
- IIP on CP	IPCUSGP
EXECUTION DELAYS % -	
- CPU	CPUDLYP
- AAP	AAPDLYP
- IIP	IIPDLYP
- I/O	IODLYP
- CAPP	CAPP
- SWIN	SWINP
- MPL	MPLP
- QMPL	QUEUEP
- Total Server Delays	SERV
- Total Storage Delays	STOP
DLY% -	
- UNKN	UNKP
- IDLE	IDLEP
CRYPTO% -	
- DLY	CRYDLYP, CAPDLYP, APDLYP, FQDLYP
- HSP	HSP

I

Table 209. Exception and Overview names in the Workload Activity Report (continued)

Field Heading or Meaning	Overview Name
TRANSACTION APPL% -	
- TOTAL CP	TAPPLCP
- TOTAL AAP/IIP ON CP	TAPPLOCP
- TOTAL AAP/IIP	TAPPLSP
- MOBILE CP	MAPPLCP
- MOBILE AAP/IIP ON CP	MAPPLOCP
- MOBILE AAP/IIP	MAPPLSP

## XCF - Cross-System Coupling Facility Activity report

The Cross-System Coupling Facility Activity report shows the XCF data from one system's processing in a sysplex. To better understand the traffic on corresponding outbound and inbound signalling paths, you might have to run RMF reports on two or more systems.

### How to request this report

Monitor III gathers data for this report automatically. If you want to suppress gathering, you have to disable writing SMF record type 74.2.

To produce this report, specify  
REPORTS(XCF)

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

### Example URL for the DDS API

<http://ddshost:8803/gpm/rmfpp.xml?reports=XCF>

### Contents of the report

The XCF Activity report is divided into three sections:

- XCF Usage by System
- XCF Usage by Member
- XCF Path Statistics

The **Usage by System** section gives information about messages sent to and received from each remote system in the sysplex, broken down by transport class. Use this section to check the class lengths and message buffer space parameters. For a sample of the XCF Activity Usage by System section, see Figure 232 on page 487.

The **Usage by Member** section gives information about messages sent to and from each remote system, broken down by remote group and member, and summarizes messages sent and received by the local system (the local system is the system on which the data was collected) broken down by local group and member. Use this section to check message traffic loads associated with groups and members, and check for groups that are candidates to be put in their own transport classes. For a sample of the XCF Activity Usage by Member section, see Figure 233 on page 488.



The **Path Statistics** section describes messages sent to and from each remote system, broken down by signalling path. Use this report to determine whether the number of XCF signalling paths are sufficient for the message traffic. For a sample of the XCF Activity Path Statistics section, see Figure 234 on page 489.

**Note:** If the XCF system, path, or member becomes inactive during the RMF interval, the appropriate counters will be reinitialized. This is indicated in the report by the message \*COUNTS RESET.

X C F A C T I V I T Y														
z/OS V2R2		SYSTEM ID TRX1			DATE 09/28/2016			INTERVAL 15.00.000			PAGE 1			
		RPT VERSION V2R2 RMF			TIME 09.15.00			CYCLE 1.000 SECONDS						
XCF USAGE BY SYSTEM														
REMOTE SYSTEMS										LOCAL				
OUTBOUND FROM TRX1					INBOUND TO TRX1					TRX1				
TO SYSTEM	TRANSPORT CLASS	BUFFER LENGTH	REQ OUT	% SML	% FIT	% BIG	% OVR	ALL PATHS UNAVAIL	REQ REJECT	FROM SYSTEM	REQ IN	REQ REJECT	TRANSPORT CLASS	REQ REJECT
TRX2	BIG	40,892	92	100	0	0	0	0	0	TRX2	843	0	BIG	0
	DB2	956	0	0	0	0	0	0	0				DB2	0
	DEFAULT	956	720	0	100	0	0	0	0				DEFAULT	0
	FEWFAST	956	0	0	0	0	0	0	0				FEWFAST	0
	JES2	956	0	0	0	0	0	0	0				JES2	0
	TCCONS	956	8	0	63	38	100	8	0				TCCONS	0
	TCGRS	956	36	0	100	0	0	0	0				TCGRS	0
	TCOPC	956	0	0	0	0	0	0	0				TCOPC	0
	TCRMF	62,464	151	30	70	0	0	0	0				TCRMF	0
	TCVLF	956	0	0	0	0	0	0	0				TCVLF	0
TOTAL			1,007							TOTAL	843			

Figure 232. XCF Activity Report - Usage by System

Table 210. Fields in the XCF Activity Report - Usage by System

Field Heading	Meaning
TO SYSTEM	The name of the system to which the messages were sent.
TRANSPORT CLASS	The name of the transport class used by XCF for the message transfer. There is one line in the report for each transport class for each target system. <b>Note:</b> If you request the XCF Activity report as a duration report, the Postprocessor inserts an * into this column, if the transport class has not been active during the complete duration interval.
BUFFER LENGTH	The internally defined message buffer size XCF uses for the transport class. The buffer size is the maximum length of the messages which can be contained in the buffers currently being used for this transport class.
REQ OUT	The total number of messages that XCF accepted for delivery to the system in the indicated transport class.
% SML	The percentage of messages sent that could have used a smaller BUFFER LENGTH than their defined BUFFER LENGTH.
% FIT	The percentage of messages sent that fit the defined BUFFER LENGTH.
% BIG	The percentage of messages sent that needed a BUFFER LENGTH larger than the defined BUFFER LENGTH.  The value is reported as '<1' if the percentage is greater than 0 but rounded to 0.
% OVR	The percentage of BIG messages sent that suffered performance degradation. If the messages are bigger than the defined transport class BUFFER LENGTH, XCF must find a buffer large enough to contain the BIG message, thus causing overhead. If enough BIG messages are sent, XCF dynamically adjusts the BUFFER LENGTH to avoid this overhead.

PP - XCF

Table 210. Fields in the XCF Activity Report - Usage by System (continued)

Field Heading	Meaning
ALL PATHS UNAVAIL	The number of messages that XCF had to migrate to a signalling path in an alternate transport class because there was no operational signalling path connected to the target system and assigned to the indicated transport class.
REQ REJECT	The number of requests for a message buffer that could not be satisfied due to constraints on the amount of message buffer space. This field appears under the INBOUND TO, the OUTBOUND FROM, and the LOCAL headings in the Usage by System section of the report. Under OUTBOUND FROM, it indicates the number of requests to send a message to a particular remote system that were rejected in a particular transport class. Under INBOUND TO, it indicates the number of time XCF could not get an inbound message buffer in anticipation of receiving a new message. Under LOCAL, it indicates the number of requests to send a message within the local system that were rejected in a particular transport class.
FROM SYSTEM	The name of system sending the message. There is one line in the report for each system in the sysplex.
REQ IN	For each system, the total number of messages that were received from the remote system in the sysplex.
TRANSPORT CLASS	The name of the transport class that XCF uses for the message transfer. There is one line in the report for each transport class for each target system.

XCF ACTIVITY									
z/OS V2R2		SYSTEM ID TRX1		DATE 09/28/2016		INTERVAL 15.00.000		PAGE 2	
		RPT VERSION V2R2 RMF		TIME 09.15.00		CYCLE 1.000 SECONDS			
XCF USAGE BY MEMBER									
MEMBERS COMMUNICATING WITH TRX1					MEMBERS ON TRX1				
GROUP	MEMBER	SYSTEM	REQ FROM TRX1	REQ TO TRX1	GROUP	MEMBER	REQ OUT	REQ IN	
BOETRX1	BOETRX1\$TRX2	TRX2	199	172	BOETRX1	BOETRX1\$TRX1	199	172	
TOTAL			199	172	TOTAL		199	172	
COFVLFNO	TRX2	TRX2	0	0	COFVLFNO	TRX1	0	0	
TOTAL			0	0	TOTAL		0	0	
CTTXNGRP	CTTXN##TRX2	TRX2	0	0	CTTXNGRP	CTTXN##TRX1	0	0	
TOTAL			0	0	TOTAL		0	0	
ESCM	M245	TRX2	0	0					
TOTAL			0	0					
EZBTCPCS	TRX2TCP/IP	TRX2	0	0	EZBTCPCS	TRX1TCP/IP	0	0	
TOTAL			0	0	TOTAL		0	0	
					IDAVQUI0	IDAVQUI0TRX1	0	0	
					TOTAL		0	0	
					IGWXSGIS	N200711006204502	0	0	
					TOTAL		0	0	
ISTCFS01	IPSVX\$\$\$DEIBMIPS	TRX2	0	0	ISTCFS01	IPSVX\$\$\$DEIBMIPS	0	0	
TOTAL			0	0	TOTAL		0	0	
ISTXCF	IPSVX\$\$\$DEIBMIPS	TRX2	140	113	ISTXCF	IPSVX\$\$\$DEIBMIPS	140	113	
TOTAL			140	113	TOTAL		140	113	
					IXCLO00F	M415	0	0	
					TOTAL		0	0	

Figure 233. XCF Activity Report - Usage by Member

Table 211. Fields in the XCF Activity Report - Usage by Member

Field Heading	Meaning
GROUP	The group name to which this member belongs.
MEMBER	The member name which was sent or received.
SYSTEM	The system name which this member resides on.
REQ FROM	The number of messages sent from the local system to the indicated member on the remote system.
REQ TO	The number of messages that the local system received from the indicated member on the remote system.
REQ OUT	The number of messages sent by the member on the local system.
REQ IN	The number of messages received by the member on the local system.

There is one line per member in the report.

```

XCF ACTIVITY
z/OS V2R2          SYSTEM ID TRX2      DATE 09/28/2016   INTERVAL 15.00.00000
                   RPT VERSION V2R2 RMF    TIME 09.15.00    CYCLE 1.000 SECONDS
PAGE 6

TOTAL SAMPLES = 900
XCF PATH STATISTICS
-----
OUTBOUND FROM TRX1
-----
T FROM/TO  Y DEVICE, OR  TRANSPORT  REQ  AVG Q
TO SYSTEM  P STRUCTURE   CLASS      OUT  LENGH  AVAIL  BUSY  RETRY
-----
TRX2  S IXCGRS      TCGRS       36  0.00   36    0    0
      S IXCPLEX_PATH1  BIG         95  0.00   95    0    0
      S IXCPLEX_PATH2  JES2        0  0.00    0    0    0
      S IXCPLEX_PATH3  TCRMF       13  0.00   13    0    0
      S IXCPLEX_PATH4  DEFAULT    839  0.00  839    0    0
      S IXCVLF        TCVLV        0  0.00    0    0    0
      C 1B62 TO 1C62  TCRMF        43  0.00   43    0    0
      C 1B63 TO 1C63  TCRMF        95  0.00   95    0    0
TOTAL                                1,121

INBOUND TO TRX1
-----
T FROM/TO  Y DEVICE, OR  REQ  BUFFERS  TRANSFER
FROM SYSTEM P STRUCTURE   IN  UNAVAIL  TIME
-----
TRX2  S IXCGRS      51    0  3.434
      S IXCPLEX_PATH1  158  0  2.791
      S IXCPLEX_PATH2   14  0  2.331
      S IXCPLEX_PATH3   14  0  0.836
      S IXCPLEX_PATH4  837  0  2.610
      S IXCVLF        15  0  0.921
      C 1C60 TO 1B60  280  0  1.152
      C 1C61 TO 1B61   77  0  0.932
TOTAL                                1,446
    
```

Figure 234. XCF Activity Report - Path Statistics (Coupling Facility and Channel-to-Channel)

Table 212. Fields in the XCF Activity Report - XCF Path Statistics

Field Heading	Meaning
TO SYSTEM	The name of the system to which the messages are sent.
TYP	Path type C Channel-to-channel S Coupling facility structure
FROM/TO DEVICE, OR STRUCTURE	<b>CTC:</b> The device number of the outbound signalling path on the local system and the device number of the inbound signalling path on the remote system that could not get message buffers (for example, 1B62 TO 1C62 in Figure 234).  <b>Coupling facility:</b> The coupling facility structure name (for example, IXCPLEX_PATH1 in Figure 234).  UNK* TO appears in this field when one or more paths are unknown at the beginning of the interval.
TRANSPORT CLASS	The name of the transport class XCF uses for the message transfer. There is one line in the report for each transport class for each target system.
REQ OUT	The number of attempts made to send a message over the indicated outbound signalling path.
AVG Q LENGTH	The average number of messages queued for data transfer over each outbound signalling path. The calculation is: $\text{AVG Q LENGTH} = \frac{\text{\# Message Entries for this Device}}{\text{\# Samples}}$
AVAIL	The number of times the signalling path was selected while available to immediately transfer a message.
BUSY	The number of times XCF selected a signalling path while a message was already in the process of being transferred.
RETRY	The number of times XCF initialized the signalling path.

Table 212. Fields in the XCF Activity Report - XCF Path Statistics (continued)

Field Heading	Meaning
FROM SYSTEM	The name of the system from which the messages are sent.
FROM/TO DEVICE, OR STRUCTURE	<p><b>CTC:</b> The device number of the remote outbound signalling path whose messages may not have been transmitted in a timely manner, and the device number of the inbound path for the system collecting the data to which the outbound device is connected.</p> <p><b>Coupling facility:</b> The coupling facility structure name.</p> <p>UNK* TO appears in this field when one or more paths are unknown at begin of the interval. The counter AVG Q LENGTH contains data, all other fields are blank.</p>
REQ IN	The number of requests received from the system on a path basis for each system.
BUFFERS UNAVAIL	The number of times that XCF was not able to get an inbound message buffer for the signalling path in anticipation of receiving a new message.
TRANSFER TIME	The average I/O transfer time in milliseconds for the most recently received signals. This metric is only available for inbound paths.

## Exception report

An Exception report presents a summary of the values that exceeded installation-defined thresholds over a specific period of time.

### How to request this report

Exception reports are generated from data contained in SMF records built during a Monitor I or a Monitor III data gatherer session. RMF compares the threshold values specified in the exception control statements with the computed value in the appropriate SMF record field. If the threshold is exceeded, RMF writes a line in the exception report.

You define the contents by one or more EXCEPT statements (see the *z/OS RMF User's Guide* for details), as follows:

```
EXCEPT(option)
```

### Contents of the report

Figure 235 on page 491 and Figure 236 on page 491 show sample exception reports.

#### Example:

If you want to produce an Exception report to display all intervals between midnight and 8 a.m. (off-shift) when the utilization of one processor was equal to or below 3%, you can use the following report option statements:

```
ETOD(0000,0800)
EXCEPT(CPU(CPUBSY,LE,3))
EXCEPT(CPU0(CPUBSY(0),LE,3))
EXCEPT(CPU1(CPUBSY(1),LE,3))
EXCEPT(CPU2(CPUBSY(2),LE,3))
EXCEPT(CPU3(CPUBSY(3),LE,3))
EXCEPT(CPU4(CPUBSY(4),LE,3))
EXCEPT(CPU5(CPUBSY(5),LE,3))
```

**Note:** The sample report assumes a 6-way processor. The first EXCEPT statement reflects the average utilization for all processors. The other EXCEPT statements reflect the average utilization for the specified processor.

R M F E X C E P T I O N R E P O R T										
z/OS V2R2				SYSTEM ID MVS1		START 09/28/2016-03.00.00		INTERVAL 00.30.00		PAGE 001
				RPT VERSION V2R2 RMF		END 09/28/2016-08.00.00		CYCLE 1.000 SECONDS		
CONTROL	DATE	TIME	INT	EXCEPTION	THRESHOLD	ACTUAL	EXCEPTION DESCRIPTION	NAME	VALUE	
STMT NAME	MM/DD	HH.MM.SS	HH.MM.SS			VALUE				
CPU1	09/28	03.00.00	00.30.00	LE 3		2.453	PERCENT CPU BUSY		-1	
CPU2	09/28	03.00.00	00.30.00	LE 3		2.315	PERCENT CPU BUSY		-2	
CPU3	09/28	03.00.00	00.30.00	LE 3		2.301	PERCENT CPU BUSY		-3	
CPU4	09/28	03.00.00	00.30.00	LE 3		2.391	PERCENT CPU BUSY		-4	
CPU5	09/28	03.00.00	00.30.00	LE 3		2.858	PERCENT CPU BUSY		-5	
CPU1	09/28	03.30.00	00.30.00	LE 3		2.637	PERCENT CPU BUSY		-1	
CPU2	09/28	03.30.00	00.30.00	LE 3		2.502	PERCENT CPU BUSY		-2	
CPU3	09/28	03.30.00	00.30.00	LE 3		2.165	PERCENT CPU BUSY		-3	
CPU4	09/28	03.30.00	00.30.00	LE 3		2.279	PERCENT CPU BUSY		-4	
CPU1	09/28	04.00.00	00.30.00	LE 3		2.690	PERCENT CPU BUSY		-1	
CPU2	09/28	04.00.00	00.30.00	LE 3		2.233	PERCENT CPU BUSY		-2	
CPU3	09/28	04.00.00	00.30.00	LE 3		2.350	PERCENT CPU BUSY		-3	
CPU4	09/28	04.30.00	00.30.00	LE 3		2.191	PERCENT CPU BUSY		-4	
CPU1	09/28	04.30.00	00.30.00	LE 3		2.830	PERCENT CPU BUSY		-1	
CPU2	09/28	04.30.00	00.30.00	LE 3		2.345	PERCENT CPU BUSY		-2	
CPU3	09/28	04.30.00	00.30.00	LE 3		2.373	PERCENT CPU BUSY		-3	
CPU4	09/28	04.30.00	00.30.00	LE 3		2.274	PERCENT CPU BUSY		-4	
CPU1	09/28	05.00.00	00.30.00	LE 3		2.783	PERCENT CPU BUSY		-1	
CPU2	09/28	05.00.00	00.30.00	LE 3		2.505	PERCENT CPU BUSY		-2	
CPU3	09/28	05.00.00	00.30.00	LE 3		2.257	PERCENT CPU BUSY		-3	
CPU4	09/28	05.00.00	00.30.00	LE 3		2.405	PERCENT CPU BUSY		-4	
CPU4	09/28	05.30.00	00.30.00	LE 3		2.548	PERCENT CPU BUSY		-4	
CPU3	09/28	06.30.00	00.30.00	LE 3		2.811	PERCENT CPU BUSY		-3	
CPU4	09/28	06.30.00	00.30.00	LE 3		2.861	PERCENT CPU BUSY		-4	
CPU4	09/28	07.00.00	00.30.00	LE 3		2.988	PERCENT CPU BUSY		-4	

Figure 235. Exception Report - Low CPU Utilization

**Example:**

If you want to produce an Exception report to display all intervals between 3 a.m. and 6 a.m., when the utilization of one processor was greater or equal 10%, you can use the following report option statements:

```
ETOD(0300,0600)
EXCEPT(CPU0(CPUBSY(0),GE,10))
```

R M F E X C E P T I O N R E P O R T										
z/OS V2R2				SYSTEM ID MVS1		START 09/28/2016-03.00.00		INTERVAL 00.30.00		PAGE 001
				RPT VERSION V2R2 RMF		END 09/28/2016-06.00.00		CYCLE 1.000 SECONDS		
CONTROL	DATE	TIME	INT	EXCEPTION	THRESHOLD	ACTUAL	EXCEPTION DESCRIPTION	NAME	VALUE	
STMT NAME	MM/DD	HH.MM.SS	HH.MM.SS			VALUE				
CPU0	09/28	03.00.00	00.30.00	GE 10		22.453	PERCENT CPU BUSY		-0	
CPU0	09/28	03.30.00	00.30.00	GE 10		22.637	PERCENT CPU BUSY		-0	
CPU0	09/28	04.00.00	00.30.00	GE 10		22.690	PERCENT CPU BUSY		-0	
CPU0	09/28	04.30.00	00.30.00	GE 10		22.830	PERCENT CPU BUSY		-0	
CPU0	09/28	05.00.00	00.30.00	GE 10		22.783	PERCENT CPU BUSY		-0	
CPU0	09/28	05.30.00	00.30.00	GE 10		21.383	PERCENT CPU BUSY		-0	

Figure 236. Exception Report - CPU Utilization

**Heading fields**

The heading fields for an Exception report identify the type of operating system, the release number and level of the operating system, the four-character SMF system ID of the system at system generation, and the RMF report level. The START field shows the date and time when the first interval in the reporting period began. The END field shows the date and time when the last interval ended. The date is in the form mm/dd/yy, and the time is in the form hh.mm.ss. The INTERVAL field shows the average length of the RMF measurement interval during the reporting period, in the form hh.mm.ss. The CYCLE field shows the length of the sampling cycle during the reporting period. When all SMF records

## PP - Exception report

have the same cycle length, that value is reported. When different cycle lengths are encountered, the Postprocessor sets the CYCLE field equal to the average of all cycle lengths encountered.

**Note:** When an Exception report consists of more than one page, the heading fields are repeated for each page. The START, END, CYCLE, and INTERVAL fields reflect the contents of the data in the entire report.

### Data fields

Two fields precede the data fields. NUMBER OF INTERVALS indicates the number of RMF measurement intervals included in the reporting period. TOTAL LENGTH OF INTERVALS indicates (in the form hh.mm.ss) the total of the reporting period.

When an Exception report consists of more than one page, the NUMBER OF INTERVALS field and the TOTAL LENGTH OF INTERVALS field reflect the contents of the entire report.

The meaning of each field in the report is described in the following text.

Table 213. Fields in the Exception Report

Field Heading	Meaning
CONTROL STMT NAME	The control statement name, if one was specified, that the Postprocessor uses to group control statements.
DATE MM/DD	The date, in the form of mm/dd, when the interval during which the exception occurred began.
TIME HH.MM.SS	The start time for the interval, during which the exception occurred, in the form of hh.mm.ss.
INT MM.SS	The actual length of the interval during which the exception occurred, in the form mm.ss.
EXCEPTION THRESHOLD	The threshold value and the relational operator specified in the EXCEPT statement.
ACTUAL VALUE	The actual value derived from the data contained in the SMF record. If RMF searches more than one resource for a value that exceeded the specified threshold level, the value is not printed. In addition, the field is not printed if all or part of a qualifier is omitted, or if a device qualified by class is specified. This also applies to Overview reports and records. <b>Note:</b> Because RMF processes the values differently, the actual values in this report may differ slightly from those in corresponding interval reports.
EXCEPTION DESCRIPTION	A short description of the exception. This is the condition name as specified in the EXCEPT control statement.
NAME VALUE	The condition-name qualifier as specified in the EXCEPT control statement.

## Improved exception reporting

The Exception report shows each exception that you have requested by an EXCEPT statement on one line. This can result in several lines being shown for each interval. Using the OVERVIEW(REPORT) statement in addition, you can create an Overview report that shows the data in a comprehensive format, similar to the Summary report. You get a listing of all exceptions for one interval on one line. This allows you to more easily analyze the performance of your system for a longer period.

### Example:

Assume that TSO001, TSO002, TSO004, TSO009, and TSO013 are your key volumes on the TSO system, and you are interested in getting all exceptions for the prime shift that you have defined either as a I/O activity rate of greater than 3 or as DASD response time greater than 25 milliseconds.

You specify the following control statements:

```
OVERVIEW(REPORT)
ETOD(0800,1800)
```

The exception-condition name DART specifies the device activity rate:

```
EXCEPT(TS0001IO(DART('TS0001'),GE,3))
EXCEPT(TS0002IO(DART('TS0002'),GE,3))
EXCEPT(TS0004IO(DART('TS0004'),GE,3))
EXCEPT(TS0009IO(DART('TS0009'),GE,3))
EXCEPT(TS0013IO(DART('TS0013'),GE,3))
```

The exception-condition name DRTAVG specifies the average response time:

```
EXCEPT(TS0001RT(DRTAVG('TS0001'),GE,25))
EXCEPT(TS0002RT(DRTAVG('TS0002'),GE,25))
EXCEPT(TS0004RT(DRTAVG('TS0004'),GE,25))
EXCEPT(TS0009RT(DRTAVG('TS0009'),GE,25))
EXCEPT(TS0013RT(DRTAVG('TS0013'),GE,25))
```

**Note:** You can get exceptions related to DASD only if you explicitly specify either device addresses or volume serial numbers. You will not get a meaningful Overview report if you just specify a generic class as DASD, because exception values will be reported only for single devices.

The *exception version* of the Overview report looks like this:

RMF OVERVIEW REPORT												
z/OS V2R2		SYSTEM ID SYS1		START 09/28/2016-08.02.26		INTERVAL 00.29.59		PAGE 001				
		RPT VERSION V2R2 RMF		END 09/28/2016-18.02.26		CYCLE 1.000 SECONDS						
NUMBER OF INTERVALS 20			TOTAL LENGTH OF INTERVALS 09.59.51									
DATE	TIME	INT	TS0001IO	TS0002IO	TS0004IO	TS0009IO	TS0013IO	TS0001RT	TS0002RT	TS0004RT	TS0009RT	TS0013RT
MM/DD	HH.MM.SS	HH.MM.SS										
09/28	08.02.26	00.29.59					4.309				28.329	
09/28	09.02.26	00.29.59		18.592					26.368			
09/28	11.32.26	00.30.00									27.843	
09/28	12.02.26	00.29.59				6.458	9.962			33.949		
09/28	12.32.26	00.30.00									26.709	
09/28	13.02.26	00.30.00										29.427
09/28	14.02.26	00.30.00							32.073			26.733
09/28	14.32.26	00.29.59	4.826							34.448		
09/28	15.02.26	00.29.59		3.134								
09/28	15.32.26	00.30.00								28.152		
09/28	16.02.26	00.30.00		13.403								
09/28	16.32.26	00.30.00										26.077

Figure 237. Overview Report - Exception Version

The reporting range covers 20 intervals, but you see in Figure 237 that only intervals with at least one exception value are listed in the report.

Table 214. Fields in the Overview Report

Field Heading	Meaning
NUMBER OF INTERVALS	The number of intervals that are being reported on.
TOTAL LENGTH OF INTERVALS	The sum of the intervals in the form HH.MM.SS, where HH is hours, MM is minutes, and SS is seconds.
DATE MM/DD	The date in the form MM/DD, where MM is the month and DD is the day.
TIME HH.MM.SS	The starting time of the interval.
INT MM.SS	The length of the interval.
FFFFFFF	The field name is the exception-condition name of the corresponding EXCEPT statement. All columns appear in the report in the same order as the EXCEPT statements are given as Postprocessor input.

### Overview report

You can use the Overview report for:

- Improved summary reporting - you can tailor the report according to your requirements
- Creating overview records - you can use the records as input for the RMF Spreadsheet Reporter or any other spreadsheet application

### How to request this report

The scope of data that can be reported in the Overview report depends on the data being gathered by Monitor I gatherer sessions.

You define the contents by one or more OVW statements (see the *z/OS RMF User's Guide* for details), as follows:

OVW(option)

**Note:** For records based on single-system reports, the option EXCEPT is still valid, but it is recommended to use OVW in general.

To produce the report, specify

OVERVIEW(REPORT)

To create data records either for spreadsheet processing or other applications, specify:

OVERVIEW(RECORD)

This report is also available in XML output format. Topic *How to work with Postprocessor XML reports* in the *z/OS RMF User's Guide* provides all required information on how to produce and view XML reports.

### Example URL for the DDS API

```
http://ddshost:8803/gpm/rmfpp.xml?overview=(DATA01(CADSTG(SSID(0600),  
DEVN(06F3))), (DB2PRD(CADRT(DEVN(0722),SSID(0700))),  
(RHT0050(CASRHT(SSID(0050)))))
```

### Report description

Internally, the Overview report and the Exception report use the same technology. Therefore, the scope of data that can be reported is the same in both reports. You find a list of all possible values in chapter *Postprocessor Exception and Overview Conditions* in the *z/OS RMF User's Guide*, based on the SMF records that are the source of the data. Furthermore, you might refer to the tables that are part of the description of all Postprocessor reports, for example Table 209 on page 484.

You cannot specify generic exception classes such as DASD without a qualifier. This is valid for exception reporting, but not for overview reporting. Here, you have to specify explicitly a qualifier, which for DASD could be either a device address or a volume serial number.

Due to the above described technology, you have to use the ETOD statement if you want to specify explicitly the time range for the Overview report.



## Improved summary reporting

The Summary report provides performance data that summarize system activity for each interval within the reporting period (see “Summary report” on page 498). The contents of the report cannot be modified.

The Overview report allows you to select the performance data you want to have shown according to your own requirements. With the suboptions NOSYSTEMS/SYSTEMS, you can select between sysplex reporting and reporting for each system that is known in the SMF records.

### Example:

You want to get an overview of the TSO activity in your sysplex for all intervals between 10am and 2pm. The following control statements assume that all TSO users run in service class TSOSERV and that you have defined three service class periods.

You specify the following control statements:

```
OVERVIEW(REPORT)
ETOD(1000,1400)
```

The exception-condition name T0TSRV specifies the total service units, the qualifier S.TSOSERV refers to service class TSOSERV, and suboption NOSYSTEMS defines sysplex reporting:

```
OVW(SERVUNIT(T0TSRV(S.TSOSERV)),NOSYSTEMS)
```

The exception-condition name RTIMETOT specifies the average response time:

```
OVW(RTIMEP1(RTIMETOT(S.TSOSERV.1)),NOSYSTEMS)
OVW(RTIMEP2(RTIMETOT(S.TSOSERV.2)),NOSYSTEMS)
OVW(RTIMEP3(RTIMETOT(S.TSOSERV.3)),NOSYSTEMS)
```

With the exception-condition name PI, you specify the performance index:

```
OVW(PIP1(PI(S.TSOSERV.1)),NOSYSTEMS)
OVW(PIP2(PI(S.TSOSERV.2)),NOSYSTEMS)
```

The exception-condition name TRANS specifies the transaction rate:

```
OVW(TRXP1(TRANS(S.TSOSERV.1)),NOSYSTEMS)
OVW(TRXP2(TRANS(S.TSOSERV.2)),NOSYSTEMS)
OVW(TRXP3(TRANS(S.TSOSERV.3)),NOSYSTEMS)
```

The *summary version* of the Overview report looks like this:

## PP - Overview report

RMF OVERVIEW REPORT											PAGE 001	
z/OS V2R2		SYSPLEX ID UTCPLXHD		START 09/28/2016-10.00.00		INTERVAL 00.10.00						
		RPT VERSION V2R2 RMF		END 09/28/2016-14.00.02								
NUMBER OF INTERVALS 24			TOTAL LENGTH OF INTERVALS 04.00.02									
DATE	TIME	INT	SERVUNIT	RTIMEP1	RTIMEP2	RTIMEP3	PIP1	PIP2	TRXP1	TRXP2	TRXP3	
MM/DD	HH.MM.SS	HH.MM.SS										
09/28	10.00.00	00.10.01	28406	0.229	14.491	18.916	0.763	2.415	12.40	0.82	0.10	
09/28	10.10.00	00.10.00	32696	0.264	6.071	20.415	0.880	1.012	15.52	1.00	0.09	
09/28	10.20.00	00.10.02	34245	0.261	4.425	17.877	0.870	0.738	15.66	1.02	0.11	
09/28	10.30.00	00.10.00	34943	0.270	5.695	19.504	0.900	0.949	15.03	1.02	0.12	
09/28	10.40.00	00.10.00	34830	0.313	12.344	20.806	1.043	2.057	15.29	1.02	0.11	
09/28	10.50.00	00.10.01	18360	0.373	6.619	23.352	1.243	1.103	9.34	0.54	0.06	
09/28	11.00.00	00.10.03	29893	3.928	25.872	57.726	13.093	4.312	12.47	0.90	0.10	
09/28	11.10.00	00.10.00	35164	0.251	8.943	15.304	0.837	1.491	15.48	1.00	0.13	
09/28	11.20.00	00.10.00	33544	0.236	4.563	17.671	0.787	0.761	15.50	1.03	0.11	
09/28	11.30.00	00.10.00	35057	0.257	5.829	13.389	0.857	0.972	15.73	1.01	0.11	
09/28	11.40.00	00.10.01	35811	0.252	10.049	15.257	0.840	1.675	15.40	1.07	0.13	
09/28	11.50.00	00.10.00	35419	0.248	4.507	17.461	0.827	0.751	15.24	1.03	0.13	
09/28	12.00.00	00.10.01	35902	0.437	6.883	18.944	1.457	1.147	15.12	1.02	0.13	
09/28	12.10.00	00.10.00	36967	0.247	9.635	16.407	0.823	1.606	16.10	1.06	0.12	
09/28	12.20.00	00.10.02	36024	0.260	4.552	18.229	0.867	0.759	15.15	1.02	0.13	
09/28	12.30.00	00.10.00	36296	0.263	5.072	20.555	0.877	0.845	15.29	1.06	0.12	
09/28	12.40.00	00.10.00	35129	0.262	10.237	16.135	0.873	1.706	15.37	1.02	0.12	
09/28	12.50.00	00.10.00	35355	0.274	8.098	15.106	0.913	1.350	15.31	1.06	0.12	
09/28	13.00.00	00.10.00	36936	0.213	3.833	10.036	0.710	0.639	16.00	1.09	0.12	
09/28	13.10.00	00.10.00	36919	0.182	5.205	9.323	0.607	0.868	16.17	1.08	0.13	
09/28	13.20.00	00.10.00	31098	0.208	5.420	9.599	0.693	0.903	14.14	0.87	0.12	
09/28	13.29.59	00.10.01	34909	1.184	13.976	25.991	3.947	2.329	14.20	1.00	0.13	
09/28	13.40.00	00.10.00	34807	0.234	3.865	11.235	0.780	0.644	15.75	1.03	0.11	
09/28	13.50.00	00.10.02	31773	0.227	8.372	12.955	0.757	1.395	14.00	0.93	0.11	

Figure 238. Overview Report - Summary Version

### Creating Overview records

You can also create records for further processing with the Spreadsheet Reporter or other applications either on the host system or on your workstation.

The Spreadsheet Reporter provides full support for converting SMF dump data, Postprocessor listings and Overview records into spreadsheets. You can use it to create and submit Postprocessor jobs directly on the workstation without a logon to the host system, and you will receive the data in the correct format back to the workstation. In addition, it provides sample spreadsheets to help you in presenting and analyzing performance data at a glance. You find a detailed description in the *z/OS RMF User's Guide*.

You get one record for each reported interval with the same information as in the printed *logical* line (this can be several *physical* lines on several pages if you define more than 11 exceptions) by specifying `OVERVIEW(RECORD)`.

A record can contain a maximum of 253 exceptions.

If you want to get both the report and the records, you can combine both control statements into `OVERVIEW(REPORT,RECORD)`.

It is recommended to use this version of the `OVERVIEW` statement, it provides the capability to check whether you really get the data that you expect.

For each report, the Postprocessor creates one Overview Header record and several (one for each interval) Overview Data records. If the input data for the Postprocessor consists of records for several systems, you get a set of records for each system.

The record mapping macro for all Overview records is `ERBOVREC`.

## Overview header record

Table 215. Overview Header Record - Prefix Section. This section is available only if you process the records in an MVS system.

Offsets	Format	Length	Name	Description
-4	binary	2	AMLEN	Record length (for records in the MVS system)
-2	binary	2	AMSGMT	

Table 216. Overview Header Record - Header Section. One per record.

Offsets	Format	Length	Name	Description							
0	0	binary	2	OVRLEN	Record length (for records on the workstation)						
2	2	binary	2	OVRSGMT							
4	4	EBCDIC	8	OVRTYPE	RMFOVREC (eye catcher)						
12	C	EBCDIC	1	OVRlvl	Record level change number						
13	D	binary	1	OVRFLG	Flags.  <table border="0"> <tr> <td><b>Bit</b></td> <td><b>Meaning When Set</b></td> </tr> <tr> <td>0</td> <td>Sysplex record</td> </tr> <tr> <td>1-7</td> <td>Reserved</td> </tr> </table>	<b>Bit</b>	<b>Meaning When Set</b>	0	Sysplex record	1-7	Reserved
<b>Bit</b>	<b>Meaning When Set</b>										
0	Sysplex record										
1-7	Reserved										
14	E	*	2	*	Reserved						
16	10	EBCDIC	3	OVRRMFV	RMF version number from SMF Record						
19	13	*	1	*	Reserved						
20	14	EBCDIC	8	OVRMVS	Version: ZVvvrmm (z/OS)						
28	1C	EBCDIC	4	OVRSID	System identification						
32	20	EBCDIC	6	OVRITIME	TOD monitor interval start: hhmms						
38	26	EBCDIC	2	OVRICENT	DATE monitor interval start: yy = high-order digits of century						
40	28	EBCDIC	5	OVRIDTYD	DATE monitor interval start: yyddd						
45	2D	*	1	*	Reserved						
46	2E	EBCDIC	4	OVRCYC	Sampling CYCLE length: tttt						
50	32	EBCDIC	8	OVRINT	Interval length: hh.mm.ss						
58	3A	EBCDIC	2	*	Reserved						
60	3C	binary	4	OVRCLGT	Total length for one SYSID: length of header record + (length of data records * number of data records)  This length does not include the prefix section of the records.						
64	40	binary	4	OVRCOLO	Offset to first report column names section						
68	44	binary	4	OVRCOLN	Number of report column names sections						
72	48	binary	4	OVRCOLL	Length of one report column names section						
78	4C	binary	4	OVRHDRN	Overview header record counter  First header record in data set has highest number, numbers will be in decreasing order.						
80	50	binary	4	OVRDATRN	Number of data records						
84	54	EBCDIC	8	OVRSPID	Sysplex Id (for sysplex records)						

Table 217. Overview Header Record - Report Column Names Section. One per interval.

Offsets	Format	Length	Name	Description	
0	0	EBCDIC	8	OVRNAME	Report column name
8	8	EBCDIC	8	OVRCOND	OVERVIEW condition name
16	10	EBCDIC	53	OVRQUAL	OVERVIEW condition qualifier
69	45	EBCDIC	1	*	Reserved

## PP - Overview report

Table 217. Overview Header Record - Report Column Names Section (continued). One per interval.

Offsets	Format	Length	Name	Description	
70	46	EBCDIC	2	OVROPER	Exception operator
72	48	EBCDIC	9	OVRTHV	Exception threshold value
81	51	EBCDIC	1	*	Reserved

### Overview data record

Table 218. Overview Data Record - Prefix Section. This section is available only if you process the records in an MVS system.

Offsets	Format	Length	Name	Description
-4	binary	2	AMDLEN	Record length (for records in the MVS system)
-2	binary	2	AMDSGMT	Zero

Table 219. Overview Data Record - Data Section. One per record.

Offsets	Format	Length	Name	Description	
0	0	binary	2	OVRDLEN	Record length (for records on the workstation)
2	2	binary	2	OVRDSGMT	Zero
4	4	binary	4	OVRDATO	Offset to first report data section
8	8	binary	4	OVRDATN	Number of report data sections
12	C	binary	4	OVRDATL	Length of one report data section
16	10	EBCDIC	5	OVRRIDAT	Reporting interval date: MM/DD
21	15	EBCDIC	1	*	Reserved
22	16	EBCDIC	8	OVRRITME	Reporting interval time hh.mm.ss
30	1E	EBCDIC	1	*	Reserved
31	1F	EBCDIC	8	OVRRIINT	Reporting interval length: hh.mm.ss
39	27	EBCDIC	1	*	Reserved

Table 220. Overview Data Record - Report Data Section. One per exception.

Offsets	Format	Length	Name	Description	
0	0	EBCDIC	1	*	Reserved
1	1	EBCDIC	9	OVRRVAL	Report actual data

## Summary report

Summary reports provide a high-level view of system activity. A summary report can consist of one line of data that summarizes system activity for each interval within the reporting period, a single line of data giving summary totals for all RMF intervals within the reporting period, or both interval summary data lines and a total summary data line. Figure 239 on page 499 shows a sample summary report, including both an interval summary line for each RMF measurement interval and a total summary line for all of the intervals.

### How to request this report

To produce this report, specify  
SUMMARY(INT | TOT)

## Generating a report

The data shown in a Summary report is derived from the SMF records created by any Monitor I sessions that ran during the reporting period. You specify the type of reporting required on Postprocessor control statements. For a description of Postprocessor control statements, see the *z/OS RMF User's Guide*.

## Special considerations of report output

The heading fields for a summary report are similar to the headings on an duration report. The START field shows when the first measurement interval began. The END field shows the date and time when the last interval ended. When a summary report consists of more than one page, the heading fields are repeated for each page. See "Single-system report header" on page 300 for more information on the heading fields.

The START, END, CYCLE, and INTERVAL fields reflect the contents of the page on which they appear. When total summary data is requested, a total summary line is generated for the intervals covered on each page, and the last page of the report shows values for START, END, CYCLE, and INTERVAL that reflect the contents of all pages in the report.

When a particular system activity is not measured during the reporting period, the columns describing that activity are omitted. For example, if a Monitor I session did not measure paging activity, the columns in the summary report that describe paging activity (SWAP RATE and DEMAND PAGING) are omitted.

The columns for JOB, TSO, STC, ASCH, and OMVS are available only if CPU activity was measured.

RMF SUMMARY REPORT																	PAGE 001
z/OS V2R2		SYSTEM ID MVS1		START 09/28/2016-08.02.26		INTERVAL 00.29.59											
		RPT VERSION V2R2 RMF		END 09/28/2016-12.02.26		CYCLE 1.000 SECONDS											
NUMBER OF INTERVALS 8				TOTAL LENGTH OF INTERVALS 03.59.56													
DATE	TIME	INT	CPU	DASD	DASD	JOB	JOB	TSO	TSO	STC	STC	ASCH	ASCH	OMVS	OMVS	SWAP	DEMAND
MM/DD	HH.MM.SS	MM.SS	BUSY	RESP	RATE	MAX	AVE	MAX	AVE	MAX	AVE	MAX	AVE	MAX	AVE	RATE	PAGING
09/28	08.02.26	29.59	22.6	12	465.9	4	2	64	51	76	73	0	0	0	0	0.00	0.00
09/28	08.32.26	30.00	28.9	17	642.2	7	4	89	77	74	71	0	0	0	0	0.00	0.00
09/28	09.02.26	29.59	26.0	19	566.7	8	5	103	97	73	71	0	0	0	0	0.00	0.00
09/28	09.32.26	30.00	27.8	12	583.5	7	4	103	100	75	70	0	0	0	0	0.00	0.00
09/28	10.02.26	29.59	25.5	20	542.6	8	6	107	102	70	68	0	0	0	0	0.00	0.00
09/28	10.32.26	29.59	43.7	11	545.3	7	4	106	99	73	70	0	0	0	0	0.00	0.00
09/28	11.02.26	30.00	48.6	13	470.1	10	6	98	94	72	69	0	0	0	0	0.00	0.00
09/28	11.32.26	30.00	27.4	13	218.1	7	4	88	85	72	70	0	0	0	0	0.00	0.00
TOTAL/AVERAGE			31.3	12	504.4	10	4	107	88	76	70	0	0	0	0	0.00	0.00

Figure 239. Summary Report

RMF omits a field (other than date, starting time, and interval time) if all values within the column are zero during the reporting period.

Table 221. Fields in the Summary Report

Field Heading	Meaning
NUMBER OF INTERVALS	The number of RMF measurement intervals included in the reporting period.
TOTAL LENGTH OF INTERVALS	The total length of the reporting period in the form hh.mm.ss.

When the report is more than one page, this field reflects the contents of the page on which it appears.

## PP - Summary report

Table 221. Fields in the Summary Report (continued)

Field Heading	Meaning	
DATE MM/DD	The date when each interval included in the summary report began.	This field is reported when you specify interval summary data. It does not appear in the total summary data line.
TIME HH.MM.SS	The start time for each interval included in the summary report.	
INT MM.SS	The actual length of each interval included in the summary report.	
CPU BUSY	The average busy percentage during the reporting period for all general purpose processors. Special purpose processors are not included in the calculation.  For systems running in a PR/SM environment, this value is the LPAR busy time percentage.	
DASD RESP	The average number of milliseconds required to complete an I/O request on all direct access storage devices included in the report.	
DASD RATE	The activity per second for all direct access storage devices included in the report. The value reported corresponds to an accumulation of each DEVICE ACTIVITY RATE field in the Direct Access Device Activity report.	
TAPE RATE	The activity per second for all magnetic tape devices included in the report. The value reported corresponds to an accumulation of each DEVICE ACTIVITY RATE field in the Magnetic Tape Device Activity report.	
JOB MAX/AVE	The maximum and average number of batch jobs that were active during each measurement interval. The values reported corresponds to the MAX/AVE number of BATCH address spaces in the CPU Activity report.	
TSO MAX/AVE	The maximum and average number of TSO/E sessions that were active during each measurement interval. The values reported corresponds to the MAX/AVE number of TSO/E address spaces in the CPU Activity report.	
STC MAX/AVE	The maximum and average number of started tasks and mount tasks that were active during each measurement interval. The value reported corresponds to the MAX/AVE number of STC address spaces in the CPU Activity report.	
ASCH MAX/AVE	The maximum/average number of APPC/MVS transaction scheduler (ASCH) address spaces that were active during each measurement interval. The value reported corresponds to the MAX/AVE number of ASCH address spaces in the CPU Activity report.	
OMVS MAX/AVE	The maximum/average number of OMVS address spaces that were active during each measurement interval. The value reported corresponds to the MAX/AVE number of OMVS address spaces in the CPU Activity report.	
SWAP RATE	The number of swaps per second for each interval. The value reported corresponds to the sum of the AUX STOR TOTAL and the EXP STOR TOTAL fields in the SWAP PLACEMENT ACTIVITY section of the Monitor I Paging Activity report.	
DEMAND PAGING	The number of demand paging requests per second for each interval. This is the demand paging rate from DASD (page fault rate).	
TOTAL/ AVERAGE	The single line that reports total summary data. The line contains either the average rate of events over the reporting period (or page), or the maximum number of events during any of the measurement intervals included in the reporting period (or page).	

## Spreadsheet reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see the *z/OS RMF User's Guide*.

---

## Appendix. Accessibility

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

If you experience difficulty with the accessibility of any z/OS information, send a detailed message to the "Contact us" web page for z/OS (<http://www.ibm.com/systems/z/os/zos/webqs.html>) or use the following mailing address.

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

---

### Accessibility features

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

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

---

### Consult assistive technologies

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

---

### Keyboard navigation of the user interface

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

- *z/OS TSO/E Primer*
- *z/OS TSO/E User's Guide*
- *z/OS V2R2 ISPF User's Guide Vol I*

---

### Dotted decimal syntax diagrams

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

Each line starts with a dotted decimal number; for example, 3 or 3.1 or 3.1.1. To hear these numbers correctly, make sure that the screen reader is set to read out



punctuation. All the syntax elements that have the same dotted decimal number (for example, all the syntax elements that have the number 3.1) are mutually exclusive alternatives. If you hear the lines 3.1 USERID and 3.1 SYSTEMID, your syntax can include either USERID or SYSTEMID, but not both.

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

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

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

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

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

**? indicates an optional syntax element**

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

**! indicates a default syntax element**

The exclamation mark (!) symbol indicates a default syntax element. A dotted decimal number followed by the ! symbol and a syntax element indicate that the syntax element is the default option for all syntax elements that share the same dotted decimal number. Only one of the syntax elements that share the dotted decimal number can specify the ! symbol. For example, if you hear the lines 2? FILE, 2.1! (KEEP), and 2.1 (DELETE), you know that (KEEP) is the



default option for the FILE keyword. In the example, if you include the FILE keyword, but do not specify an option, the default option KEEP is applied. A default option also applies to the next higher dotted decimal number. In this example, if the FILE keyword is omitted, the default FILE(KEEP) is used. However, if you hear the lines 2? FILE, 2.1, 2.1.1! (KEEP), and 2.1.1 (DELETE), the default option KEEP applies only to the next higher dotted decimal number, 2.1 (which does not have an associated keyword), and does not apply to 2? FILE. Nothing is used if the keyword FILE is omitted.

**\* indicates an optional syntax element that is repeatable**

The asterisk or glyph (\*) symbol indicates a syntax element that can be repeated zero or more times. A dotted decimal number followed by the \* symbol indicates that this syntax element can be used zero or more times; that is, it is optional and can be repeated. For example, if you hear the line 5.1\* data area, you know that you can include one data area, more than one data area, or no data area. If you hear the lines 3\* , 3 HOST, 3 STATE, you know that you can include HOST, STATE, both together, or nothing.

**Notes:**

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

**+ indicates a syntax element that must be included**

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



---

## Notices

This information was developed for products and services offered in the U.S.A. or elsewhere.

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

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

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A

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

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

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

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

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

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

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

Site Counsel  
IBM Corporation  
2455 South Road  
Poughkeepsie, NY 12601-5400  
USA

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

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

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

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

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

#### COPYRIGHT LICENSE:

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

---

## Policy for unsupported hardware

Various z/OS elements, such as DFSMS, HCD, JES2, JES3, and MVS™, contain code that supports specific hardware servers or devices. In some cases, this device-related element support remains in the product even after the hardware devices pass their announced End of Service date. z/OS may continue to service element code; however, it will not provide service related to unsupported hardware devices. Software problems related to these devices will not be accepted

for service, and current service activity will cease if a problem is determined to be associated with out-of-support devices. In such cases, fixes will not be issued.

---

## Minimum supported hardware

The minimum supported hardware for z/OS releases identified in z/OS announcements can subsequently change when service for particular servers or devices is withdrawn. Likewise, the levels of other software products supported on a particular release of z/OS are subject to the service support lifecycle of those products. Therefore, z/OS and its product publications (for example, panels, samples, messages, and product documentation) can include references to hardware and software that is no longer supported.

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

---

## Programming interface information

This book is intended to help the customer to use RMF reports, and contains a detailed description of all reports.

This book documents intended Programming Interfaces that allow the customer to write programs to obtain the services of RMF. This information is identified where it occurs, either by an introductory statement to a topic or section or by this marking:

————— **Programming Interface Information** —————

————— **End Programming Interface Information** —————

---

## Trademarks

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

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

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

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

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

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



---

## Glossary

This glossary contains chiefly definitions of terms used in this book, but some more general RMF and MVS terms are also defined.

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

### **APPC/MVS**

Advanced program-to-program communication

### **ASCH address space**

APPC transaction scheduler address space

### **AS** *Address space*

### **address space**

That part of MVS main storage that is allocated to a job.

### **auxiliary storage (AUX)**

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

### **background session**

In RMF, a monitor session that is started and controlled from the operator console. Contrast with *interactive session*

### **balanced systems**

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

### **basic mode**

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

### **bottleneck**

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

### **callable services**

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

### **captured storage**

See shared page group.

### **capture ratio**

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

### **central processor (CP)**

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

### **central processor complex (CPC)**

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

### **channel path**

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

### **CICS** Customer Information Control System

### **CIM provider**

A CIM provider is the link between the CIM server and the system interfaces. It allows the CIM server to access and manage the resources. Each CIM provider exposes the resources it represents in a standard way, using a small number of classes from the CIM schema or derived from the CIM schema. RMF monitoring providers are CIM providers implemented by RMF.

### **contention**

Two or more incompatible requests for the same resource. For example, contention occurs if a user requests a resource and specifies exclusive use, and another user requests the same resource, but specifies shared use.

### **coupling facility**

See *Cross-system Extended Services/Coupling Facility*.

### **CP** *Central processor*

### **criteria**

Performance criteria set in the WFEX report options. You can set criteria for all report classes (PROC, SYSTEM, TSO, and so on).

### **CPU speed**

Measurement of how much work your CPU can do in a certain amount of time.

### **cross-system coupling facility (XCF)**

A component of MVS that provides

functions to support cooperation between authorized programs running within a *sysplex*.

**Cross-system Extended Services/Coupling Facility (XES/CF)**

Provides services for MVS systems in a sysplex to share data on a coupling facility (CF).

**CS** Central storage

**Customer Information Control System (CICS)**

An IBM licensed program that enables transactions entered at remote terminals to be processed concurrently by user-written application programs. It includes facilities for building, using, and maintaining data bases.

**cycle** In RMF, the time at the end of which one sample is taken. Varies between 50 ms and 9999 ms. See also *sample*.

**data sample**

See *sample*

**DCM** See *Dynamic Channel Path Management*

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

**direct access storage device (DASD)**

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

**DLY** Delay

**DP** Dispatching priority

**dynamic channel path management**

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

**EMIF** ESCON multiple image facility

**enclave**

An enclave is a group of associated

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

**EPDM**

Enterprise Performance Data Manager/MVS

**execution velocity**

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

**exception reporting**

In RMF, the reporting of performance measurements that do not meet user-defined criteria. Shows potential performance problems explicitly, thus avoiding the need for constant monitoring.

**generalized trace facility (GTF)**

A service program that records significant system events, such as supervisor calls and start I/O operations, for the purpose of problem determination.

**GO mode**

In RMF, the Monitor III mode in which the screen is updated with the interval you specified in your session options. The terminal cannot be used for anything else when it is in GO mode. See also *mode*.

**graphic mode**

In RMF Monitor III, the mode which presents the performance data from the system in graphic format using the GDDM product. Contrast with *tabular mode*.

**GTF** generalized trace facility

**high-speed buffer (HSB)**

A cache or a set of logically partitioned blocks that provides significantly faster access to instructions and data than provided by central storage.

**HS** hiperspace

**HSB** High-speed buffer

**HSM** Hierarchical Storage Manager

**IBM System z Application Assist Processor (zAAP)**

A special purpose processor configured for running Java programming on selected zSeries machines.



**IBM System z Integrated Information Processor (zIIP)**

A special purpose processor designed to help free-up general computing capacity and lower overall total cost of computing for selected data and transaction processing workloads for business intelligence (BI), ERP and CRM, and selected network encryption workloads on the mainframe.

**IMS** Information Management System

**Information Management System (IMS)**

A database/data communication (DB/DC) system that can manage complex databases and networks. Synonymous with IMS/VS.

**interactive session**

In RMF, a monitor display-session that is controlled from the display terminal. Contrast with *background session*.

**JES** Job Entry Subsystem

**LCU** Logical control unit. Logical control units are also called 'Control Unit Headers' (CUH). For details about LCU/CUH please refer to the applicable *System z Input/Output Configuration Program User's Guide for ICP IOCP* (SB10-7037).

**logically partitioned (LPAR) mode**

A central processor mode that is available on the Configuration frame when using the PR/SM feature. It allows an operator to allocate processor unit hardware resources among logical partitions. Contrast with *basic mode*.

**logical partition (LP)**

A subset of the processor hardware that is defined to support an operating system. See also *logically partitioned (LPAR) mode*.

**LP** Logical partition

**LPAR** Logically partitioned (mode)

**LPAR cluster**

An LPAR cluster is the subset of the systems that are running as LPARs on the same CEC. Based on business goals, WLM can direct PR/SM to enable or disable CP capacity for an LPAR, without human intervention.

**migration rate**

The rate (pages/second) of pages being moved from expanded storage through central storage to auxiliary storage.

**mintime**

The smallest unit of sampling in Monitor III. Specifies a time interval during which the system is sampled. The data gatherer combines all samples gathered into a set of samples. The set of samples can be summarized and reported by the reporter.

**mode** Monitor III can run in various modes: GO mode (see *GO mode*) and STOP mode, which is the default mode. See also *graphic mode* and *tabular mode*.

**MPL** Multiprogramming level

**OMVS**

Reference to z/OS UNIX System Services

**partitioned data set (PDS)**

A data set in direct access storage that is divided into partitions, called members, each of which can contain a program, part of a program, or data.

**PDS** partitioned data set

**performance management**

The activity which monitors and allocates data processing resources to applications according to goals defined in a service level agreement or other objectives.

The discipline that encompasses collection of performance data and tuning of resources.

**PR/SM**

Processor Resource/Systems Manager

**Processor Resource/Systems Manager (PR/SM)**

The feature that allows the processor to run several operating systems environments simultaneously and provides logical partitioning capability. See also *LPAR*.

**range** The time interval you choose for your report.

**Resident time**

The time the address space was swapped in, in units of seconds.

**RMF monitoring provider**

see CIM provider

**sample**

Once in every cycle, the number of jobs waiting for a resource, and what job is using the resource at that moment, are

gathered for all resources of a system by Monitor III. These numbers constitute one sample.

**SCP** System control program

**seek** The DASD arm movement to a cylinder. A seek can range from the minimum to the maximum seek time of a device. In addition, some I/O operations involve multiple imbedded seeks where the total seek time can be more than the maximum device seek time.

**service class**

In Workload Manager, a subdivision of a *workload*. Performance goals and capacity boundaries are assigned to service classes.

**service level agreement (SLA)**

A written agreement of the information systems (I/S) service to be provided to the users of a computing installation.

**Service Level Reporter (SLR)**

An IBM licensed program that provides the user with a coordinated set of tools and techniques and consistent information to help manage the data processing installation. For example, SLR extracts information from SMF, IMS, and CICS logs, formats selected information into tabular or graphic reports, and gives assistance in maintaining database tables.

**service rate**

In the system resources manager, a measure of the rate at which system resources (services) are provided to individual jobs. It is used by the installation to specify performance objectives, and used by the workload manager to track the progress of individual jobs. Service is a linear combination of processing unit, I/O, and main storage measures that can be adjusted by the installation.

**shared page groups**

An address space can decide to share its storage with other address spaces using a function of RSM. As soon as other address spaces use these storage areas, they can no longer be tied to only one address space. These storage areas then reside as *shared page groups* in the system. The pages of shared page groups can reside in central, expanded, or auxiliary storage.

**SLA** service level agreement

**SLIP** serviceability level indication processing

**SLR** Service Level Reporter

**SMF** System management facility

**SMF buffer**

A wrap-around buffer area in storage, to which RMF data gatherers write performance data, and from which the Postprocessor extracts data for reports.

**speed** See *workflow*

**SRB** Service request block

**SRM** System resource manager

**SSCH** Start subchannel

**system control program (SCP)**

Programming that is fundamental to the operation of the system. SCPs include MVS, VM, and VSE operating systems and any other programming that is used to operate and maintain the system. Synonymous with *operating system*.

**sysplex**

A complex consisting of a number of coupled MVS systems.

**tabular mode**

In RMF, the mode in which Monitor III displays performance data in the form of lists. Contrast with *graphic mode*.

**TCB** Task control block

**threshold**

The exception criteria defined on the report options screen.

**throughput**

A measure of the amount of work performed by a computer system over a period of time, for example, number of jobs per day.

**TPNS** Teleprocessing network simulator

**TSO** Time Sharing Option, see *Time Sharing Option/Extensions*

**Time Sharing Option Extensions (TSO/E)**

In MVS, a time-sharing system accessed from a terminal that allows user access to MVS system services and interactive facilities.

**UIC** Unreferenced interval count

**uncaptured time**

CPU time not allocated to a specific address space.

**using** Jobs getting service from hardware resources (PROC or DEV) are *using* these resources.

**velocity**

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

**VTOC** Volume table of contents

**workflow**

The workflow of an address space represents how a job uses system resources and the speed at which the job moves through the system in relation to the maximum average speed at which the job could move through the system.

The workflow of resources indicates how efficiently users are being served.

**workload**

A logical group of work to be tracked, managed, and reported as a unit. Also, a logical group of service classes.

**WLM** Workload Manager

**XCF** Cross-system coupling facility

**XES/CF**

See *Cross-system Extended Services/Coupling Facility*.

**zAAP** see IBM System z Application Assist Processor.

**zIIP** see IBM System z Integrated Information Processor.



---

# Index

## Special characters

- % CU BUSY (control unit busy %) field
  - in Monitor III IOQUEUE report 114
- % Delayed for (percent delayed for) field
  - in Monitor III DELAY report 66
- % DP BUSY (director port busy %) field
  - in Monitor III IOQUEUE report 114
- % Frames (percentage frames) field
  - in Monitor III STORR report 168
- % of TRX (percentage of transactions) field
  - in Monitor III SYSRTD report 185

## A

- abnormally ended transactions field
  - in Monitor III SYSWKM report 203
- accessibility 501
  - contact IBM 501
  - features 501
- action field
  - in Monitor III WFEX report options action panel 217
- activated at field
  - in Monitor III SYSSUM report 191
- active frames field
  - in Monitor III STORF report 163
- active percentage field
  - in Monitor III DEVJ report (Active) 119
- active policy field
  - in Monitor III SYSSUM report 191
- activity rate field
  - in the Monitor III DEVN report (Rate) 74
- actual field
  - in Monitor III SYSWKM report 202
- address space identifier (ASID) field
  - in Monitor III STORC report 159
- Address Space Resource Data (ARD/ARDJ)
  - Monitor II report
    - command 246
    - contents of 246
    - field descriptions 247
    - how to request 246
    - purpose of 246
- Address Space SRM Data (ASRM/ASRMJ)
  - Monitor II report
    - command 253
    - contents of 253
    - how to request 253
    - purpose of 253
- Address Space State Data (ASD/ASDJ)
  - Monitor II report
    - command 250
    - contents of 250
    - how to request 250
    - purpose of 250

- address spaces serving this service class field
  - in Monitor III SYSWKM report 204
- Advanced Encryption Standard (AES) 361, 362, 364
- AES (Advanced Encryption Standard) 361, 362, 364
- alert field
  - in Monitor III WFEX report definition and criteria panel 219
- amount of common storage not released field
  - in Monitor III STORCR report 161
- amount used field
  - in Monitor III STORC report 159
- AP 481
- APPL% 475
- Application Assist Processor
  - See IBM System z Application Assist Processor (zAAP)
- ARD/ARDJ (Address Space Resource Data)
  - Monitor II report
    - command 246
    - contents of 246
    - field descriptions 247
    - how to request 246
    - purpose of 246
- ASD/ASDJ (Address Space State Data)
  - Monitor II report
    - command 250
    - contents of 250
    - how to request 250
    - purpose of 250
- ASID (address space identifier) field
  - in Monitor III STORC report 159
  - in Monitor III SYSWKM report 205
- ASM 416
- ASRM/ASRMJ (Address Space SRM Data)
  - Monitor II report
    - command 253
    - contents of 253
    - how to request 253
    - purpose of 253
- assistive technologies 501
- automatic customization 224
  - where to specify 224
  - workflow/exception (WFEX) options 211
- auxiliary slots field
  - in Monitor III STORF report 163
- Auxiliary Storage Manager 416
- available at end of range field
  - in Monitor III STORC report 158
- average active users field
  - in Monitor III STORR report 170
- average CPU utilization field
  - in Monitor III SYSINFO report (average CPU UTIL) 178

- average CPU utilization field (*continued*)
  - in Monitor III WFEX report (Average CPU Util) 212
- average CPU Utilization field
  - in the Monitor III SYSTREND report (CPU %) 197
- average CSA to SQA conversion field
  - in Monitor III STORC report 158
- average execution time (Avg. Exec. Time) field
  - in Monitor III SYSWKM report 202
- average number delayed for field
  - in Monitor III STORS report 172
  - in Monitor III SYSINFO report 180
- average response time (Avg. Resp. Time) field
  - in Monitor III SYSRTD report 185
  - in Monitor III SYSSUM report (active time) 195
  - in Monitor III SYSSUM report (queue time) 194
  - in Monitor III SYSSUM report (total time) 195
  - in Monitor III SYSWKM report 202
- average SRB time field
  - in the Monitor III SYSTREND report 197
- average TCB time field
  - in the Monitor III SYSTREND report 197
- average use summary field
  - in Monitor III STORC report 158
- average users delayed field
  - in Monitor III DEVJ report 120

## B

- blocked workload 344
  - threshold for promotion 349
- buffer counts by pool 145

## C

- C (class) field
  - in Monitor III STOR report 155
  - in Monitor III STORF report 162
  - in Monitor III WFEX report options action panel 217
- CACHDET (cache activity details)
  - how to request 35
  - Monitor III report
    - field descriptions 37
- CACHDET (Cache Detail)
  - Monitor III report
    - command abbreviation 21
    - purpose of 35
- CACHE (cache subsystem activity)
  - Postprocessor report
  - Subsystem Summary report 303

- CACHE (Cache Subsystem Activity)
  - Postprocessor report
    - command 302
    - device-level reporting 311
    - how to request 302
    - purpose of 301
    - report levels 302
    - spreadsheet and overview
      - reference 314
    - Subsystem Device Overview report 309
    - Subsystem Status report 304
    - subsystem summary report 303
    - subsystem-level reporting 304
    - top-20 device lists 303
- cache activity details
  - Monitor III report (CACHDET) 21, 35
    - field descriptions 37
- cache activity summary
  - Monitor III report (CACHSUM) 21, 39
    - field descriptions 41
- Cache Detail
  - Monitor III report
    - command abbreviation 21
    - purpose of 35
- Cache Detail Report
  - how to request 35
- Cache Subsystem Activity (CACHE)
  - Postprocessor report
    - command 302
    - device-level reporting 311
    - how to request 302
    - purpose of 301
    - report levels 302
    - spreadsheet and overview
      - reference 314
    - Subsystem Device Overview report 309
    - Subsystem Overview report 304
    - subsystem summary report 303
    - subsystem-level reporting 304
    - top-20 device lists 303
- Cache Summary (CACHSUM)
  - Monitor III report
    - command abbreviation 21
    - purpose of 39
- CACHSUM (cache activity summary)
  - Monitor III report
    - field descriptions 41
    - how to request 39
- CACHSUM (Cache Summary)
  - Monitor III report
    - command abbreviation 21
    - purpose of 39
- CAP 481
- capacity field
  - in Monitor III SYSSUM report (actual) 195
  - in Monitor III SYSSUM report (max) 195
  - in Monitor III SYSSUM report (min) 195
- Capacity report
  - contents and field descriptions 61
- capping delay (Capp) field
  - in Monitor III SYSWKM report 205
- central storage frames field
  - PAGING report 423
- central storage movement rates 422
- central storage request rates 422
- Central Storage/Processor/SRM (SRCS)
  - Monitor II report
    - command 292
    - contents of 293
    - field descriptions 293
    - how to request 292
    - purpose of 292
- CF (Coupling Facility Activity)
  - Postprocessor report
    - activity section 325
    - CF to CF activity section 332
    - command 316
    - contents of 317
    - how to request 316
    - purpose of 316
    - spreadsheet and overview
      - reference 333
    - subchannel activity section 329
    - usage summary section 318
- CF Systems report (Monitor III CFSYS report) 21, 52
- CFACT (coupling facility activity)
  - Monitor III report
    - field descriptions 46
    - report options panel 49
    - sample 44
- CFACT (Coupling Facility Activity)
  - Monitor III report
    - command 43
    - command abbreviation 21
    - contents of 43
    - how to request 43
    - purpose of 42
    - special considerations 43
- CFOVER (Coupling Facility Overview)
  - Monitor III report
    - command 50
    - command abbreviation 21
    - how to request 50
    - purpose of 50
- CFOVER Report
  - contents of 50
- CFSYS (coupling facility systems)
  - Monitor III report
    - field descriptions 53
- CFSYS (Coupling Facility Systems)
  - Monitor III report
    - command 52
    - command abbreviation 21
    - how to request 52
    - purpose of 52
- CHAN (Channel Path Activity)
  - Postprocessor report
    - command 335
    - contents of 336
    - dynamic channel path management (DCM) 334
    - how to request 335
    - purpose of 334
    - spreadsheet and overview
      - reference 339
- CHANNEL (Channel Path Activity)
  - Monitor II report
    - command 255
    - how to request 255
    - purpose of 255
  - Monitor III report
    - command 56
    - command abbreviation 21
    - customizing with Monitor III Utility 59
    - how to request 56
    - purpose of 56
- Channel Path
  - generation 58
- Channel Path Activity (CHAN)
  - Postprocessor report
    - command 335
    - contents of 336
    - dynamic channel path management (DCM) 334
    - how to request 335
    - purpose of 334
    - spreadsheet and overview
      - reference 339
- Channel Path Activity (CHANNEL)
  - Monitor II report
    - command 255
    - how to request 255
    - purpose of 255
  - Monitor III report
    - command abbreviation 21
    - customizing with Monitor III Utility 59
    - purpose of 56
- Channel Path Activity report (Monitor III CHANNEL report) 56
- channel paths field
  - in Monitor III IOQUEUE report 113
- CHPID taken field
  - in Monitor III IOQUEUE report 113
- class (C) field
  - in Monitor III STOR report 155
  - in Monitor III STORF report 162
  - in Monitor III WFEX report options action panel 217
- class field
  - in Monitor III GROUP report 103
  - in Monitor III SYSWKM report 202
- CML and LOCAL lock 438
- CMS/CMSEQDQ/CMSLatch/CMSSMF locks 438
- common storage
  - Monitor III report (STORC)
    - field descriptions - system information 158
- Common Storage (STORC)
  - Monitor III report
    - command 156, 157
    - how to request 157
    - purpose of 156
- common storage remaining
  - Monitor III report (STORCR)
    - field descriptions 161
- Common Storage Remaining (STORCR)
  - Monitor III report
    - command 160
    - how to request 160



- Common Storage Remaining (STORCR)
    - (continued)
    - Monitor III report (continued)
      - purpose of 160
  - Concepts, Monitor III
    - shortened intervals, reason for 16
  - connect time percentage field
    - in Monitor III DEV report (CON %) 72
    - in Monitor III DEVJ report 119
  - contact
    - z/OS 501
  - contention rate field
    - in Monitor III IOQUEUE report 113
  - control unit busy field
    - in Monitor III IOQUEUE report (% CU BUSY) 114
  - control unit field
    - in Monitor III IOQUEUE report 113
    - in Monitor III STORR report 168
  - Coupling facility
    - Postprocessor report
      - subchannel activity section, field description 329
  - Coupling Facility Activity (CF)
    - Postprocessor report
      - activity section 325
      - CF to CF activity section 332
      - command 316
      - contents of 317
      - how to request 316
      - purpose of 316
      - Spreadsheet and Overview
        - reference 333
      - subchannel activity section 329
      - usage summary section 318
  - Coupling Facility Activity (CFACT)
    - Monitor III report
      - command 43
      - command abbreviation 21
      - contents of 43
      - field descriptions 46
      - how to request 43
      - purpose of 42
      - report options panel 49
      - special considerations 43
  - Coupling Facility Overview (CFOVER)
    - Monitor III report
      - command abbreviation 21
      - purpose of 50
  - Coupling Facility Systems report
    - CFSYS report (Monitor III)
      - command abbreviation 21
      - purpose of 52
  - CPC (CPC Capacity)
    - Monitor III report
      - command 60
      - contents of 60
      - field descriptions 61
      - how to request 60
      - purpose of 60
  - CPC Capacity (CPC)
    - CPC report (Monitor III)
      - contents of 60
      - field descriptions 61
  - CPC Capacity report
    - CPC report (Monitor III)
      - purpose of 60
  - CPC Capacity report (Monitor III report) 21, 60
  - CPU (CPU Activity)
    - Postprocessor report
      - Blocked Workload Analysis 344
      - command 341
      - contents of 341, 343, 344
      - duration report 360
      - how to request 341
      - purpose of 340
      - Spreadsheet and Overview
        - reference 349, 355, 357, 359
      - System Address Space 343
      - using information in CPU Activity report 344
      - using information in Group Capacity report 357
      - using information in LPAR Cluster report 356
      - using information in Partition Data report 350
      - Work Unit Analysis 343
  - CPU Activity (CPU)
    - Postprocessor report
      - Blocked Workload Analysis 344
      - command 341
      - contents of 341, 343, 344
      - duration report 360
      - how to request 341
      - purpose of 340
      - Spreadsheet and Overview
        - reference 349, 355, 357, 359
      - System Address Space 343
      - using information in CPU Activity report 344
      - using information in Group Capacity report 357
      - using information in LPAR Cluster report 356
      - using information in Partition Data report 350
      - Work Unit Analysis 343
  - CPU utilization
    - LPAR view 243, 342
    - MVS view 243, 342
  - CPU utilization field
    - in Monitor III SYSINFO report (average CPU UTIL) 178
  - CPU Utilization field
    - in Monitor III WFEX report (Average CPU Util) 212
    - in the Monitor III SYSTREND report (CPU %) 197
  - criteria set
    - in Monitor III WFEX report definition and criteria panel 219
  - cross system coupling facility job delays (XCF)
    - Monitor III report
      - field descriptions 124
  - Cross-System Coupling Facility Activity (XCF)
    - Postprocessor report
      - command 486
  - Cross-System Coupling Facility Activity (XCF) (continued)
    - Postprocessor report (continued)
      - contents of 486
      - how to request 486
      - purpose of 486
  - Cross-System Coupling Facility Delays (XCF)
    - Monitor III report
      - command 226
      - command abbreviations 23
      - contents of 226
      - field descriptions 227
      - how to request 226
      - purpose of 226
      - report options 227
  - CRYPTO (Crypto Hardware Activity)
    - Postprocessor report
      - command 362
      - contents of 362
      - how to request 362
      - purpose of 361
      - spreadsheet and overview
        - reference 366
  - Crypto Hardware Activity (CRYPTO)
    - Postprocessor report
      - command 362
      - contents of 362
      - how to request 362
      - purpose of 361
      - spreadsheet and overview
        - reference 366
    - cryptographic accelerator 361
    - cryptographic coprocessor 361
    - cryptographic processors
      - adjunct processor 481
      - asynchronous message processor 481
  - CU (control unit) field
    - in the Monitor III DEVT report 81
- ## D
- Data Encryption Standard (DES) 361, 362, 364
  - data field
    - in Monitor III SYSRTD report 185
  - Data Index (DI)
    - Monitor III report
      - command 28
      - command abbreviations 21
      - condensed information 30
      - contents of 28
      - cursor-sensitive control 33
      - field descriptions 31
      - how to request 28
      - messages 32
      - options 33
      - preallocated data sets 29
      - purpose of 28
      - sample 29
  - data set delays
    - DSNJ (Monitor III report) 85
    - DSNV (Monitor III report) 87
  - Data Set Delays - Job (DSNJ)
    - Monitor III report
      - command 86
      - how to request 86

- Data Set Delays - Volume (DSNV)
  - Monitor III report
    - command 87
    - how to request 87
- Data Set Delays (DSND)
  - Monitor III report
    - command 83
    - command abbreviations 21
    - how to request 83
    - purpose of 82
- DCM
  - Monitor II
    - IOQUEUE report 271
  - Monitor III
    - IOQUEUE report 111
- DCM (dynamic channel path management)
  - in Monitor II Channel Path Activity report 255
  - in Monitor III Channel Path Activity report 56
  - in Postprocessor Channel Path Activity report 334
  - in Postprocessor I/O Queuing Activity report 403
- DDNAME
  - in Monitor III DI report 31
- DEL Q LNGTH (delayed queue length) field
  - in Monitor III IOQUEUE report 113
- DELAY
  - Monitor III report
    - command 63
    - command abbreviation 21
    - how to request 63
    - job selection panel 69
    - purpose of 63
    - sample 64
- Delay DB% (device busy delay percentage) field
  - in Monitor III DEVJ report 119
- delay monitoring 4
- delay percentage field
  - in Monitor III DELAY report (DLY %) 65
  - in Monitor III DEV report (DLY %) 71
  - in Monitor III DEVR report (DLY %) 78
  - in Monitor III ENQ report (DLY %) 96
  - in Monitor III ENQR report (Delayed) 100
  - in Monitor III HSM report (DLY %) 110
  - in Monitor III HSMJ report 121
  - in Monitor III JESJ report 121
  - in Monitor III STOR report (DLY %) 155
  - in Monitor III XCF report (DLY %) 227
  - in Monitor III XCFJ report 124
- DELAY Report
  - Monitor III Utility fields 67
- delayed name field
  - in Monitor III ENQR report 100
- DELAYJ (Job Delay)
  - Monitor III report
    - command abbreviations 21
- delta mode 244
- DES (Data Encryption Standard) 361, 362, 364
- description field
  - in Monitor III GROUP report 103
- DEV (device delays)
  - Monitor III report
    - options panel 72
    - sample 71
- DEV (Device Delays)
  - Monitor III report
    - command 70
    - command abbreviations 21
    - how to request 70
    - purpose of 70
- DEV/CU (device/control unit) field
  - in Monitor III DEVR report 78
- DEV/DEVV (Device Activity)
  - Monitor II report
    - command 259
    - contents of 261
    - field descriptions 262
    - how to request 259
    - purpose of 258
    - report options 265
    - special considerations 259
  - Postprocessor report
    - average cannot be calculated 370
    - byte-multiplexor-channel-attached device 368
    - command 368
    - connect/disconnect time overflow 371
    - contents of 368, 369, 370, 371
    - data inaccuracies in duration report 377
    - device changed or deleted 370
    - device data incomplete or missing 369
    - device in use 370
    - device not available during entire interval 369
    - Direct Access Device Activity report 369
    - hardware data not available 370
    - how to request 368
    - NMBR and SG suboptions 368
    - overflow condition occurred 370
    - overflow in non-Monitor II reports 370
    - pending time overflow 371
    - purpose of 367
    - required data not available 370
    - spreadsheet and overview reference 376
    - using information in the report 367
- Device Activity (DEV/DEVV) (continued)
  - Monitor II report (continued)
    - purpose of 258
    - report options 265
    - special considerations 259
  - Device Activity (DEVICE)
    - Postprocessor report
      - average cannot be calculated 370
      - byte-multiplexor-channel-attached device 368
      - command 368
      - contents of 368, 369, 370, 371, 376, 377
      - data inaccuracies in duration report 377
      - device changed or deleted 370
      - device data incomplete or missing 369
      - device in use 370
      - device not available during entire interval 369
      - Direct Access Device Activity report 369
      - hardware data not available 370
      - how to request 368
      - NMBR and SG suboptions 368
      - overflow condition occurred 370
      - overflow in non-Monitor II reports 370, 371
      - pending time overflow 371
      - purpose of 367
      - required data not available 370
      - spreadsheet and overview reference 376
      - using information in the report 367
  - device activity (DEVN)
    - Monitor III report
      - sample 73
  - Device Activity (DEVN)
    - Monitor III report
      - command 73
      - how to request 73
      - purpose of 72
  - Device Activity Trend (DEVT)
    - Monitor III report
      - command 80, 81
      - description 81
      - how to request 80, 81
      - purpose of 80, 81
  - device busy delay field
    - in Monitor III DEVJ report (Delay DB%) 119
  - device delays
    - Monitor III report (DEV)
      - options panel 72
      - sample 71
  - Device Delays - Job (DEVJ)
    - Monitor III report
      - command abbreviations 21
      - how to request 118
      - variation of JOB report 21, 118
  - Device Delays (DEV)
    - Monitor III report
      - command abbreviations 21
      - purpose of 70



- device field
    - in Monitor III DEVJ report 119
  - device number field
    - in the Monitor III DEVT report 81
  - Device Resource Delays (DEVR)
    - Monitor III report
      - command 75
      - command abbreviations 21
      - how to request 75
      - purpose of 75
  - device type field
    - in Monitor III STORR report 168
  - device using percentage field
    - in Monitor III PROCJ report 122
  - device/control unit (DEV/CU) field
    - in Monitor III DEVR report 78
  - DEVJ (Device Delays - Job)
    - Monitor III report
      - command abbreviations 21
      - how to request 118
      - variation of JOB report 21, 118
  - DEVJ (device delays job) 21, 118
  - DEVN (device activity)
    - Monitor III report
      - sample 73
  - DEVN (Device Activity)
    - Monitor III report
      - command 73
      - how to request 73
      - purpose of 72
  - DEVR (Device Resource Delays)
    - Monitor III report
      - command 75
      - command abbreviations 21
      - how to request 75
      - purpose of 75
  - DEVT (Device Activity Trend)
    - Monitor III report
      - command 80, 81
      - how to request 80, 81
      - purpose of 80
  - DI (Data Index)
    - Monitor III report
      - command 28
      - command abbreviations 21
      - condensed information 30
      - contents of 28
      - cursor-sensitive control 33
      - field descriptions 31
      - how to request 28
      - messages 32
      - options 33
      - preallocated data sets 29
      - purpose of 28
      - sample 29
  - direct memory access (DMA) 430
  - director port busy field
    - in Monitor III IOQUEUE report (% DP BUSY) 114
  - disconnect time field
    - in Monitor III DEVJ report 119
    - in Monitor III STORR report (DSC %) 169
  - Disk Space (SPACED)
    - Monitor III report
      - command 152
      - how to request 152
  - Disk Space (SPACED) (*continued*)
    - Monitor III report (*continued*)
      - purpose of 152
  - DLY % (delay percentage) field
    - in Monitor III DELAY report 65
    - in Monitor III DEV report 71
    - in Monitor III DEVR report 78
    - in Monitor III ENQ report 96
    - in Monitor III HSM report 110
    - in Monitor III STOR report 155
    - in Monitor III XCF report 227
  - DMA (direct memory access) 430
  - DMA (direct-memory-access) 133
  - DOMINO (Lotus Domino Server)
    - Postprocessor report
      - command 378
      - contents of 378
      - how to request 378
      - purpose of 377
  - DSC % (disconnect time percentage) field
    - in Monitor III STORR report 169
  - DSD (Detailed Storage Delays)
    - Monitor III report
      - purpose of 81
  - DSINDEX
    - Monitor III report
      - command abbreviations 21
      - purpose of 28
  - DSND (data set delays)
    - Monitor III report
      - field descriptions 84
      - report options panel 84
      - sample 84
  - DSND (Data Set Delays)
    - Monitor III report
      - command 83
      - command abbreviations 21
      - how to request 83
      - purpose of 82
  - DSNJ (data set delays - Job report) 21
  - DSNJ (Data Set Delays - Job) 85
    - Monitor III report
      - command 86
      - how to request 86
  - DSNV (data set delays - Volume report) 21
    - Monitor III report
      - command 87
      - field descriptions 88
      - purpose of 87
      - report options panel 89
      - sample 88
  - DSNV (Data Set Delays - Volume)
    - Monitor III report
      - command 87
      - how to request 87
  - duration field
    - in Monitor III SYSSUM report 195
  - duration reports 298
  - dynamic channel path management
    - Monitor II
      - IOQUEUE report 271
    - Monitor III
      - IOQUEUE report 111
  - dynamic channel path management (DCM)
    - in Monitor II Channel Path Activity report 255
    - in Monitor III Channel Path Activity report 56
    - in Postprocessor Channel Path Activity report 334
    - in Postprocessor I/O Queuing Activity report 403
- ## E
- earliest field
    - in the Monitor III DEVT report 81
    - in the Monitor III SYSTREND report 197
  - ECKD 388
  - elapsed time (ELAP time) field
    - in Monitor III STORC report 159
  - ENCLAVE (enclave activities)
    - Monitor III report
      - command 89, 90
      - field descriptions 90
      - how to request 90
      - sample 89
  - ENCLAVE (enclave activity) 21
  - enclave activities
    - Monitor III report
      - command 90
      - how to request 90
    - Monitor III report (ENCLAVE)
      - command 89
      - field descriptions 90
      - sample 89
  - Enclave Classification Attributes
    - pop-up 91
  - Enclave Details pop-up 91
  - enclave support
    - DELAY report (Monitor III) 68
    - ENCLAVE report (Monitor III) 89
    - PROC report (Monitor III) 139
    - SYSINFO report (Monitor III) 177
  - encryption 364
  - Enhanced DAT architecture 419
  - ENQ (Enqueue Activity)
    - Postprocessor report
      - command 382
      - contents of 383
      - how to request 382
      - purpose of 381
      - report levels 382
      - spreadsheet and overview reference 385
  - ENQ (enqueue delays) 21
    - Monitor III report
      - sample 95
  - ENQ (Enqueue Delays)
    - Monitor III report
      - command 95
      - how to request 95
      - purpose of 95
  - ENQJ (enqueue delays job) 21
  - ENQR (enqueue resource delays)
    - Monitor III report
      - field descriptions 100
      - sample 99

- ENQR (Enqueue Resource Delays)
  - Monitor III report
    - command 99
    - how to request 99
    - purpose of 99
- ENQR (resource-oriented enqueue delays) 21
- Enqueue Activity (ENQ)
  - Postprocessor report
    - command 382
    - contents of 383
    - how to request 382
    - purpose of 381
    - report levels 382
    - spreadsheet and overview reference 385
- enqueue delays
  - Monitor III report (ENQ)
    - sample 95
  - Monitor III report (SYSENG)
    - field descriptions 175
- Enqueue Delays (ENQ)
  - Monitor III report
    - command 95
    - how to request 95
    - purpose of 95
- enqueue resource delays (ENQR)
  - Monitor III report
    - field descriptions 100
    - sample 99
- Enqueue Resource Delays (ENQR)
  - Monitor III report
    - command 99
    - how to request 99
    - purpose of 99
- enterprise disk system 386
- Enterprise Disk Systems (ESS)
  - Postprocessor report
    - command 387
    - contents of 387
    - ESS statistics 386, 387
    - how to request 387
    - purpose of 386
- ESS (Enterprise Disk Systems)
  - Postprocessor report
    - command 387
    - contents of 387
    - ESS statistics 386, 387
    - how to request 387
    - purpose of 386
- ESS rank statistics
  - solid state drive 391
- EXCEPT (Exception)
  - Postprocessor report
    - command 490
    - contents of 490, 491, 492
    - data fields 492
    - heading fields 491

- Exception (EXCEPT) *(continued)*
  - Postprocessor report *(continued)*
    - how to request 490
    - improved exception reporting 492
    - purpose of 490
- Exceptions
  - in WFEX (Workflow/Exceptions)
    - Monitor III report 210
- EXCP rate field
  - in Monitor III GROUP report 104
- execution velocity (actual) field
  - in Monitor III SYSSUM report 192
- execution velocity (goal) field
  - in Monitor III SYSSUM report 191
- extended count key data 388

## F

- FCD ( FICON Director Activity)
  - Postprocessor report
    - command 392
    - how to request 392
    - purpose of 392
- FICON Director Activity (FCD)
  - Postprocessor report
    - command 392
    - how to request 392
    - purpose of 392
- fixed 1 MB memory objects 427
- frames online field
  - in Monitor III STORR report 168

## G

- global resource serialization (GRS) 435
- glossary 509
- goal field
  - in Monitor III SYSWKM report 202
- goal type field
  - in Monitor III SYSSUM report 195
- goals versus actual field
  - in Monitor III SYSSUM report 191
- GROUP (group response time) 21
- GROUP (Group Response Time)
  - Monitor III report
    - command 101
    - how to request 101
    - purpose of 101
- Group Capacity Report 357
- Group Response Time (GROUP)
  - Monitor III report
    - command 101
    - how to request 101
    - purpose of 101
- GRS (global resource serialization) 435
- GRS enqueue request details 438
- GRS enqueue requests summary 436
- GRS latches 436, 438

## H

- Hardware Group Report
  - using the information in 359
- help panels 19
- heterogeneous report class period 189
- heterogeneous report class periods 470

- HFS (HFS statistics)
  - Postprocessor report
    - field descriptions 395, 397
    - sample 395, 397
- HFS (Hierarchical File System Statistics)
  - Monitor II report
    - command 266
    - contents of 267
    - field descriptions 267
    - how to request 266
    - purpose of 266
    - report options 268
  - Postprocessor report
    - command 394
    - how to request 394
    - purpose of 394
- Hierarchical File System Statistics (HFS)
  - Monitor II report
    - command 266
    - contents of 267
    - field descriptions 267
    - how to request 266
    - purpose of 266
    - report options 268
  - Postprocessor report
    - command 394
    - how to request 394
    - purpose of 394
- hierarchical storage manager delays (HSM)
  - Monitor III report
    - report options panel 111
    - sample 110
- Hierarchical Storage Manager Delays (HSM)
  - Monitor III report
    - command 109
    - how to request 109
- High Performance FICON for System z 339
  - Monitor III 59, 258, 339
- HiperDispatch
  - parked time 346
- holding field
  - in Monitor III ENQJ report 120
  - in Monitor III PROCJ report 122
- holding name field
  - in Monitor III ENQR report 100
- holding name/SYS field
  - in Monitor III ENQ report 98
- holding percentage field
  - in Monitor III ENQ report 98
- holding status field
  - in Monitor III ENQ report 98
  - in Monitor III ENQR report 100
- homogeneous report class period 188
- homogeneous report class periods 470
- HSM (hierarchical storage manager delays) 22
  - Monitor III report
    - report options panel 111
    - sample 110
- HSM (Hierarchical Storage Manager Delays)
  - Monitor III report
    - command 109
    - how to request 109

HSM (Hierarchical Storage Manager Delays) (*continued*)  
 Monitor III report (*continued*)  
 purpose of 109

HSMJ (hierarchical storage manager delays job) 22

HTTP (HTTP Server)  
 Postprocessor report  
 command 398  
 how to request 398  
 purpose of 398

HTTP Server (HTTP)  
 Postprocessor report  
 command 398  
 how to request 398  
 purpose of 398

**I**

I/O Queuing Activity (IOQ)  
 Postprocessor report  
 command 404  
 contents of 406  
 data gathering  
 considerations 404  
 how to request 404  
 messages 405  
 missing data in report fields 404  
 purpose of 403  
 spreadsheet and overview  
 reference 413  
 using information in report 404

I/O Queuing Activity (IOQUEUE)  
 Monitor II report  
 command 272  
 contents of 273  
 evaluating details of cumulative mode output 271  
 field descriptions 274  
 how to request 272  
 purpose of 271  
 report options 275  
 special considerations 272

Monitor III report  
 command 111, 112  
 how to request 111, 112

IBM System z Application Assist Processor (zAAP)  
 definition 510  
 how referred to xvii  
 in CPC Capacity Report 62  
 in ENCLAVE Report 91  
 in PROC Report 137  
 in Processor Usage report 139  
 in SYSINFO Report 178

IBM System z Integrated Information Processor (zIIP)  
 definition 511  
 how referred to xvii  
 in CPC Capacity Report 62  
 in ENCLAVE Report 91  
 in PROC Report 137  
 in Processor Usage report 139  
 in SYSINFO Report 178

ID field  
 in Monitor III STORCR report 161

idling percentage (IDL %) field  
 in Monitor III DELAY report 65

IFL (Integrated Facility for Linux)  
 in CPC Capacity Report 62

ILOCK (IRLM Long Lock Detection)  
 Monitor II report  
 command 269  
 contents of 270  
 field descriptions 270  
 how to request 269  
 purpose of 269  
 special considerations 269

importance (I) field  
 in Monitor III SYSSUM report 191

indicator field  
 in Monitor III WFEX report options action panel 217

installed at field  
 in Monitor III SYSSUM report 191

Integrated Facility for Linux (IFL)  
 in CPC Capacity Report 62

Integrated Information Processor  
*See* IBM System z Integrated Information Processor (zIIP)

interval reports 298

IOQ (I/O Queuing Activity)  
 Postprocessor report  
 command 404  
 contents of 406  
 data gathering considerations 404  
 how to request 404  
 messages 405  
 missing data in report fields 404  
 purpose of 403  
 spreadsheet and overview  
 reference 413  
 using information in report 404

IOQUEUE (I/O queuing activity) 22

IOQUEUE (I/O Queuing Activity)  
 Monitor II report  
 command 272  
 contents of 273  
 evaluating details of cumulative mode output 271  
 field descriptions 274  
 how to request 272  
 purpose of 271  
 report options 275  
 special considerations 272

Monitor III report  
 command 111, 112  
 how to request 112  
 purpose of 111

IOS queue time field  
 in the Monitor III DEVN report (IosQ) 74

IPL definitions field  
 in Monitor III STORC report 158

IRLM Long Lock Detection (ILOCK)  
 Monitor II report  
 command 269  
 contents of 270  
 field descriptions 270  
 how to request 269  
 purpose of 269  
 special considerations 269

**J**

JES (Job Entry Subsystem Delays)  
 Monitor III report  
 command 114  
 how to request 114  
 purpose of 114

JES delays  
 Monitor III reporter session  
 command 22

JESJ (JES delays for specified job) 22

JOB (Job Delay)  
 Monitor III report  
 command 117  
 command abbreviations 22  
 Device Delay variation 118  
 how to request 117  
 purpose of 117

Job Delay (JOB)  
 Monitor III report  
 command 117  
 command abbreviations 22  
 Device Delay variation 118  
 how to request 117  
 purpose of 117

Job Delay report  
 Monitor III Utility fields 125

job ended date field  
 in Monitor III STORCR report 161

job ended time field  
 in Monitor III STORCR report 161

Job Entry Subsystem Delays (JES)  
 Monitor III report  
 command 114  
 how to request 114  
 purpose of 114

job field  
 in Monitor III ENQJ report 120  
 in Monitor III HSMJ report 121  
 in Monitor III JESJ report 121  
 in Monitor III PROCJ report 122

job report menu 25

job report options panel 125

Job Usage (USAGE)  
 Monitor III report  
 command 206  
 contents of 206  
 cursor-sensitive control on 209  
 field descriptions 207  
 how to request 206  
 utility fields 208

jobname field  
 in Monitor III DEV report 71  
 in Monitor III DEVJ report (JOB) 118  
 in Monitor III DEVR report 78  
 in Monitor III ENQ report 95  
 in Monitor III HSM report 110  
 in Monitor III MSGJ report 122  
 in Monitor III PROC report 137  
 in Monitor III STOR report 155  
 in Monitor III STORCR report 161  
 in Monitor III STORF report 162  
 in Monitor III SYSWKM report 204

## K

- keyboard
  - navigation 501
  - PF keys 501
  - shortcut keys 501

## L

- label field
  - in Monitor III WFEX report definition and criteria panel 219
  - in Monitor III WFEX report options action panel 217
- latest field
  - in the Monitor III DEVT report 81
  - in the Monitor III SYSTREND report 197
- LCU field
  - in Monitor III IOQUEUE report 113
- Library List (LLI)
  - Monitor II report
    - command 276
    - Contents of the report - APF Library List 277
    - Contents of the report - Link Library List 276
    - Contents of the report - LPA Library List 277
    - field descriptions 278
    - how to request 276
    - purpose of 276
  - LLI (Library List)
    - Monitor II report
      - command 276
      - contents of the report - APF Library List 277
      - contents of the report - Link Library List 276
      - contents of the report - LPA Library List 277
      - field descriptions 278
      - how to request 276
      - purpose of 276
    - Monitor II report
      - field descriptions 278
- local page data set slots 424
- LOCKSP (Spin Lock)
  - Monitor III report
    - command 126
    - how to request 126
    - purpose of 126
- LOCKSP (Spin locks) 22
- LOCKSU (Suspend Lock)
  - Monitor III report
    - command 128
    - contents of 128
    - how to request 128
    - purpose of 128
- LOCKSU (Suspend locks) 22
- long-term overview reporting
  - with the Postprocessor 297
- long-term performance analysis
  - RMF XP
    - long-term performance analysis 5
  - with RMF XP 5

## Lotus Domino Server (DOMINO)

- Postprocessor report
  - command 378
  - how to request 378
- Postprocessor reportS
  - contents of 378
  - purpose of 377
- LPAR cluster
  - description 356
  - Postprocessor report 356
- LPAR mode
  - CPU utilization 342

## M

- main delay path field
  - in Monitor III XCF report 227
- main delay volume(s) field
  - in Monitor III DEV report 72
- memory objects
  - fixed 1 MB 427
  - shared 1 MB 427
- MNTJ (operator mount job delays) 22
- model number field
  - in Monitor III SYSINFO report 178
- Monitor I
  - Real-time reporting with 295
- Monitor I session
  - system activity measured 3
- Monitor II background session
  - system activity measured 3
- Monitor II display session
  - system activity measured 3
- Monitor III
  - shortened intervals, reason for 16
  - sysplex support in different time zones 17
- Monitor III data gatherer session
  - delay information 4
  - system activity measured 3
- Monitor III menu
  - Overview Report Selection Menu 25
- Monitor III report commands
  - header for single-system reports 23
  - header for sysplex reports 23
  - overview 20
- Monitor III reporter session
  - commands 23
  - delay information 4
  - starting 18
  - system activity measured 3
- Monitor III Utility
  - customizing CHANNEL report with 59
- MSGJ (operator message job delays) 22
- MSU/h 60, 61
- Multiple Lock Structure 150
- MVS view
  - CPU utilization 243

## N

- name field
  - in Monitor III SYSINFO report 179
  - in Monitor III SYSSUM report 191

## name field (continued)

- in Monitor III WFEX report 212, 213, 226
- in Monitor III WFEX report definition and criteria panel 220
- navigation
  - keyboard 501
- No. of completed transactions (For nnnn TRX) field
  - in Monitor III SYSWKM report 202
- Notices 505
- number field
  - in Monitor III DEVJ report 119

## O

- OMVS
  - address spaces 129
  - OPD report 129
- OMVS (OMVS Kernel Activity)
  - Postprocessor report
    - command 414
    - contents of 414
    - field descriptions 415
    - how to request 414
    - purpose of 413
- OMVS Kernel Activity (OMVS)
  - Postprocessor report
    - command 414
    - contents of 414
    - field descriptions 415
    - how to request 414
    - purpose of 413
- OMVS Process Data (OPD)
  - Monitor III report
    - command 129, 130
    - contents of 130
    - how to request 130
    - purpose of 129
    - report options 132
- OPD (OMVS process data) 22
- OPD (OMVS Process Data)
  - Monitor III report
    - command 129, 130
    - contents of 130
    - how to request 130
    - purpose of 129
    - report options 132
- OPT (OPT Settings)
  - Monitor II report
    - command 278
    - contents of 279
    - field descriptions 279
    - how to request 278
    - purpose of 278
- OPT Settings (OPT)
  - Monitor II report
    - command 278
    - contents of 279
    - field descriptions 279
    - how to request 278
    - purpose of 278
- OVERVIEW
  - Postprocessor report
    - command 494
    - creating Overview records 496
    - description of 494



OVERVIEW (*continued*)  
 Postprocessor report (*continued*)  
   how to request 494  
   improved summary reporting 495  
   Overview data record 498  
   Overview header record 497  
   purpose of 494  
 overview condition names 301  
 overview conditions 301  
 overview names 301  
 Overview Report Selection Menu 25

## P

Page Data Set Activity (PAGESP)  
 Postprocessor report  
   command 416  
   contents of 417  
   how to request 416  
   overview names in 418  
   purpose of 416  
   using information in the report 416  
 Page Data Set Activity (PGSP)  
 Monitor II report  
   command 280  
   contents of 281  
   field descriptions 281  
   how to request 280  
   purpose of 280  
   report options 282  
   special considerations 280  
 PAGESP (Page Data Set Activity)  
 Postprocessor report  
   command 416  
   contents of 417  
   how to request 416  
   overview names in 418  
   purpose of 416  
   using information in the report 416  
 PAGING (Paging Activity)  
 Postprocessor report  
   Central Storage Movement Rates 422  
   central storage paging rates 419  
   Central Storage Request Rates 422  
   command 418  
   fixed 1 MB memory objects 427  
   frame counts 424  
   high virtual storage frames 427  
   how to request 418  
   local page data set slots 424  
   memory objects 427  
   purpose of 418  
   SCM paging blocks 424  
   shared 1 MB memory objects 427  
   shared frames 424  
   slot counts 424  
   spreadsheet and overview reference 422, 424, 426, 429  
 Paging Activity (PAGING)  
 Postprocessor report  
   Central Storage Movement Rates 422  
   central storage paging rates 419

Paging Activity (PAGING) (*continued*)  
 Postprocessor report (*continued*)  
   Central Storage Request Rates 422  
   command 418  
   fixed 1 MB memory objects 427  
   frame counts 424  
   high virtual storage frames 427  
   how to request 418  
   local page data set slots 424  
   memory objects 427  
   purpose of 418  
   SCM paging blocks 424  
   shared 1 MB memory objects 427  
   shared frames 424  
   slot counts 424  
   spreadsheet and overview reference 422, 424, 426, 429  
 Paging Activity (SPAG)  
 Monitor II report  
   command 290  
   contents of 291  
   field descriptions 291  
   how to request 290  
   purpose of 290  
 Paging Activity report  
   request rates and frames 423  
 parked time 346  
 Partition Data Report  
   z/VM guest 340  
 path attributes 113, 274  
 path number field  
   in Monitor III XCFJ report 124  
 PAV (parallel access volume)  
   in Monitor I DEVICE report 373  
 PCI functions 133, 430  
 PCIE (PCIE Activity Report)  
 Monitor III report  
   command 133  
   contents of 133  
   how to request 133  
   purpose of 133  
 PCIE (PCIE activity) 22  
 PCIE (PCIE Activity)  
 Postprocessor report  
   command 430  
   contents of 430  
   how to request 430  
   purpose of 430  
   spreadsheet and overview 432  
 PCIE Activity (PCIE)  
 Postprocessor report  
   command 430  
   contents of 430  
   how to request 430  
   purpose of 430  
   spreadsheet and overview 432  
 PCIE Activity Report  
   cursor sensitive control 135  
   PCIE Activity Report pop-up panels 135  
 PCIE Activity Report (PCIE)  
 Monitor III report  
   command 133  
   contents of 133  
   how to request 133  
   purpose of 133

peak allocation values field  
   in Monitor III STORC report 158  
 peer-to-peer remote copy 388  
 pend time percentage field  
   in Monitor III DEVJ report 119  
   in Monitor III DEVR report (PND % Reasons) 78  
   in Monitor III STORR report (PND %) 169  
   in the Monitor III DEVN report (PND % Reasons) 74  
 Pending (pend time percentage) field  
   in the Monitor III DEVN report 74  
 percent delayed for field  
   in Monitor III DELAY report (% Delayed for) 66  
   in Monitor III STOR report 156  
 percent used field  
   in Monitor III STORC report 159  
 percentage frames (% Frames) field  
   in Monitor III STORR report 168  
 performance index 193  
 period field  
   in Monitor III GROUP report 103  
   in Monitor III SYSWKM report 202  
 peripheral-component-interconnect (PCI) devices 133, 430  
 PGSP (Page Data Set Activity)  
 Monitor II report  
   command 280  
   contents of 281  
   field descriptions 281  
   how to request 280  
   purpose of 280  
   report options 282  
   special considerations 280  
 phase (P) field  
   in Monitor III SYSWKM report 203  
 PND % (pend time percentage) field  
   in Monitor III DEVR report 78  
   in Monitor III STORR report 169  
 possible cause or action field  
   in Monitor III WFEX report 215  
 Postprocessor  
   long-term overview reporting with 297  
 PPRC 388  
 preallocated datasets  
   in DSINDEX report 29  
 preferred path 113, 274  
 primary delay category field  
   in Monitor III HSMJ report 121  
   in Monitor III JESJ report 121  
 primary delay reason field  
   in Monitor III HSMJ report 121  
   in Monitor III JESJ report 121  
 primary reason field  
   in Monitor III DELAY report 67  
 primary response time component field  
   in Monitor III GROUP report 103  
 probable cause field  
   in Monitor III DEVJ report 118  
 PROC (processor delays) 22  
 PROC (Processor Delays)  
   Monitor III report  
   command 136  
   contents of 137

PROC (Processor Delays) (*continued*)  
 Monitor III report (*continued*)  
 field descriptions 137  
 how to request 136  
 purpose of 136  
 report options 139  
 utility fields 139

Processor Delays (PROC)  
 Monitor III report  
 command 136  
 contents of 137  
 field descriptions 137  
 how to request 136  
 purpose of 136  
 report options 139  
 utility fields 139

Processor Usage (PROCU)  
 Monitor III report  
 command 139  
 contents of 140  
 field descriptions 140  
 how to request 139  
 purpose of 139  
 report options 141  
 utility fields 141

processor using percentage field  
 in Monitor III PROCJ report 122  
 in Monitor III SYSWKM report 205

PROCJ (processor job delays) 22

PROCU (processor usage) 22

PROCU (Processor Usage)  
 Monitor III report  
 command 139  
 contents of 140  
 field descriptions 140  
 how to request 139  
 purpose of 139  
 report options 141  
 utility fields 141

Programming interface information 507

promotion threshold 349

## Q

QSCJ (quiescent jobs) 22

qualifier field  
 in Monitor III WFEX report options  
 action panel 217

quiesce delay (Quies) field  
 in Monitor III SYSWKM report 205

## R

Range field  
 in the Monitor III DEVT report 81  
 in the Monitor III SYSTREND  
 report 197

Real-time reporting  
 with Monitor I 295

reason field  
 in Monitor III WFEX report 213

Reasons field  
 in Monitor III DEVR report 78  
 in the Monitor III DEVN report 74

refresh field  
 in Monitor III SYSSUM report 191

report class period  
 heterogeneous 189  
 homogeneous 188

report class periods  
 heterogeneous 470  
 homogeneous 470

report mode  
 delta 244

report options (Monitor III) 17

Resource Group Data (RG)  
 Monitor III report  
 command 141  
 how to request 141  
 purpose of 141

resource group field  
 in Monitor III SYSSUM report 195

resource name field  
 in Monitor III ENQR report 99

Resource Report Selection Menu 26

resource waiting percentage field  
 in Monitor III ENQ report 96

resource waiting status field  
 in Monitor III ENQ report 96

response time  
 average 482  
 midpoint 482

response time (actual) field  
 in Monitor III SYSSUM report 192

response time (goal) field  
 in Monitor III SYSSUM report 192

response time breakdown field (delayed  
 by)  
 in Monitor III SYSWKM report 204

response time distribution (SYSRTD)  
 Monitor III report  
 option panel 185  
 sample 182

Response Time Distribution (SYSRTD)  
 Monitor III report  
 command 181  
 how to request 181  
 purpose of 181

response time field  
 in Monitor III SYSINFO report (RESP  
 Time) 179  
 in Monitor III SYSRTD report 185  
 in the Monitor III DEVN report  
 (RspT) 74

RG (Resource Group Data)  
 Monitor III report  
 command 141  
 how to request 141  
 purpose of 141

RLSDS (VSAM RLS activity by data  
 set) 22

RLSLRU (VSAM LRU Overview)  
 Monitor III report  
 command 143  
 how to request 143  
 purpose of 143

RLSLRU (VSAM RLS activity LRU  
 statistics) 22

RLSLRU report  
 buffer counts by pool 145

RLSSC (VSAM RLS activity by storage  
 class) 22

RLSSC/RLSDS (VSAM RLS Activity)  
 Monitor III report  
 command 145  
 how to request 145  
 purpose of 145  
 RLSDS command 145  
 RLSSC command 145

RMF (Resource Measurement Facility)  
 help panels 19  
 tutorial 19

RMF command 18

row and position field  
 in Monitor III WFEX report options  
 action panel 217

## S

SAID 388

sample to cycles ratio  
 changing 456

SCM  
 in Postprocessor PAGESP report 417  
 Storage Class Memory 417

SCM (SCM activity) 22

SCM (Storage Class Memory Activity  
 Report)  
 Monitor III report  
 command 151  
 contents of 151  
 how to request 151  
 purpose of 151

Postprocessor report  
 command 433  
 contents of 434  
 how to request 433  
 purpose of 433

SCM (storage class memory) 318

SCM paging blocks 424

SCSI 388

SDELAY (Serialization Delay)  
 XML Postprocessor report  
 command 435  
 contents of 436  
 how to request 435  
 purpose of 435  
 serialization delay details 438  
 serialization delay summary 436

SDEVICE (Shared Device Activity)  
 Postprocessor report  
 command 443  
 field descriptions 444  
 how to request 443  
 purpose of 442

Shared Direct Access Device  
 Activity report 443

Shared Magnetic Tape Device  
 Activity report 444

spreadsheet and overview  
 reference 449

using information in the  
 report 443

SDS (Sysplex Data Server)  
 Monitor II report  
 command 282  
 contents of 283  
 field descriptions 283  
 how to request 282

- SDS (Sysplex Data Server) *(continued)*
  - Monitor II report *(continued)*
    - purpose of 282
  - selecting a report 19
  - selection criteria
    - in the Monitor III DEVN report 74
  - sending comments to IBM xix
- SENQ (System Enqueue Contention)
  - Monitor II report
    - command 284
    - contents of 285
    - field descriptions 286
    - how to request 284
    - purpose of 284
    - report options 287
    - types of 284
- SENQR (System Enqueue Reserve)
  - Monitor II report
    - command 288
    - contents of 289
    - field descriptions 289
    - how to request 288
    - purpose of 288
    - report options 290
- Serialization Delay (SDELAY)
  - XML Postprocessor report
    - command 435
    - contents of 436
    - how to request 435
    - purpose of 435
    - serialization delay details 438
    - serialization delay summary 436
- server field
  - in Monitor III ENQJ report 120
- service class field
  - in Monitor III SYSWKM report 205
- service definition field
  - in Monitor III SYSSUM report 191
- service field
  - in Monitor III SYSWKM report 205
- Service Policy page 470
- SERVICE TIME 474
- session options xvii
- SHA-1 hash algorithm 365
- SHA-256 hash algorithm 365
- Shared Device Activity (SDEVICE)
  - Postprocessor report
    - command 443
    - field descriptions 444
    - how to request 443
    - purpose of 442
  - Shared Direct Access Device Activity report 443
  - Shared Magnetic Tape Device Activity report 444
  - spreadsheet and overview reference 449
  - using information in the report 443
- shared field
  - in Monitor III DEVJ report 119
- shared frames, PAGING report 424
- shortcut keys 501
- shortened intervals, reason for 16
- SHR field
  - in Monitor III CHANNEL report 58
- slot counts
  - available slots, PAGING report 426
  - bad slots, PAGING report 426
  - NON-VIO slots, PAGING report 426
  - total slots, PAGING report 426
  - VIO slots, PAGING report 426
- slots
  - PAGING report 424
- small computer system interface 388
- software pricing
  - Postprocessor CPU report 350
- solid state drive 391
- sort criteria
  - in the Monitor III DEVN report 74
- SPACED (disk space) 22
- SPACED (Disk Space)
  - Monitor III report
    - command 152
    - how to request 152
    - purpose of 152
- SPACEG (storage group space) 22
- SPACEG (Storage Space)
  - Monitor III report
    - command 153
    - how to request 153
    - purpose of 153
- SPAG (Paging Activity)
  - Monitor II report
    - command 290
    - contents of 291
    - field descriptions 291
    - how to request 290
    - purpose of 290
- Speed (Workflow)
  - in WFEX (Workflow/Exceptions) Monitor III report 210
- speed field
  - in Monitor III WFEX report 212
- speedometer
  - in Monitor III WFEX report 226
- Spin Lock (LOCKSP)
  - Monitor III report
    - command 126
    - how to request 126
    - purpose of 126
- Spin Lock Report
  - commands 126
- SQA + ESQA (system queue area)
  - overflow field
    - in Monitor III STORR report 168
- SRB field
  - in the Monitor III SYSTREND report 197
- SRCS (Central Storage/Processor/SRM)
  - Monitor II report
    - command 292
    - contents of 293
    - field descriptions 293
    - how to request 292
    - purpose of 292
- starting and ending an RMF reporter session 18
- status field
  - in Monitor III ENQJ report 120
- STOR (storage delays) 22
  - Monitor III report
    - command 153
    - report options panel 156
- STOR (Storage Delays)
  - Monitor III report
    - command 154
    - how to request 154
    - purpose of 154
- STOR Report
  - Monitor III Utility fields 156
- storage class memory (SCM) 318
- Storage Class Memory (SCM) 417
- Storage Class Memory (SCM) Activity Report (SCM)
  - Postprocessor report
    - command 433
    - how to request 433
- Storage Class Memory Activity Report (SCM)
  - Monitor III report
    - command 151
    - contents of 151
    - how to request 151
    - purpose of 151
  - Postprocessor report
    - contents of 434
    - purpose of 433
- storage delay summary
  - Monitor III report (STORS)
    - field descriptions 172
    - options panel 173
    - sample 171
- Storage Delay Summary (STORS)
  - Monitor III report
    - command 170
    - how to request 170
    - purpose of 170
- storage delays
  - Monitor III report (STOR)
    - report options panel 156
- Storage Delays (STOR)
  - Monitor III report
    - command 154
    - how to request 154
    - purpose of 154
- storage frames
  - Monitor III report (STORF)
    - report options panel 163
- Storage Frames (STORF)
  - Monitor III report
    - command 162
    - how to request 162
    - purpose of 162
- Storage Memory Objects (STORM)
  - Monitor III report
    - command 164
    - how to request 164
    - purpose of 164
- Storage Resource Delays (STORR)
  - Monitor III report
    - command 166
    - contents of 167
    - field descriptions 168
    - how to request 166
    - purpose of 166
    - utility fields 170
- Storage Space (SPACEG)
  - Monitor III report
    - command 153
    - how to request 153

Storage Space (SPACEG) (*continued*)  
 Monitor III report (*continued*)  
 purpose of 153  
 storage summary field  
 in Monitor III STORS report 171  
 STORC (common storage) 22  
 Monitor III report  
 field descriptions - system  
 information 158  
 STORC (Common Storage)  
 Monitor III report  
 command 156, 157  
 how to request 157  
 purpose of 156  
 STORC report  
 unallocated common area 159  
 STORC Report  
 Monitor III Utility fields 160  
 STORCR (common storage  
 remaining) 22  
 Monitor III report  
 field descriptions 161  
 STORCR (Common Storage Remaining)  
 Monitor III report  
 command 160  
 how to request 160  
 purpose of 160  
 STORF (storage frames) 22  
 Monitor III report  
 report options panel 163  
 STORF (Storage Frames)  
 Monitor III report  
 command 162  
 how to request 162  
 purpose of 162  
 STORJ (storage job delays) 22  
 STORM  
 field descriptions 164  
 STORM (storage memory objects) 164  
 STORM (Storage Memory Objects)  
 Monitor III report  
 command 164  
 how to request 164  
 purpose of 164  
 STORM (storage usage by memory  
 objects) 22  
 STORR (resource-oriented storage  
 delays) 22  
 STORR (Storage Resource Delays)  
 Monitor III report  
 command 166  
 contents of 167  
 field descriptions 168  
 how to request 166  
 purpose of 166  
 utility fields 170  
 STORS (storage delay summary)  
 Monitor III report  
 options panel 173  
 sample 171  
 STORS (Storage Delay Summary)  
 Monitor III report  
 command 170  
 how to request 170  
 purpose of 170  
 STORS (storage delays summary) 22  
 STORS Report  
 Monitor III Utility fields 172  
 subsystem data field  
 in Monitor III SYSRTD report 185  
 subsystem menu 27  
 subsystem reports 27  
 subsystem type field  
 in Monitor III SYSWKM report 203  
 SUMMARY  
 Postprocessor report  
 command 498  
 generating 499  
 how to request 498  
 purpose of 498  
 special considerations 499  
 using the Spreadsheet  
 Reporter 500  
 Summary of changes xxv  
 Suspend Lock (LOCKSU)  
 Monitor III report  
 command 128  
 contents of 128  
 how to request 128  
 purpose of 128  
 Suspend Lock Report  
 commands 128  
 switched time field  
 in Monitor III SYSWKM report 204  
 SYSENQ (sysplex enqueue delays) 22  
 Monitor III report  
 field descriptions 175  
 SYSENQ (Sysplex Enqueue Delays)  
 Monitor III report  
 command 174  
 how to request 174  
 purpose of 174  
 SYSINFO (system information) 23  
 Monitor III report  
 sample 178  
 SYSINFO (System Information)  
 Monitor III report  
 command 176  
 field descriptions 178  
 how to request 176  
 purpose of 176  
 utility fields 180  
 Sysplex Data Server (SDS)  
 Monitor II report  
 command 282  
 contents of 283  
 field descriptions 283  
 how to request 282  
 purpose of 282  
 Sysplex Enqueue Delays (SYSENQ)  
 Monitor III report  
 command 174  
 how to request 174  
 purpose of 174  
 Sysplex Report Selection Menu 24  
 sysplex summary  
 Monitor III report (SYSSUM)  
 sample 190  
 Sysplex Summary (SYSSUM)  
 Monitor III report  
 command 185, 186  
 how to request 186  
 purpose of 185  
 SYSRTD (response time distribution) 23  
 Monitor III report  
 option panel 185  
 sample 182  
 SYSRTD (Response Time Distribution)  
 Monitor III report  
 command 181  
 how to request 181  
 purpose of 181  
 SYSSUM (sysplex summary) 23  
 Monitor III report  
 sample 190  
 SYSSUM (Sysplex Summary)  
 Monitor III report  
 command 185, 186  
 how to request 186  
 purpose of 185  
 SYSSUM Report  
 field descriptions 191  
 field descriptions- continuation 194  
 Monitor III Utility fields 195  
 system activities measured  
 Monitor I session 3  
 Monitor II session 3  
 Monitor III session 3  
 Postprocessor 3  
 System Enqueue Contention (SENQ)  
 Monitor II report  
 command 284  
 contents of 285  
 field descriptions 286  
 how to request 284  
 purpose of 284  
 report options 287  
 types of 284  
 System Enqueue Reserve (SENQR)  
 Monitor II report  
 command 288  
 contents of 289  
 field descriptions 289  
 how to request 288  
 purpose of 288  
 report options 290  
 system field  
 in Monitor III ENQJ report 120  
 in Monitor III SYSWKM report 185  
 system information  
 Monitor III report (SYSINFO)  
 sample 178  
 Monitor III report (SYSTREND)  
 field descriptions 197  
 System Information (SYSINFO)  
 Monitor III report  
 command 176  
 field descriptions 178  
 how to request 176  
 purpose of 176  
 utility fields 180  
 system name field  
 in Monitor III SYSWKM report 205  
 system queue area (SQA + ESQA)  
 overflow field  
 in Monitor III STORR report 168  
 system suspend locks 436  
 System Trend (SYSTREND)  
 Monitor III report  
 command 196



System Trend (SYSTREND) (*continued*)  
 Monitor III report (*continued*)  
   how to request 196  
   purpose of 196  
 System z Application Assist Processor  
   *See* IBM System z Application Assist  
   Processor (zAAP)  
 System z Integrated Information  
 Processor  
   *See* IBM System z Integrated  
   Information Processor (zIIP)  
 SYSTREND (system trend)  
 Monitor III report  
   field descriptions 197  
 SYSTREND (System Trend)  
 Monitor III report  
   command 196  
   how to request 196  
   purpose of 196  
 SYSWKM (work manager delays) 23  
 Monitor III report  
   report options panel 205  
   sample 199  
 SYSWKM (Work Manager Delays)  
 Monitor III report  
   command 198  
   how to request 198  
   purpose of 198

## T

TCB field  
 in the Monitor III SYSTREND  
 report 197  
 text field  
 in Monitor III WFEX report definition  
 and criteria panel 219  
 time field  
 in Monitor III SYSSUM report 192  
 Time field  
 in the Monitor III DEVT report  
 (Time) 81  
 in the Monitor III SYSTREND report  
 (Time) 197  
 Time Stamp field  
 in the Monitor III DEVT report  
 (Time) 81  
 in the Monitor III SYSTREND report  
 (Time) 197  
 time zones  
 Monitor III sysplex support 17  
 total utilization field  
 in Monitor II channel path activity  
 (CHANNEL) report 59  
 TRACE - Trace Activity report  
 purpose of 450  
 TRACE (Trace Activity)  
 Postprocessor report  
   command 450  
   contents of 450  
   how to request 450  
   standard deviation output 454  
   using information in the  
   report 450  
 Trace Activity (TRACE)  
 Postprocessor report  
   command 450

Trace Activity (TRACE) (*continued*)  
 Postprocessor report (*continued*)  
   contents of 450  
   how to request 450  
   standard deviation output 454  
   using information in the  
   report 450  
 trademarks 507  
 TRANS/SEC (transactions per second)  
 field  
 in Monitor III SYSINFO report 179  
 Transaction ended rate field  
 in Monitor III GROUP report 104  
 transaction rate (Trx Rate) field  
 in Monitor III SYSRTD report 185  
 transaction response time field  
 in Monitor III SYSINFO report (RESP  
 Time) 179  
 transactions per second (TRANS/SEC)  
 field  
 in Monitor III SYSINFO report 179  
 tutorial 19  
 type (T) field  
 in Monitor III SYSSUM report 191  
 Type field  
 in the Monitor III DEVT report 81

## U

UKN % (unknown percentage) field  
 in Monitor III DELAY report 65  
 unallocated common area 159  
 unknown percentage (UKN %) field  
 in Monitor III DELAY report 65  
 USAGE (Job usage report) 23  
 USAGE (Job Usage)  
 Monitor III report  
   command 206  
   contents of 206  
   cursor-sensitive control on 209  
   field descriptions 207  
   how to request 206  
   utility fields 208  
 Usage report  
   purpose of 206  
 user interface  
   ISPF 501  
   TSO/E 501  
 user menu 27  
 user reports 27  
 users field  
 in Monitor III SYSINFO report 179  
 users/active field  
 in Monitor III WFEX report 212  
 using percentage field  
 in Monitor III DELAY report (USG  
 %) 65  
 in Monitor III DEV report (USG  
 %) 71  
 in Monitor III HSMJ report 121  
 in Monitor III JESJ report 121  
 in Monitor III PROC report (USG  
 %) 138

## V

vector utilization field  
 in Monitor III GROUP report (Vector  
 UTIL) 104  
 velocity field  
 in Monitor III SYSWKM report 205  
 Virtual Storage Activity (VSTOR)  
 High Virtual Memory Usage  
 section 463  
 Postprocessor report  
   allocated storage 454  
   changing sample to cycles  
   ratio 456  
   command 455  
   Common Storage Detail  
   section 458  
   Common Storage Summary  
   section 456  
   contents of 455  
   data gathering considerations 455  
   free storage 454  
   how to request 455  
   overview reference 458  
   partial private area data 461  
   Private Area Detail section 462  
   Private Area Summary  
   section 459  
   Private Storage Map 460, 461  
   purpose of 454  
   using information in 454  
 VMGUEST Monitor I gatherer  
 option 340  
 VolSer field  
 in the Monitor III DEVT report 81  
 volume serial field  
 in Monitor III STORR report 168  
 VSAM LRU Overview (RLSLRU)  
 Monitor III report  
   command 143  
   how to request 143  
   purpose of 143  
 VSAM RLS Activity (RLSSC/RLSDS)  
 Monitor III report  
   command 145  
   how to request 145  
   purpose of 145  
   RLSDS command 145  
   RLSSC command 145  
 VSTOR (Virtual Storage Activity)  
 Postprocessor report  
   allocated storage 454  
   changing sample to cycles  
   ratio 456  
   command 455  
   Common Storage Detail  
   section 458  
   Common Storage Summary  
   section 456  
   contents of 455  
   data gathering considerations 455  
   free storage 454  
   High Virtual Memory Usage  
   section 463  
   how to request 455  
   overview reference 458  
   partial private area data 461  
   Private Area Detail section 462

VSTOR (Virtual Storage Activity)  
(*continued*)  
  Postprocessor report (*continued*)  
    Private Area Summary  
      section 459  
  Private Storage Map 460, 461  
  purpose of 454  
  using information in 454

## W

WFEX (workflow/exceptions) 23  
  Monitor III report  
    field descriptions 226  
    graphic report 225  
    sample 212  
    summary 225  
WFEX (Workflow/Exceptions)  
  Monitor III report  
    command 210, 211  
    how to request 211  
    purpose of 210  
WFL % (workflow percentage) field  
  in Monitor III DELAY report 65  
  in Monitor III SYSINFO report 179  
  the Monitor III GROUP report 104  
WLMGL (Workload Activity)  
  Postprocessor report  
    command 463  
    contents of 464, 467, 468, 469, 470  
    field descriptions 471  
    how to request 463  
    option list 464  
    Policy Summary report 469  
    purpose of 463  
    Report Class Period report 470  
    Report Class report 469  
    Service Class Period report 464  
    Service Class report 467  
    spreadsheet and overview  
      reference 483  
    WLMGL option list 464  
    Workload Group and Service Class  
      Period report 468  
    Workload Group report 468  
work manager delays (SYSWKM)  
  Monitor III report  
    field descriptions 202  
    report options panel 205  
    sample 199  
Work Manager Delays (SYSWKM)  
  Monitor III report  
    command 198  
    how to request 198  
    purpose of 198  
workflow percentage field  
  in Monitor III DELAY report (WFL  
  %) 65  
  in Monitor III HSMJ report 121  
  in Monitor III JESJ report 121  
  in Monitor III SYSINFO report (WFL  
  %) 179  
  the Monitor III GROUP report (WFL  
  %) 104  
workflow/exceptions (WFEX)  
  Monitor III report  
    field descriptions 226

workflow/exceptions (WFEX) (*continued*)  
  Monitor III report (*continued*)  
    graphic report 225  
    sample 212  
    summary 225  
Workflow/Exceptions (WFEX)  
  Monitor III report  
    command 210, 211  
    how to request 211  
    purpose of 210  
Workload Activity (WLMGL)  
  Postprocessor report  
    command 463  
    contents of 464, 467, 468, 469, 470  
    field descriptions 471  
    how to request 463  
    option list 464  
    Policy Summary report 469  
    purpose of 463  
    Report Class Period report 470  
    Report Class report 469  
    Service Class Period report 464  
    Service Class report 467  
    spreadsheet and overview  
      reference 483  
    WLMGL option list 464  
    Workload Group and Service Class  
      Period report 468  
    Workload Group report 468

## X

XCF (Cross-System Coupling Facility  
Activity)  
  Postprocessor report  
    command 486  
    contents of 486  
    how to request 486  
    purpose of 486  
XCF (Cross-System Coupling Facility  
Delays)  
  Monitor III report  
    command 226  
    command abbreviations 23  
    contents of 226  
    field descriptions 227  
    how to request 226  
    purpose of 226  
    report options 227  
XCFJ (cross system coupling facility job  
delays) 23  
  Monitor III report  
    field descriptions 124

## Y

yellow and red (Yel Red) fields  
  in Monitor III WFEX report definition  
  and criteria panel 220

## Z

z/OS Unix  
  address spaces 129  
  OPD report 129

z/VM guest  
  Partition Data Report 340  
zAAP  
  See IBM System z Application Assist  
  Processor (zAAP)  
zFS File System (ZFSFS)  
  Monitor III report  
    command 228, 229  
    command abbreviation 23  
    contents of the report 229  
    field descriptions 230  
    how to request 228  
    purpose of 228  
zFS Kernel (ZFSKN)  
  Monitor III report  
    command 232  
    command abbreviation 23  
    contents of the report 233  
    field descriptions 233  
    how to request 232  
    purpose 232  
zFS Overview (ZFSOVW)  
  Monitor III report  
    command 234  
    contents of the report 234  
    field descriptions 234  
    how to request 234  
    navigating to details in 235  
    purpose of 23, 233  
ZFSFS (zFS File System)  
  Monitor III report  
    command 228, 229  
    command abbreviation 23  
    contents of the report 229  
    field descriptions 230  
    how to request 228  
    purpose of 228  
ZFSKN (zFS Kernel)  
  Monitor III report  
    command 232  
    command abbreviation 23  
    contents of the report 233  
    field descriptions 233  
    how to request 232  
    purpose 232  
ZFSOVW (zFS Overview)  
  Monitor III report  
    command 234  
    command abbreviation 23  
    contents of the report 234  
    field descriptions 234  
    how to request 234  
    navigating to details in 235  
    purpose of 233  
zHPF 339  
  Monitor III 59, 258, 339  
zIIP  
  See IBM System z Integrated  
  Information Processor (zIIP)





Product Number: 5650-ZOS

Printed in USA

SC34-2665-03

