z/OS

Resource Measurement Facility Report Analysis

SC33-7991-05

z/OS

Resource Measurement Facility Report Analysis

SC33-7991-05

Note
Before using this information and the product it supports, be sure to read the general information under "Notices" on page C-1.

Fifth Edition, October 2003

This edition applies to Version 1 Release 4 of z/OS (5694-A01), z/OS.e Version 1 Release 4 (5655-G52), and to all subsequent releases and modifications until otherwise indicated in new editions.

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

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

```
IBM DEUTSCHLAND ENTWICKLUNG GMBH
eSERVER PERFORMANCE MANAGEMENT DEVELOPMENT
SCHOENAICHER STRASSE 220
71032 BOEBLINGEN
GERMANY
```

If you prefer to send comments electronically, use one of the following methods:
FAX (RMF Development): Your International Access Code +49+7031+16+4240
Internet: rmf@de.ibm.com

Internet
Visit our homepage at http://www.ibm.com/servers/eserver/zseries/zos/rmf/

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

Make sure to include the following in your comment or note:
- Title and order number of this book
- Page number or topic related to your comment

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

© Copyright International Business Machines Corporation 1990, 2003. All rights reserved.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures vii

Tables xi

About this document xv
Who should use this document. xv
How this document is organized xv
The z/OS RMF library xvi
Related information xvi
Using LookAt to look up message explanations xvi

Summary of changes xvii
What's new in z/OS V1R4 xvii
What's new in SC33-7991-05 (z990 Exploitation) xvii
What's new in SC33-7991-03 and SC33-7991-04 (z990 Compatibility) xvii
What's new in SC33-7991-02 xviii
History of changes xviii
What's new in z/OS V1R2 xviii
What's new in z/OS V1R1 xix

Chapter 1. Introducing RMF data gathering and reporting 1-1
Gathering data 1-2
Short-term data collection with Monitor III . . 1-2
Snapshot monitoring with Monitor II 1-2
Long-term data gathering with Monitor I and Monitor III 1-2
Long-term overview reporting with the Postprocessor 1-3
Report analysis with the Spreadsheet Reporter . . 1-3
Monitoring on the workstation. 1-3
What you can gather and report 1-4
Activity monitoring 1-4
Delay monitoring 1-5
Reporting of other SMF data 1-5

Chapter 2. Interactive performance analysis with Monitor III 2-1
Using Monitor III reports 2-2
System activities measured 2-2
Where to start 2-2
Using cursor-sensitive control 2-5
Choosing graphic and tabular reports 2-6
Common Monitor III report measurements . . 2-6
Monitor III gatherer/reporter concepts. . . . 2-10
Monitor III report options 2-12
Monitor III sysplex support in different time zones 2-13
The reporter session 2-14
Starting and stopping a Monitor III reporter session 2-14
The Primary Menu 2-15
Selecting a report 2-15
Monitor III report commands 2-16
Header for single-system reports 2-20
Header for sysplex reports. 2-20
The Sysplex Report Selection Menu. 2-21
The Overview Report Selection Menu 2-22
The Job Report Selection Menu 2-23
The Resource Report Selection Menu 2-24
The Subsystem Report Selection Menu. 2-25
The User Report Selection Menu. 2-25
The Data Index 2-26
Purpose of the Data Index 2-26
Contents of the Data Index 2-26
CACHDET - Cache Detail report 2-33
Purpose of the report 2-33
Contents of the report 2-33
CACHSUM - Cache Summary report 2-38
Purpose of the report 2-38
Contents of the report 2-38
CFACT - Coupling Facility Activity report . . . 2-42
Purpose of the report 2-42
Special considerations 2-42
Contents of the report 2-43
CFOVER - Coupling Facility Overview report . . 2-49
Purpose of the report 2-49
Contents of the report 2-49
CFSYS - Coupling Facility Systems report. . . . 2-51
Purpose of the report 2-51
Contents of the report 2-51
CHANNEL - Channel Path Activity report . . . 2-55
Purpose of the report 2-55
Special considerations of report output . . . 2-55
Contents of the report 2-56
CPC - CPC Capacity report 2-59
Purpose of the report 2-59
Contents of the report 2-60
DELAY - Delay report 2-62
Purpose of the report 2-62
Contents of the report 2-62
DEV - Device Delays report 2-71
Purpose of the report 2-71
Contents of the report 2-71
DEVN - Device Activity report 2-74
Purpose of the report 2-74
Contents of the report 2-74
DEVR - Device Resource Delays report 2-77
Purpose of the report 2-77
Contents of the report 2-77
DEVT - Device Activity Trend report 2-82
Purpose of the report 2-82
Contents of the report 2-82
DSD - Detailed Storage Delays report 2-84
Purpose of the report 2-84
Contents of the report 2-84
DSND - Data Set Delays report 2-86
Purpose of the report 2-86
Special considerations 2-86

© Copyright IBM Corp. 1990, 2003

Contents of the report 2-87
DSNJ - Data Set Delays - Job report. 2-89
Purpose of the report 2-89
Contents of the report 2-89
DSNV - Data Set Delays - Volume report 2-91
Purpose of the report 2-91
Contents of the report 2-91
ENCLAVE - Enclave report 2-94
Purpose of the report 2-94
Contents of the report 2-94
ENQ - Enqueue Delays report 2-99
Purpose of the report 2-99
Contents of the report 2-99
ENQR - Enqueue Resource Delays report . . . 2-103
Purpose of the report 2-103
Contents of the report 2-103
GROUP - Group Response Time report 2-106
Purpose of the report 2-106
Contents of the report 2-106
HSM - Hierarchical Storage Manager Delays report 2-115
Purpose of the report 2-115
Contents of the report 2-115
IOQUEUE - I/O Queuing Activity report . . . 2-117
Purpose of the report 2-117
Special considerations of report output . . . 2-117
Contents of the report 2-118
JES - Job Entry Subsystem Delays report. . . . 2-121
Purpose of the report 2-121
Contents of the report 2-121
JOB - Job report 2-124
Purpose of the report 2-124
Contents of the report 2-124
Displaying the Job report 2-124
The top and bottom parts of the Job report 2-125
Device Delay variation 2-125
Enqueue Delay variation 2-128
HSM and JES variations 2-129
Operator Message and Mount Delay variations 2-130
Processor Delay variation. 2-130
Quiesce variation 2-131
Storage Delay variation 2-132
XCF variation 2-133
MSI - Migration SYSINFO report including execution velocity 2-136
Purpose of the report 2-136
Contents of the report 2-136
OPD - OMVS Process Data report 2-138
Purpose of the report 2-138
Contents of the report 2-139
PROC - Processor Delays report 2-143
Purpose of the report 2-143
Contents of the report 2-143
RG - Resource Group Data report 2-146
Purpose of the report 2-146
Contents of the report 2-146
RLSLRU - VSAM LRU Overview report 2-148
Purpose of the report 2-148
Contents of the report 2-148
RLSSC/RLSDS - VSAM RLS Activity report 2-150
Purpose of the report 2-150
Contents of the report 2-151
STOR - Storage Delays report 2-155
Purpose of the report 2-155
Contents of the report 2-155
STORC - Common Storage report 2-158
Purpose of the report 2-158
Contents of the report 2-159
STORCR - Common Storage Remaining report 2-163
Purpose of the report 2-163
Contents of the report 2-163
STORF - Storage Frames report. 2-165
Purpose of the report 2-165
Contents of the report 2-165
STORR - Storage Resource Delays report . . . 2-168
Purpose of the report 2-168
Contents of the report 2-168
STORS - Storage Delay Summary report. . . . 2-173
Purpose of the report 2-173
Contents of the report 2-173
SYSENQ - Sysplex Enqueue Delays report . . . 2-177
Purpose of the report 2-177
Contents of the report 2-177
SYSINFO - System Information report 2-179
Purpose of the report 2-179
Contents of the report 2-179
SYSRTD - Response Time Distribution report 2-184
Purpose of the report 2-184
Contents of the report 2-184
SYSSUM - Sysplex Summary report 2-189
Purpose of the report 2-189
Contents of the report 2-190
SYSTREND - System Trend report 2-201
Purpose of the report 2-201
Contents of the report 2-201
SYSWKM - Work Manager Delays report . . . 2-204
Purpose of the report 2-204
Contents of the report 2-205
WFEX - Workflow/Exceptions report 2-213
Purpose of the report 2-213
Using the information given in the report 2-213
Contents of the report 2-214
Workflow/Exceptions graphic report 2-216
XCF - Cross-System Coupling Facility Delays report 2-232
Purpose of the report 2-232
Contents of the report 2-232

Chapter 3. Snapshot reporting with Monitor II 3-1
Monitor II sessions. 3-2
Structure of Monitor II reports 3-3
Contents of the Monitor II report header . . . 3-3
Different formats of Monitor II reports 3-4
Different modes of Monitor II reports 3-5
ARD/ARDJ - Address Space Resource Data report 3-8
Purpose of the report 3-8
Using the information given in the report . . . 3-8
Contents of the report 3-8
ASD/ASDJ - Address Space State Data report . . 3-13
Purpose of the report 3-13

Using the information given in the report. . . 3-13
Contents of the report 3-13
Report options. 3-16
ASRM/ASRMJ - Address Space SRM Data report 3-17
Purpose of the report 3-17
Using the information given in the report. . . 3-17
Contents of the report 3-17
Report options. 3-19
CHANNEL - Channel Path Activity report . . . 3-20
Purpose of the report 3-20
Using the information given in the report. . . 3-20
Special considerations of report output . . . 3-20
Contents of the report 3-21
DDMN - Domain Activity report 3-24
Purpose of the report 3-24
Using the information given in the report. . . 3-24
Contents of the report 3-24
DEV/DEVV - Device Activity report 3-26
Purpose of the report 3-26
Using the information given in the report. . . 3-26
Special considerations of report output . . . 3-27
Contents of the report 3-28
HFS - Hierarchical File System Statistics report 3-34
Purpose of the report 3-34
Contents of the report 3-34
ILOCK - IRLM Long Lock Detection report . . . 3-37
Purpose of the report 3-37
Special considerations 3-37
Contents of the report 3-38
IOQUEUE - I/O Queuing Activity report. . . . 3-40
Purpose of the report 3-40
Using the information given in the report. . . 3-40
Special considerations of report output . . . 3-41
Contents of the report 3-42
LLI - Library List report 3-45
Purpose of the report 3-45
Contents of the report - Link Library List. . . 3-45
Contents of the report - LPA Library List . . . 3-46
Contents of the report - APF Library List . . . 3-46
PGSP - Page/Swap Data Set Activity report . . . 3-48
Purpose of the report 3-48
Using the information given in the report. . . 3-48
Special considerations of report output . . . 3-48
Contents of the report 3-49
SDS - Sysplex Data Server report 3-51
Purpose of the report 3-51
Contents of the report 3-51
SENQ - System Enqueue Contention report . . . 3-53
Purpose of the report 3-53
Using the information given in the report. . . 3-53
Different types of SENQ reports 3-53
Contents of the report 3-54
SENQR - System Enqueue Reserve report. . . . 3-58
Purpose of the report 3-58
Using the information given in the report. . . 3-58
Contents of the report 3-58
SPAG - Paging Activity report 3-60
Purpose of the report 3-60
Using the information given in the report. . . 3-60
Contents of the report 3-60
SRCS - Central Storage/Processor/SRM report 3-62
Using the information given in the report. . . 3-62
Contents of the report 3-62
TRX - Transaction Activity report 3-64
Purpose of the report 3-64
Using the information given in the report. . . 3-64
Special considerations of report output . . . 3-64
Contents of the report 3-65

Chapter 4. Real-time reporting with Monitor I 4-1

Chapter 5. Long-term overview reporting with the Postprocessor . . . 5-1
Interval and duration reports 5-2
Interval reports 5-2
Duration reports 5-2
Single-system report header. 5-4
Sysplex report header. 5-5
Duration report header 5-5
Overview condition names 5-5
CACHE - Cache Subsystem Activity report . . . 5-6
Purpose of the report 5-6
Using the information given in the report . . . 5-6
Different report levels. 5-6
Contents of the report 5-7
Device-level reporting 5-17
Spreadsheet and Overview reference 5-19
CF - Coupling Facility Activity report 5-22
Contents of the report 5-22
Spreadsheet and Overview reference 5-34
CHAN - Channel Path Activity report 5-35
Purpose of the report 5-35
Using the information given in the report. . . 5-36
Contents of the report 5-37
Spreadsheet and Overview reference 5-40
CPU - CPU Activity report 5-41
Purpose of the report 5-41
Contents of the report 5-41
Using the information in the CPU report . . . 5-43
Spreadsheet and Overview reference 5-45
Using the information in the Partition Data report. 5-46
Spreadsheet and Overview reference 5-50
Using the information in the LPAR Cluster report. 5-50
Spreadsheet and Overview reference 5-52
Duration report 5-53
CRYPTO - CRYPTO Hardware Activity report 5-54
Purpose of the report 5-54
Contents of the report 5-55
Spreadsheet and Overview reference 5-57
DEVICE - Device Activity report. 5-59
Purpose of the report 5-59
Using the information given in the report. . . 5-59
Contents of the report 5-59
Spreadsheet and Overview reference 5-68
DOMINO - Lotus Domino Server report 5-71
Purpose of the report 5-71
Contents of the report 5-71
ENQ - Enqueue Activity report 5-75

Purpose of the report 5-75
Using the information given in the report. . . 5-75
Different report levels 5-75
Contents of the report 5-76
Messages 5-79
FCD - FICON Director Activity report 5-80
Purpose of the report 5-80
Contents of the report 5-80
HFS - Hierarchical File System Statistics report 5-82
Purpose of the report 5-82
Using the information given in the report. . . 5-82
Contents of the report 5-83
Special considerations 5-86
HTTP - HTTP Server report 5-87
Purpose of the report 5-87
Contents of the report 5-87
IOQ - I/O Queuing Activity report 5-92
Purpose of the report 5-92
Using the information given in the report. . . 5-92
Data gathering considerations 5-93
Missing data in report fields 5-93
Messages 5-93
Contents of the report 5-94
Spreadsheet and Overview reference 5-99
OMVS - OMVS Kernel Activity report 5-101
Contents of the report 5-101
PAGESP - Page/Swap Data Set Activity report 5-104
Purpose of the report 5-104
Using the information given in the report 5-104
Contents of the report 5-105
Overview reference 5-106
PAGING - Paging Activity report 5-107
Purpose of the report 5-107
Using the information given in the report 5-107
Contents of the report 5-107
SDEVICE - Shared Device Activity report . . . 5-125
Purpose of the report 5-125
Using the information given in the report 5-125
Shared Direct Access Device Activity report 5-125
Shared Magnetic Tape Device Activity report 5-126
Spreadsheet and Overview reference 5-133
TRACE - Trace Activity report 5-134
Purpose of the report 5-134
Using the information given in the report 5-134
Contents of the report 5-134
VSTOR - Virtual Storage Activity report 5-140
Purpose of the report 5-140
Using the information given in the report 5-140
Contents of the report 5-141
Data gathering considerations 5-141
Overview reference 5-144
Common Storage Detail section. 5-144
Private Area Summary section 5-145
Private Area Detail section 5-147
WKLD - Workload Activity report - Compatibility Mode 5-149
Purpose of the report 5-149
Generating a report. 5-149
Using the information given in the report 5-150
Contents of the report 5-150
WLMGL - Workload Activity report - Goal Mode 5-153
The WLMGL option list 5-153
Contents of the report 5-153
Service Class Period report 5-155
Service Class report. 5-156
Workload Group report 5-156
Policy Summary report 5-156
Workload Group and Service Class Period report 5-157
Report Class report 5-157
Report Class Period report 5-157
Service Policy page 5-158
Field descriptions of all reports 5-159
Spreadsheet and Overview reference 5-171
XCF - Cross-System Coupling Facility Activity report 5-174
Contents of the report 5-174
Spreadsheet reference 5-178
Exception report 5-179
Contents of the report 5-179
Generating a report. 5-179
Report description 5-179
Improved exception reporting 5-182
Overview report 5-185
Report description 5-185
Plot report 5-191
Contents of the report 5-191
Generating a report. 5-191
Report description 5-191
Types of plot reports 5-192
Summary report 5-202
Contents of the report 5-202
Generating a report. 5-202
Special considerations of report output . . . 5-202
Spreadsheet reference 5-204

Appendix. Accessibility A-1
Using assistive technologies A-1
Keyboard navigation of the user interface A-1

Glossary B-1

Notices C-1
Programming Interface Information C-2
Trademarks C-2

Index X-1

Figures

2-1. Suggested Sequence for Using Monitor III Reports 2-4
2-2. Graphic Delay Report for all Jobs 2-6
2-3. Monitor III Primary Menu 2-15
2-4. Header of Monitor III Single-System Reports 2-20
2-5. Header of Monitor III Sysplex Reports 2-20
2-6. Monitor III Sysplex Report Selection Menu 2-21
2-7. Monitor III Overview Report Selection Menu 2-22
2-8. Monitor III Job Report Selection Menu 2-23
2-9. Monitor III Resource Report Selection Menu 2-24
2-10. Monitor III Subsystem Report Selection Menu 2-25
2-11. Monitor III User Report Selection Menu 2-25
2-12. Data Index 2-27
2-13. Data Index with Preallocated Data Sets - Detailed View 2-28
2-14. Data Index - Condensed Version 2-28
2-15. Data Index with Preallocated Data Sets - Condensed Version. 2-29
2-16. Data Index Options Panel 2-31
2-17. Data Index - Sort Order Descend 2-32
2-18. CACHDET Report 2-33
2-19. CACHDET Report - Volume Details 2-34
2-20. CACHDET Report - SSID Details 2-34
2-21. CACHDET Report Options Panel 2-37
2-22. CACHSUM Report. 2-38
2-23. CACHSUM Report - SSID Details 2-39
2-24. CFACT Report 2-43
2-25. CFACT Report - Details for a List Structure 2-44
2-26. CFACT Report - Details for a Lock Structure 2-44
2-27. CFACT Report - Details for a Cache Structure 2-45
2-28. Coupling Facility Report Options Panel 2-48
2-29. CFOVER Report. 2-49
2-30. CFSYS Report 2-51
2-31. CFSYS Report - Subchannel Details 2-52
2-32. CHANNEL Report 2-56
2-33. CPC Report 2-60
2-34. DELAY Report 2-62
2-35. DELAY Report - Details for Enclave Classification Data 2-67
2-36. DELAY Report Options Panel 2-68
2-37. DELAY Report Job Selection/Exclusion Panel 2-69
2-38. DEV Report 2-71
2-39. DEV Report Options Panel 2-73
2-40. DEVN Report 2-74
2-41. DEVR Report 2-77
2-42. DEVR Report Options Panel 2-81
2-43. DEVT Report. 2-82
2-44. Modified STOR Report Showing all Storage Delays in Detail 2-84
2-45. DSND Report 2-87
2-46. DSND Report Options Panel 2-88
2-47. DSNJ Report 2-89
2-48. DSNV Report 2-91
2-49. DSNV Report Options Panel 2-92
2-50. ENCLAVE Report 2-94
2-51. ENCLAVE Report - Enclave Details 2-95
2-52. ENCLAVE Report Options Panel 2-98
2-53. ENQ Report 2-99
2-54. ENQR Report 2-103
2-55. ENQR Report Options Panel 2-105
2-56. GROUP Report 2-106
2-57. GROUP Report - Response Time Components 2-107
2-58. GROUP Report Options Panel 2-114
2-59. HSM Report 2-115
2-60. IOQUEUE Report 2-118
2-61. JES Report 2-121
2-62. Top Part of Job Report 2-125
2-63. Bottom Part of Job Report 2-125
2-64. Device Delay Variation of the Job Report 2-126
2-65. Enqueue Variation of the Job Report 2-128
2-66. HSM Variation of Job Report 2-129
2-67. Operator Message Variation of the Job Report 2-130
2-68. Processor Delay Variation of the Job Report 2-131
2-69. Quiesce Variation of the Job Report 2-131
2-70. Storage Delay Variation of Job Report 2-132
2-71. XCF Variation of Job Report 2-133
2-72. Job Report Options Panel 2-134
2-73. Migration SYSINFO Report 2-136
2-74. OPD Report. 2-139
2-75. OPD Report - Details for Server Process 2-140
2-76. OPD Report - Details for Non-Server Process 2-140
2-77. OPD Report Options Panel. 2-142
2-78. PROC Report 2-143
2-79. RG Report 2-146
2-80. RLSLRU Report 2-148
2-81. VSAM LRU Overview - Buffer Counts by Pool 2-149
2-82. VSAM RLS Activity by Storage Class - Sysplex Totals View 2-151
2-83. VSAM RLS Activity by Storage Class - System/CF Structure View 2-151
2-84. VSAM RLS Activity by Data Set - Sysplex Totals View 2-152
2-85. VSAM RLS Activity by Data Set - System View 2-153
2-86. STOR Report 2-155
2-87. STORC Report 2-159
2-88. STORC Report Options Panel 2-162
2-89. STORCR Report 2-163
2-90. STORF Report 2-165
2-91. STORR Report 2-168

© Copyright IBM Corp. 1990, 2003

2-92. STORS Report 2-173
2-93. STORS Report Options Panel 2-175
2-94. SYSENQ Report 2-177
2-95. SYSINFO Report 2-179
2-96. SYSRTD Report - With Response Time Data 2-184
2-97. SYSRTD Report - Without Response Time Data 2-185
2-98. SYSSUM Report - GO Mode 2-190
2-99. SYSSUM Report - STOP Mode 2-191
2-100. SYSSUM Report for Workload Groups 2-194
2-101. SYSTREND Report 2-202
2-102. SYSWKM Report for Subsystem CICS 2-205
2-103. SYSWKM Report for Subsystem IMS 2-205
2-104. SYSWKM Report Options Panel 2-212
2-105. WFEX Report 2-214
2-106. WFEX Graphic Report 2-216
2-107. WFEX Report Options Action Panel 2-221
2-108. WFEX Definition and Criteria Panel 2-223
2-109. WFEX Definition of UNAVAIL Panel 2-227
2-110. WFEX Criteria Names for Class: Job 2-228
2-111. WFEX Criteria Names for Class: Service Class 2-229
2-112. WFEX Criteria Names for Class: Processor 2-229
2-113. WFEX Criteria Names for Class: Report Performance Group 2-230
2-114. WFEX Criteria Names for Class: Storage 2-230
2-115. WFEX Criteria Names for Class: Device 2-230
2-116. XCF Report 2-232
3-1. Header of a Monitor II ISPF Session Report 3-3
3-2. Header of a Monitor II TSO Session Report 3-3
3-3. Format of Printed Reports from a Monitor II Display Session 3-6
3-4. Format of Printed Reports from a Monitor II Background Session 3-7
3-5. ARD Report 3-9
3-6. ARD Report in z/Architecture 3-9
3-7. ARDJ Report 3-10
3-8. ARD, ASD, and ASRM Report Options Panel 3-12
3-9. ARDJ, ASDJ, and ASRMJ Report Options Panel 3-12
3-10. ASD Report 3-14
3-11. ASDJ Report 3-14
3-12. ASRM Report 3-18
3-13. ASRMJ Report 3-18
3-14. CHANNEL Report 3-21
3-15. DDMN Report 3-24
3-16. DEV Report 3-28
3-17. DEVV Report 3-28
3-18. DEV Report Options Panel 3-32
3-19. DEVV Report Options Panel 3-32
3-20. HFS Report 3-34
3-21. HFS Report Options Panel 3-36
3-22. ILOCK Report 3-38
3-23. IOQUEUE Report 3-42
3-24. IOQUEUE Report Options Panel 3-44
3-25. LLI Report - Link Library List 3-45
3-26. LLI Report - LPA Library List 3-46
3-27. LLI Report - APF Library List 3-46
3-28. PGSP Page Report 3-49
3-29. SDS Report 3-51
3-30. SENQ Summary Report 3-54
3-31. SENQ GRS Report 3-55
3-32. SENQ Detail Report 3-55
3-33. SENQ Report Options Panel 3-57
3-34. SENQR Report 3-58
3-35. SENQR Report Options Panel 3-59
3-36. SPAG Report 3-60
3-37. SRCS Report 3-62
3-38. TRX Report 3-65
3-39. TRX Report Options Panel 3-66
5-1. Header of a Single-System Report 5-4
5-2. Header of a Sysplex Report 5-5
5-3. Header of a Single-System Duration Report 5-5
5-4. Cache Subsystem Summary Report - Subsystem Summary 5-8
5-5. Cache Subsystem Summary Report - Top-20 Device Lists 5-8
5-6. Cache Subsystem Activity Report - Status and Overview 5-9
5-7. Cache Subsystem Activity Report - Device Overview 5-15
5-8. Cache Subsystem Activity Report - Device-level Reporting 5-17
5-9. Cache Subsystem Activity Report - Cache Offline 5-19
5-10. Structure of the Coupling Facility Report 5-22
5-11. Coupling Facility Activity Report - Usage Summary 5-23
5-12. Coupling Facility Activity Report - Structure Activity 5-28
5-13. Coupling Facility Activity Report - Subchannel Activity 5-31
5-14. CF to CF Activity 5-33
5-15. Channel Path Activity Report 5-37
5-16. CPU Activity Report 5-43
5-17. CPU Activity Report - Partition Data Report Section 5-47
5-18. CPU Activity Report - LPAR Cluster Report Section 5-51
5-19. Crypto Hardware Activity Report 5-56
5-20. Direct Access Device Activity Report 5-63
5-21. Communication Equipment Activity Report 5-63
5-22. Magnetic Tape Device Activity Report 5-64
5-23. Lotus Domino Server Report - Summary 5-71
5-24. Lotus Domino Server Report - Details 5-72
5-25. Enqueue Detail Activity Report 5-77
5-26. FICON Director Activity Report 5-80
5-27. HFS Global Statistics Report. 5-83
5-28. HFS File System Statistics Report 5-84
5-29. HTTP Server Report - Summary 5-87
5-30. HTTP Server Report - Details 5-88
5-31. I/O Queuing Activity Report 5-94
5-32. IOQ Report - z/OS Running as VM Guest System 5-95
5-33. OMVS Kernel Activity Report 5-101
5-34. PAGESP Report 5-105
5-35. PAGING Report - Central Storage Paging Rates Section 5-108

5-36. PAGING Report - Expanded Storage Movement Rates/Frame and Slot Counts Sections 5-113
5-37. PAGING Report - Central Storage Movement Rates/Frame and Slot Counts Sections 5-116
5-38. PAGING Report - Swap Placement Section 5-120
5-39. Shared DASD Activity Report 5-126
5-40. Shared Magnetic Tape Device Activity Report 5-126
5-41. TRACE Report 5-136
5-42. VSTOR Report - Common Storage Summary Section 5-142
5-43. VSTOR Report - Common Storage Detail Section 5-144
5-44. VSTOR Report - Private Area Summary Section 5-145
5-45. VSTOR Report - Private Area Detail Section 5-148
5-46. WKLD Report 5-152
5-47. WLMGL Report Sections 5-154
5-48. WLMGL - Service Class Period Report 5-155
5-49. WLMGL - Service Class Report 5-156
5-50. WLMGL - Workload Group and Service Class Period Report 5-157
5-51. WLMGL Report - Service Policy Page 5-158
5-52. XCF Activity Report - Usage by System 5-175
5-53. XCF Activity Report - Usage by Member 5-176
5-54. XCF Activity Report - Path Statistics (Coupling Facility Structure) 5-177
5-55. XCF Activity Report - Path Statistics (Channel-to-Channel) 5-177
5-56. Exception Report - Low CPU Utilization 5-180
5-57. Exception Report - TSO Response Time 5-181
5-58. Overview Report - Exception Version 5-183
5-59. Overview Report - Summary Version 5-187
5-60. Plot Report 5-192
5-61. Plot Report - Measurement Scale Entries 5-201
5-62. Summary Report 5-203

Tables

1. RMF Library xvi
1-1. Activities that RMF Monitors and SMF Record Types 1-4
2-1. Report Commands 2-17
2-2. User-Written Report Commands 2-19
2-3. Monitor III Heading Information 2-20
2-4. Fields in the Data Index 2-29
2-5. Fields in the CACHDET Report 2-35
2-6. Fields in the CACHDET Report - Volume and SSID Details 2-35
2-7. Fields in the CACHSUM Report 2-39
2-8. Fields in the CACHSUM Report - SSID Details 2-40
2-9. Fields in the CFACT Report 2-45
2-10. Fields in the CFOVER Report 2-49
2-11. Fields in the CFSYS Report 2-52
2-12. Fields in the CFSYS Report - Subchannel Details 2-54
2-13. Fields in the CHANNEL Report 2-56
2-14. Additional Fields in the CHANNEL Report 2-58
2-15. Fields in the CPC Report 2-60
2-16. Fields in the DELAY Report 2-63
2-17. Additional Fields in the DELAY Report 2-66
2-18. Fields in the DEV Report. 2-72
2-19. Fields in the DEVN Report 2-75
2-20. DEVN Report - Cursor-sensitive Control for Select and Sort 2-76
2-21. Fields in the DEVR Report 2-78
2-22. Additional Fields in the DEVR Report 2-80
2-23. Fields in the DEVT Report 2-83
2-24. Fields in the DSD Report. 2-85
2-25. Fields in the DSND Report 2-87
2-26. Fields in the DSNJ Report 2-90
2-27. Fields in the DSNV Report 2-92
2-28. Fields in the ENCLAVE Report 2-95
2-29. Fields in the ENCLAVE Report - Enclave Details 2-96
2-30. Fields in the ENQ Report 2-100
2-31. Fields in the ENQR Report. 2-104
2-32. Fields in the GROUP Report 2-107
2-33. Fields in the HSM Report 2-116
2-34. Fields in the IOQUEUE Report 2-118
2-35. Fields in the JES Report 2-122
2-36. Fields in the Job Report 2-125
2-37. Fields in the Device Delay Variation of the Job Report 2-126
2-38. Fields in the Enqueue Variation of the Job Report 2-128
2-39. Fields in the HSM/JES Variation of the Job Report 2-129
2-40. Fields in the Operator Message Variation of the Job Report 2-130
2-41. Fields in the Processor Delay Variation of the Job Report 2-131
2-42. Fields in the Quiesce Variation of the Job Report 2-132
2-43. Fields in the Storage Delays Variation of the Job Report 2-132
2-44. Fields in the XCF Variation of the Job Report 2-133
2-45. Additional Fields in the Job Report 2-134
2-46. Fields in the MSI Report 2-137
2-47. Fields in the OPD Report 2-139
2-48. Fields in the OPD Details Report 2-141
2-49. Fields in the PROC Report 2-144
2-50. Additional Fields in the PROC Report 2-145
2-51. Fields in the RG Report 2-147
2-52. Fields in the RLSLRU Report 2-149
2-53. Fields in the Buffer Counts by Pool Panel 2-149
2-54. Fields in the VSAM RLS Activity Report 2-153
2-55. Fields in the STOR Report 2-156
2-56. Additional Fields in the STOR Report 2-157
2-57. Fields in the STORC Report - System Information Section 2-159
2-58. Fields in the STORC Report - Job Information Section 2-160
2-59. Additional Fields in the STORC Report 2-161
2-60. Fields in the STORCR Report 2-164
2-61. Fields in the STORF Report 2-165
2-62. Additional Fields in the STORF Report 2-167
2-63. Fields in the STORR Report - Storage Summary Section 2-169
2-64. Fields in the STORR Report - Page/Swap Activity Section 2-170
2-65. Additional Fields in the STORR Report 2-172
2-66. Fields in the STORS Report 2-174
2-67. Additional Fields in the STORS Report 2-175
2-68. Fields in the SYSENQ Report 2-178
2-69. Fields in the SYSINFO Report 2-180
2-70. Additional Fields in the SYSINFO Report 2-182
2-71. Fields in the SYSRTD Report 2-188
2-72. Fields in the SYSSUM Report 2-194
2-73. Fields in the SYSSUM Report - Continuation 2-198
2-74. Additional Fields in the SYSSUM Report 2-199
2-75. Fields in the SYSTREND Report 2-202
2-76. SYSTREND Report - Cursor-sensitive Control for Navigation 2-203
2-77. Fields in the SYSWKM Report 2-208
2-78. Fields in the Speed Section of the WFEX Report 2-217
2-79. Fields in the Exceptions Section of the WFEX Report 2-218
2-80. Fields in the Graphic WFEX Report 2-220
2-81. Fields in the WFEX Action Panel 2-222
2-82. Fields in the WFEX Definition and Criteria Panel 2-224
2-83. Fields in the XCF Report 2-233
3-1. Monitor II Display Session Areas 3-4
3-2. Fields in the ARD and ARDJ Report 3-10
3-3. Fields in the ASD and ASDJ Report 3-14
3-4. Fields in the ASRM and ASRMJ Report 3-18

© Copyright IBM Corp. 1990, 2003

3-5. Fields in the CHANNEL Report 3-21
3-6. Fields in the DDMN Report 3-25
3-7. Fields in the DEV and DEVV Report 3-29
3-8. Fields in the HFS Report 3-34
3-9. Fields in the HFS Report Options Panel 3-36
3-10. Fields in the ILOCK Report 3-38
3-11. Fields in the IOQUEUE Report 3-42
3-12. Fields in the LLI Report 3-47
3-13. Fields in the PGSP Report 3-49
3-14. Fields in the SDS Report 3-52
3-15. Fields in the SENQ Report 3-55
3-16. Fields in the SENQR Report 3-58
3-17. Fields in the SPAG Report 3-61
3-18. Fields in the SRCS Report 3-63
3-19. Fields in the TRX Report 3-65
5-1. Interval and Duration Reports 5-3
5-2. Monitor II Interval Reports 5-3
5-3. Fields in the Cache Subsystem Summary Report 5-8
5-4. Fields in the Cache Subsystem Activity Report - Header 5-10
5-5. Fields in the Cache Subsystem Activity Report - Subsystem Status 5-10
5-6. Fields in the Cache Subsystem Activity Report - Subsystem Overview 5-12
5-7. Fields in the Cache Subsystem Activity Report - Device Overview 5-15
5-8. Fields in the Cache Subsystem Activity Report - RAID Rank Activity 5-16
5-9. Fields in the Cache Subsystem Activity Report - Cache Device Status 5-18
5-10. Overview Names in the Cache Subsystem Activity Report 5-19
5-11. Fields in the Coupling Facility Activity Report - Usage Summary 5-24
5-12. Fields in the Coupling Facility Activity Report - Structure Activity 5-29
5-13. Fields in the Coupling Facility Activity Report - Subchannel Activity 5-31
5-14. Fields in the CF to CF Activity Section 5-33
5-15. Overview Names in the Coupling Facility Activity Report 5-34
5-16. Fields in the Channel Path Activity Report 5-37
5-17. Overview Names in the Channel Path Activity Report 5-40
5-18. Fields in the CPU Activity Report 5-43
5-19. Overview Names in the CPU Activity Report 5-45
5-20. Fields in the Partition Data Report 5-48
5-21. Overview Condition Names Applicable to the Partition Data Report 5-50
5-22. Fields in the LPAR Cluster Report 5-51
5-23. Overview Condition Names Applicable to the LPAR Cluster Report 5-52
5-24. Fields in the CRYPTO Hardware Activity Report 5-56
5-25. Overview Names in the CRYPTO Hardware Activity Report 5-57
5-26. Fields in the Device Activity Reports 5-64
5-27. Overview Names in the DASD Activity Report 5-69
5-28. Fields in the Domino Server Summary Report 5-72
5-29. Fields in the Domino Server Details Report 5-73
5-30. Fields in the Enqueue Activity Report 5-77
5-31. Overview Names in the Enqueue Activity Report 5-78
5-32. Fields in the FICON Director Activity Report 5-80
5-33. Overview Names in the FICON Director Activity Report 5-81
5-34. Fields in HFS Global Statistics Report 5-83
5-35. Fields in the HFS File System Statistics Report 5-84
5-36. Fields in the HTTP Server Summary Report 5-88
5-37. Fields in the HTTP Server Details Report 5-89
5-38. Fields in the I/O Queuing Activity Report 5-95
5-39. Overview Names in the I/O Queuing Activity Report 5-99
5-40. Fields in the OMVS Kernel Activity Report 5-102
5-41. Fields in the Page/Swap Data Set Activity Report 5-105
5-42. Overview Names in the Page/Swap Data Set Activity Report 5-106
5-43. Fields in the Paging Activity Report - Central Storage Paging Rates Section . . 5-109
5-44. Overview Names in the Paging Activity Report - Central Storage Paging Rates Section 5-112
5-45. Fields in the Paging Activity Report - Expanded Storage Movement Rates Section 5-114
5-46. Spreadsheet Range Names and Overview Names in the Paging Activity Report - Expanded Storage Movement Rates Section 5-115
5-47. Fields in the Paging Activity Report - Central Storage Movement Rates Section . 5-116
5-48. Overview Names in the Paging Activity Report - Central Storage Movement Rates Section 5-117
5-49. Fields in the Paging Activity Report - Frame and Slots Counts Section 5-118
5-50. Overview Names in the Paging Activity Report - Frame and Slot Counts Section . 5-119
5-51. Common Swap Reasons in the Paging Activity Report 5-121
5-52. Uncommon Swap Reasons in the Paging Activity Report 5-121
5-53. Fields in the Paging Activity Report - Swap Placement Activity Section 5-123
5-54. Overview Names in the Paging Activity Report - Swap Placement Activity Section . 5-124
5-55. Fields in the Shared Device Activity Reports 5-127
5-56. Overview Names in the Shared DASD Activity Report 5-133
5-57. Fields in the Trace Activity Report 5-136
5-58. Variables in the Trace Activity Report 5-136

5-59. Fields in the Virtual Storage Activity Report - Common Storage Summary . . 5-143
5-60. Overview Names in the Virtual Storage Activity Report 5-144
5-61. Fields in the Virtual Storage Activity Report - Common Storage Detail Section . 5-144
5-62. Fields in the Virtual Storage Activity Report - Private Area Summary 5-145
5-63. Fields in the Virtual Storage Activity Report - Private Storage Map 5-146
5-64. Fields in the Virtual Storage Activity Report - Bottom Half. 5-146
5-65. Fields in the Workload Activity Report 5-159
5-66. Fields in the WLMGL Report - POLICY 5-170
5-67. Exception and Overview Names in the Workload Activity Report 5-171
5-68. Fields in the XCF Activity Report - Usage by System 5-175
5-69. Fields in the XCF Activity Report - Usage by Member 5-177
5-70. Fields in the XCF Activity Report - XCF Path Statistics 5-178
5-71. Fields in the Exception Report 5-182
5-72. Fields in the Overview Report 5-184
5-73. Overview Header Record - Prefix Section 5-188
5-74. Overview Header Record - Header Section 5-188
5-75. Overview Header Record - Report Column Names Section 5-189
5-76. Overview Data Record - Prefix Section 5-189
5-77. Overview Data Record - Data Section 5-189
5-78. Overview Data Record - Report Data Section 5-189
5-79. Fields in the Summary Report 5-203

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*.

For more information on how to go about the task of performance management, see *z/OS RMF Performance Management Guide*.

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"
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"
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"
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"
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 the Chapter 5, "Long-term overview reporting with the Postprocessor"

Chapter 5, "Long-term overview reporting with the Postprocessor"
This chapter describes the reports you can request using the Postprocessor. The descriptions include report examples and detailed descriptions of the report fields.

© Copyright IBM Corp. 1990, 2003

The z/OS RMF library

This table shows the full titles and order numbers of the books in the RMF library for z/OS.

Table 1. RMF Library

Title	Order Number
Books available as Hardcopy and Softcopy	
z/OS *z/OS RMF User's Guide*	SC33-7990
z/OS *z/OS RMF Report Analysis*	SC33-7991
z/OS *z/OS RMF Performance Management Guide*	SC33-7992
z/OS *z/OS RMF Messages and Codes*	SC33-7993
z/OS *z/OS RMF Programmer's Guide*	SC33-7994
z/OS *z/OS RMF Reference Summary*	SX33-9033
Softcopy documentation as part of the *z/OS Collection* (SK3T-4269 (CD-Rom) and SK3T-4271 (DVD))	
z/OS *z/OS RMF NewsFLASH*	SC33-7995

Related information

For additional information on z/OS, see the *z/OS Information Roadmap*, SA22-7500.

Using LookAt to look up message explanations

LookAt is an online facility that allows you to look up explanations for z/OS messages and system abends.

Using LookAt to find information is faster than a conventional search because LookAt goes directly to the explanation.

LookAt can be accessed from the Internet or from a TSO command line.

You can use LookAt on the Internet at:

`www.ibm.com/servers/eserver/zseries/zos/bkserv/lookat/lookat.html`

To use LookAt as a TSO command, LookAt must be installed on your host system. You can obtain the LookAt code for TSO from the LookAt Web site by clicking on the **News and Help** link or from the *z/OS Collection*, SK3T-4269.

To find a message explanation from a TSO command line, simply enter: **lookat** *message-id* as in the following:

```
lookat erb100i
```

This results in direct access to the message explanation for message ERB100I.

To find a message explanation from the LookAt Web site, simply enter the message ID and select the release you are working with.

Summary of changes

What's new in z/OS V1R4

This document contains information previously presented in *z/OS RMF Report Analysis*, SC33-7991-01, which supports the Resource Measurement Facility.

For z/OS V1R4, there are several updates of this document, SC33-7990-02 through SC33-7991-05, which are summarized in the following sections.

This document includes terminology, maintenance, and editorial changes. The following information describes the enhancements that are being distributed with z/OS Version 1 Release 4. All technical changes or additions to the text are indicated by a vertical line to the left of the change.

What's new in SC33-7991-05 (z990 Exploitation)

The new functionality for z990 exploitation introduced in the following subsections is available as SPE and needs to be installed as APAR OW56656.

Enhanced reporting on I/O queuing activity

Performance monitoring on I/O queuing has been enhanced in the Monitor III, Monitor II and Postprocessor I/O Queuing reports. The following data is now available:

- Average control unit busy delay time
- Average initial command response time
- Average channel subsystem delay time

In addition, new overview and exception conditions are available for the Postprocessor.

Support of new cryptographic hardware

With the availability of the new cryptographic hardware (PCIX cryptographic card), the Postprocessor Crypto Hardware Activity report is enhanced to provide performance monitoring for the PCIX Cryptographic Coprocessor.

What's new in SC33-7991-03 and SC33-7991-04 (z990 Compatibility)

- The following delay metrics or pending reasons have been removed from the Device Activity reports:
  - director port busy
  - control unit busy
- The following delay metric or pending reason has been added to the Device Activity reports:
  - initial command response time
- The overview and exception conditions for device activity have changed. The following conditions are no longer available:
  - CUBDL (Average control unit busy delay time)
  - DPBDL (Average director port busy delay time)

  while the following new condition is supported:

© Copyright IBM Corp. 1990, 2003

– CMRDL (Average command response delay time)

This new functionality is available as SPE and needs to be installed as APAR OW54347.

What's new in SC33-7991-02

Note:
An appendix with z/OS product accessibility information has been added.

Support of the cryptographic hardware

The **Crypto Hardware Activity** report, introduced with RMF for z/OS 1.2, now includes measurements for the cryptographic coprocessor facility (CCF). In addition, new overview conditions are available for the Postprocessor.

This functionality is available as SPE and needs to be installed as APAR OW49808.

Support of FICON switch cascading

The **FICON Director Activity** report is enhanced to indicate that another switch is connected to a port in case of FICON cascaded switches.

This functionality is available as SPE and needs to be installed as APAR OW52396.

State samples breakdown in the WLMGL report

Up to this release, state samples have been reported as a percentage of average transaction response time (response time breakdown). The response time is calculated when a transaction completes. This can result in percentages greater than 100 when samples are included for long running transactions which have not completed in the gathering interval.

Percentages greater than 100 in the breakdown section are now avoided by showing the state values as percentages of the total transaction samples (state samples breakdown) instead of percentages of response time.

This functionality is available as SPE and needs to be installed as APAR OW52227.

History of changes

What's new in z/OS V1R2

Support of the cryptographic hardware

With the availability of the new cryptographic hardware (PCI Cryptographic Coprocessors and PCI Cryptographic Accelerators) on an LPAR basis, RMF now provides performance monitoring with the new Postprocessor **Crypto Hardware Activity** report which is based on new SMF records type 70 subtype 2.

In addition, new overview criteria are available for the Postprocessor, and Crypto Using and Delay values are shown in the Postprocessor **WLMGL** report.

This report is available as SPE and needs to be installed as APAR OW49808.

Reporting of report class periods

Statistics on response times is a key performance metric used for groups of work having response time objectives. WLM report classes allow aggregation of performance data so that you can evaluate performance of applications

independent of the definition how they are managed. Now, reporting of report classes is available with the same granularity as for service classes.

This can be seen in the Postprocessor **WLMGL** report, and in several Monitor III reports: **GROUP**, **STORS**, **SYSINFO**, **SYSRTD**, **SYSSUM**, and **SYSWKM**.

Enhanced reporting for the coupling facilities

CF duplexing ensures high application availability in a Parallel Sysplex. The performance management aspects which have to be covered for CF duplexing are provided by RMF in the Postprocessor **CF Activity** report about new peer CF connectivity. This enables you to evaluate and monitor your CF configuration, and you can apply the necessary changes to tune accommodation of new structure instances resulting from system-managed duplexing.

Software pricing online support

A new **CPC Capacity** report provides the capability to monitor values that are relevant for software pricing and also monitors partition related processor activities.

This report gives you an indication whether the defined capacity limits are correctly set. You can see the actual CPU consumption, the four-hours average, as well as the defined capacity limit. These values show whether the capacity settings are appropriate to the workload running on the system.

This report is available as SPE and needs to be installed as APAR OW49807.

Support of HiperSockets

HiperSockets can be used for a very fast communication among operating systems images running in different partitions of a processor. Reporting is done in all Channel Activity reports.

Enhanced reporting of IOP utilization

The Monitor I I/O Queuing report has been enhanced by providing more information about the utilization of I/O processors. This will help customers in analyzing their IOP capacity and in planning for future growth.

These enhancements are available as SPE which needs to be installed as APAR OW49806.

What's new in z/OS V1R1

The following information describes the enhancements that are being distributed with z/OS Version 1 Release 1. It is indicated by a vertical line to the left of the text.

There is one release RMF 2.10 that can run in OS/390 2.10 as well as in z/OS V1R1. Nevertheless, some new functions are available only when z/OS is running on a zSeries 900 server.

Intelligent Resource Director (IRD)

The Workload Manager is extended to work with PR/SM on zSeries 900 (z900) servers 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. This combination of WLM working with PR/SM on a z900 server is called IRD.

IRD is made up of three parts which work together to help increase your business' productivity:

- **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 Postprocessor **CPU Activity (Partition Data)** report which provides LPAR views as well as aggregated views on LPAR cluster level.

  In addition, new overview criteria are available for the Postprocessor.

- **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.

  All channel and I/O queuing reports have been extended to differentiate static channels from floating channels.

  In addition, new overview criteria are available for the Postprocessor.

- Channel Subsystem Priority Queuing

  This topic is not reflected directly in any RMF report.

This powerful triad of functions can increase your productivity by putting your most business-critical work first.

zSeries software pricing

Workload License Charges (WLC) is a pricing model that provides more flexibility, supports the variability and growth of e-business workloads, and improves the cost of computing. Customers will have the flexibility to configure their system to match workload requirements and be charged for the products used by these workloads at less than the full machine capacity, allowing customers to pay for what they need.

The Postprocessor **CPU Activity (Partition Data)** report will show CPU resource consumption within an LPAR in terms of millions of service units (MSUs) and the corresponding LPAR MSU defined capacity. This helps you to understand how much of the defined capacity an LPAR is consuming.

In addition, new overview criteria are available for the Postprocessor.

FICON director

RMF offers new reporting capabilities for the FICON director. Due to the different technology and implementation compared to ESCON, the new Postprocessor **FICON Director Activity** report will provide information about director and port activities. This will assist you in analyzing performance problems and in capacity planning.

Goal mode enhancements

The Workload Manager has implemented some enhancements to help increase production use of goal mode. They are

- Classification by System Name
- Separation of Production from Test CICS/IMS Regions
- Address Space Storage Isolation
- Pro-active CPU Protection

This is done by new specifications (CPU protection and storage protection) in the service policy.

RMF supports these enhancements in the **WLMGL** Postprocessor report and in several online reports:

- Monitor II ARD/ARDJ report
- Monitor III SYSWKM report
- Monitor III DELAY report
- Monitor III JOB report
- Monitor III STORF report
- Monitor III ENCLAVE report - Details

Chapter 1. Introducing RMF data gathering and reporting

RMF Performance Monitoring

This book 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

© Copyright IBM Corp. 1990, 2003

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 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.

Plot report: Presents a graphic summary of a number of key system performance indicators.

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*, and you find many report samples in the *z/OS RMF Performance Management 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

RMF Performance Monitoring (RMF PM) gives you the capability to construct monitoring scenarios and use them whenever necessary. This is done on the Windows (95, 98, NT) or OS/2 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 the *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 1-1. Activities that RMF Monitors and SMF Record Types

Gathering				Activity	Reporting			
Short-term Mon III		Snapshot Mon II	Long-term Mon I		Interactive Mon III	Snapshot Mon II	Real-time Mon I	Long-term Post-processor
SMF	VSAM	SMF	SMF					
	★	79.1/2/5		Address space	★	★		★
	★		74.5	Cache	★			★
	★	79.12	73	Channel path	★	★	★	★
74.4	★			Coupling facility	★			★
			70.2	Cryptographic hardware			★	★
	★	79.9	74.1	Device	★	★	★	★
	★	79.10		Domain	★	★	★	★
	★			Enclave	★			
	★	79.7	77	Enqueue	★	★	★	★
			74.7	FICON director				★
		79.15		IRLM long locks		★		
	★	79.13/14	78.3	I/O queuing	★	★	★	★
	★		70.1	License manager	★		★	★
		79.11	75	Page data set		★	★	★
		79.4	71	Paging		★	★	★
	★	79.3	70.1	Processor	★	★	★	★
		79.6		Reserve		★		★
72.2/4	★	79.3		Storage	★	★		★
			76	System counters			★	★
		79.8		Transactions		★		★
74.3/6	★			UNIX	★	★		★
	★		78.2	Virtual storage	★		★	★
	★ ★		72.1 72.3	Workload Performance groups Service classes	★ ★		★	★ ★

Table 1-1. Activities that RMF Monitors 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					
74.2	★			XCF	★			★

Delay monitoring

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

Delayed address spaces and groups

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

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

For any service class, report class and workload group (or domain and performance 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

Reporting of other SMF data

In addition to reporting of data that have been gathered by the RMF monitors, 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 the 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

All About Monitor III

This chapter:

- 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

© Copyright IBM Corp. 1990, 2003

Using Monitor III reports

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
- Domains and performance groups
- 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 domain, performance group, or service class period Monitor III provides a breakdown of response time and displays the information on the **Group Response Time** report.

See the "DELAY - Delay report" on page 2-62 for more details.

Where to start

You can have two possible environments:

- A sysplex environment (your system is running in goal mode)
- A single-system environment (your system is running in compatibility mode)

If you have a sysplex environment

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.

With the introduction of goal mode, it is necessary to:

- Report goal versus actual values in the new notion of MVS workload management
- 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 some information from that report is needed for additional investigations, or navigation should continue 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″ in to any of the listed systems. Placing the cursor on the system-ID in one of the rows and pressing ENTER, gives 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 the next chapter.

If you have a single-system environment

If you do not have a sysplex, and your system runs in compatibility mode, you can start with the Workflow/Exceptions (WFEX) report.

In this 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 2-5 describes how to invoke reports using cursor sensitive control.

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 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 (DEVR) report and the resource-oriented enqueue (ENQR) report. For storage problems involving paging or swapping delays, you can use the resource-oriented storage (STORR) delay report. 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 report.

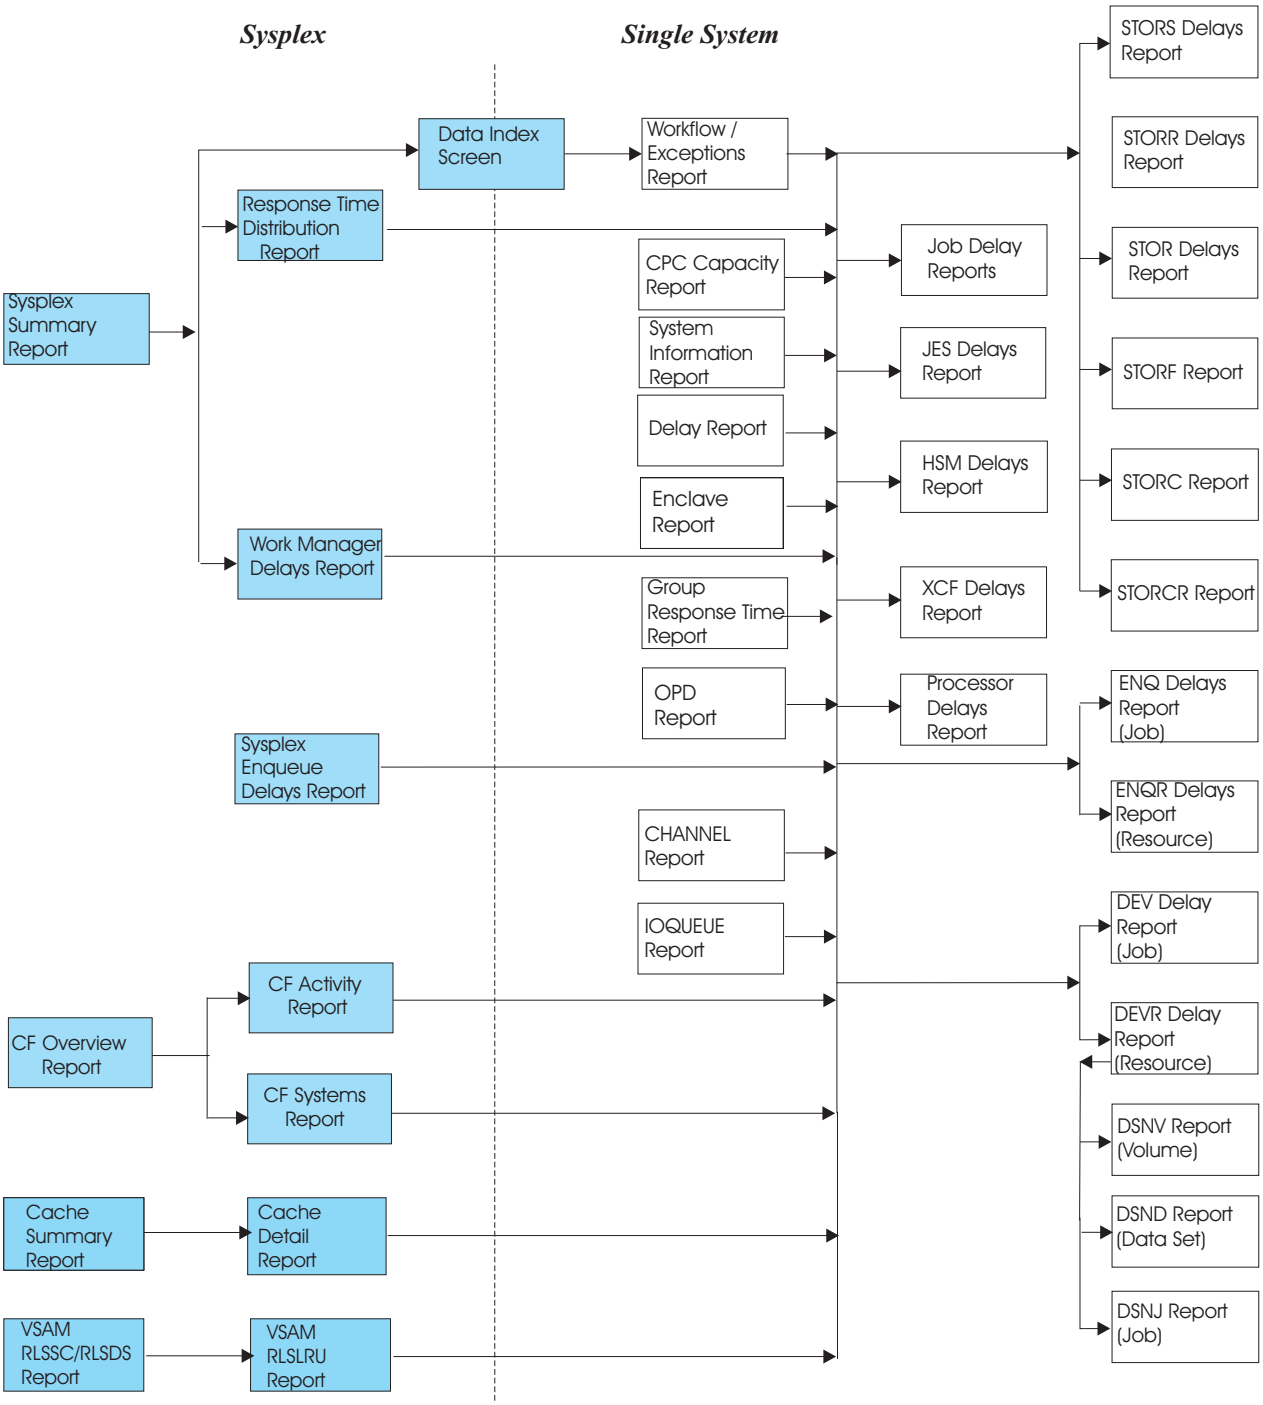

Figure 2-1. Suggested Sequence for Using Monitor III Reports

For a summary of common system storage consumption, use the Common Storage Summary (STORC) report. 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 2-1 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 panel 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 navigate easily 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.

Choosing graphic and tabular reports

Most Monitor III reports are available in both tabular and color graphic form. Tabular reports provide detailed information about jobs and resources. Graphic reports provide quick information about potential problem areas. Use the TOGGLE command (PF6) or the GRAPHIC and TABULAR commands to display the different forms. Figure 2-2 shows a sample graphic report.

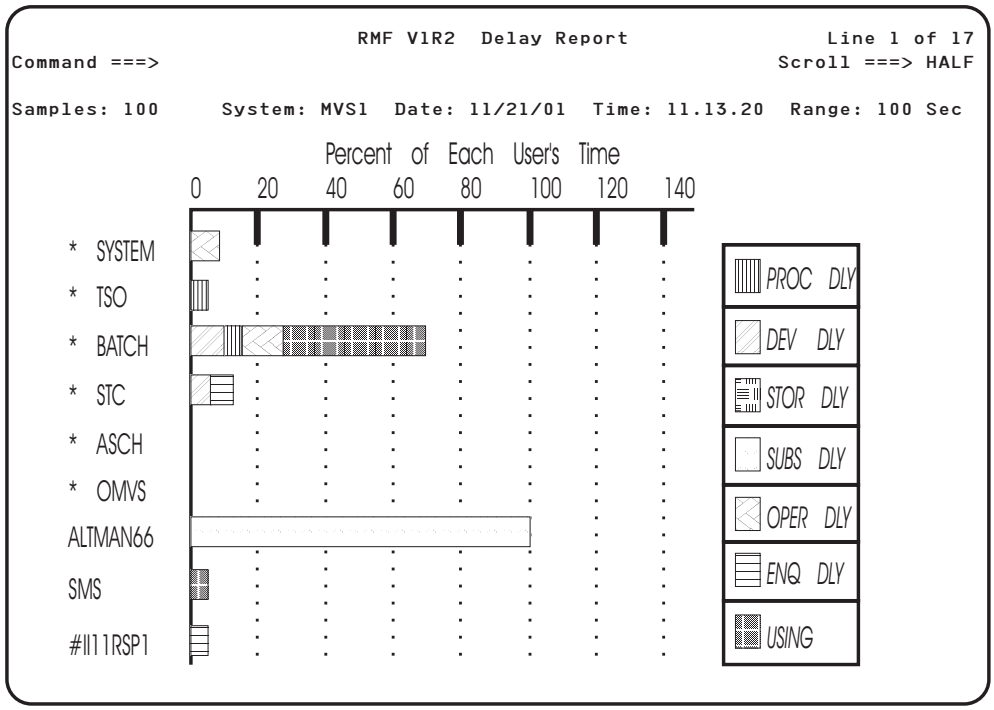

Figure 2-2. Graphic Delay Report for all Jobs

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. An address space is considered using a processor(s) 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 a device(s) 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. 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. However, the address space with the first ready unit of work is not considered delayed because it would have been using the processor currently being used by Monitor III (if Monitor III was not using the processor). 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

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 six 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%.

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 queueing 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

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

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

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

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 should run when ready, without being delayed for resources.

The calculation of the execution velocity is:

Execution Velocity

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

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

Monitor III gatherer/reporter concepts

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 will be performed based on this MINTIME interval and is defined by the Range value. Range can be set either on the Session Options panel 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

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 will be gathered in the following intervals:

Data Gathered

```
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, when you press Enter your report heading will change 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 mode change
A report interval containing a mode change will have both compatibility and goal mode data.

A policy switch (in goal mode)
A report interval containing a policy switch can have data with different gatherer options.

An IPS- or ICS-member switch (in compatibility mode)
A report interval containing an IPS- or ICS-member switch can have data with different gatherer options.

A system IPL

A change of the gatherer CYCLE time

RMF cannot combine data of two different modes, or 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:

Data Gathered

```
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 report 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 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

- 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.

Notes:

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 V1R2 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               (RES)
  4 SUBS            Subsystem information for HSM, JES, and XCF           (SUB)

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


                    O OPTIONS    T TUTORIAL    X EXIT

       5694-A01 (C) Copyright IBM Corp. 1986, 2001. All Rights Reserved
                     Licensed Materials - Property of IBM
```

Figure 2-3. 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 select the Delay report, enter the report command DELAY on any selection or command line:

```
Command ===> delay
```

Or enter a 1 on the Primary menu:

```
Selection ===> 1
```

and then a 3 on the Overview Report menu:

```
Selection ===> 3
```

Or select the report from any other report panel by using the ISPF 'jump' function:

```
Command ===> =1.3
```

Or select the report from any other report panel by using the ISPF 'jump' function:

```
Command ===> =7
```

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 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

The following table 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. Abbreviated commands are shown in Table 2-1 on page 2-17.

You can enter the commands on any command line.

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

cfname A coupling facility name

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

s-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

stclass 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

Note: For systems running in compatibility mode, the terms s-class, s/r-class, and wlm have to be replaced by DMNnnn (domain) or PGNnnn (performance group).

Table 2-1. 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	class,s-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	class,s-class	Device delays report for all jobs or specified job groups	DD, DVD

Table 2-1. Report Commands (continued)

Command	Parameters	Displays	Abbreviation
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	class,s-class	Enqueue delays report for all jobs or specified job groups	ED
ENQJ	jobname	Enqueue delays variation of job report for specified jobname	EJ
ENQR	resource	Enqueue delays for all or specified resources	ER
GROUP	s/r-class, period	Group response time breakdown	GP, GRP, GD, RT, GRT
HSM	class,s-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	class,s-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
MNTJ	jobname	Operator delays variation for mount request of job report for specified jobname	MTJ
MSGJ	jobname	Operator delays variation for message request of job report for specified jobname	MSJ
OPD		OMVS process data	
PROC	class,s-class	Processor delays report for all jobs or specified job groups	PD
PROCJ	jobname	Processor delays variation of job report for specified job	PJ
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	stclass	VSAM RLS activity by storage class	RLS
STOR	class,s-class	Storage delays report for all jobs or specified job group	SD

Table 2-1. Report Commands (continued)

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

Table 2-2 contains commands for the examples of user-written reports that were delivered with RMF.

Table 2-2. User-Written Report Commands

Command	Parameters	Displays	Abbreviation
DEVN		Device activity	DA
DEVT		Device trend	DT
DSD		Detailed storage delays	
MSI		Migration SYSINFO including execution velocity	
RG		Resource group data	
SYSTREND		System trend	ST

Header for single-system reports

Figure 2-4 shows the common header for single-system Monitor III reports.

```
                     RMF V1R2  TITLE                        Line 1 of 30
Command ===>                                               Scroll ===> HALF

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

Figure 2-4. Header of Monitor III Single-System Reports

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

Table 2-3. 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 report header for sysplex reports is very similar to the header for single-system reports. It differs only in three fields:

1. The sysplex reports show the sysplex name, where the single-system reports show the SMF system identification. In the example shown below, **SYSPLEXN** is the 8 character sysplex name.
2. 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. Finally, the number of systems participating in the report is shown. So it is possible to see 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 2-5.

```
                   RMF V1R2  TITLE - SYSPLEXN            Line 1 of 30
Command ===>                                             Scroll ===> HALF

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

Figure 2-5. 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 SYSENQ   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)

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

Figure 2-6. 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 2-26.

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)
         5 GROUP    Group response time breakdown                (RT)
         6 ENCLAVE  Enclave resource consumption and delays    (ENCL)
         7 OPD      OMVS process data
```

Figure 2-7. Monitor III Overview Report Selection Menu

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

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 a **2** from the Primary Menu or enter JOBS on any command line. Use this menu to choose the specific job you want to analyze and the type of delay you want reported.

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

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

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

  Jobname ===> ________

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


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

Figure 2-8. 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)
  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)
                 8 STORR    Storage usage for each resource      (SR)
                 9 STORS    Storage summary for each group       (SS)
                10 STORC    Common storage summary               (SC)
                11 STORCR   Common storage remaining            (SCR)
  I/O Subsystem 12 CHANNEL  Channel path activity                (CH)
                13 IOQUEUE  I/O queuing activity                 (IQ)
```

Figure 2-9. 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 a **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 2-10. 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 a **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 2-11. Monitor III User Report Selection Menu

The Data Index

How to Request the Data Index

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
```

Purpose of 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.

Contents of the Data Index

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 2-12 on page 2-27 shows a sample Data Index that is using data from active Monitor III data gatherers in a sysplex.

```
                        RMF V1R2  Data Index - RMFPLEX1          Line 1 of 22
Command ===>                                                   Scroll ===> HALF

Samples: 118     System: MVS2  Date: 11/29/02  Time: 10.12.00  Range: 120   Sec

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

MVS1   11/29/02 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   11/29/02 09.11.00 SYS00002 RMF.MONITOR3.DATASET1.MVS2
                09.14.00
MVS2   11/29/02 10.03.00 SYS00003 RMF.MONITOR3.DATASET2.MVS2
                10.12.00             * * *      Currently active      * * *
MVS2   11/29/02 10.03.00
                10.12.00             * * *      In-storage buffer     * * *

MVS3   11/29/02 09.11.00 SYS00002 RMF.MONITOR3.DATASET1.MVS3
                09.14.00
MVS3   11/29/02 10.03.00 SYS00003 RMF.MONITOR3.DATASET2.MVS3
                10.12.00             * * *      Currently active      * * *
MVS3   11/29/02 10.03.00
                10.12.00             * * *      In-storage buffer     * * *
TEST
                                     * * *      No response        * * *
```

Figure 2-12. 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*.

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 2-13 on page 2-28 shows a sample Data Index that is using data from preallocated data sets to a reporter session.

```
                        RMF V1R2  Data Index - RMFPLEX           Line 1 of 18
Command ===>                                                   Scroll ===> HALF

Samples: 37      System: RMFB  Date: 11/29/02  Time: 13.09.00  Range: 60    Sec

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

RMFA   05/09/02 04.17.00 RMFDS01  RMF.RMFA.RMFDS01
                06.26.00
RMFA   05/12/02 04.20.00 RMFDS00  RMF.RMFA.RMFDS00
                06.28.00
RMFA   05/12/02 11.28.00 RMFDS02  RMF.RMFA.RMFDS02
                13.07.00
RMFB   05/09/02 04.16.00 RMFDS04  RMF.RMFB.RMFDS01
                06.24.00
RMFB   05/12/02 04.26.00 RMFDS05  RMF.RMFB.RMFDS00
                06.36.00
RMFB   05/12/02 11.58.00 RMFDS03  RMF.RMFB.RMFDS02
                13.10.00
RMFC   05/09/02 04.17.00 RMFDS07  RMF.RMFC.RMFDS01
                06.26.00
RMFC   05/12/02 04.27.00 RMFDS06  RMF.RMFC.RMFDS00
                06.37.00
RMFC                     RMFDS07  RMF.RMFC.RMFDS02
                                      * * *          Empty         * * *
```

Figure 2-13. 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 2-14 and Figure 2-15 on page 2-29 show what the data displayed in Figure 2-12 on page 2-27 and Figure 2-13, respectively, look like if the DDNAMES/DDNAMES option is used to condense the information displayed.

```
                        RMF V1R2  Data Index - RMFPLEX1           Line 1 of 4
Command ===>                                                   Scroll ===> HALF

Samples: 118     System: RMFE  Date: 11/29/02  Time: 10.12.00  Range: 120   Sec

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

MVS1   11/29/02 10.03.20          11/29/02 10.12.00
MVS2   11/29/02 09.11.00          11/29/02 10.12.00
MVS3   11/29/02 09.11.00          11/29/02 10.12.00
TEST                                   * * *   No response        * * *
```

Figure 2-14. Data Index - Condensed Version

```
                        RMF V1R2  Data Index - RMFPLEX           Line 1 of 3
Command ===>                                                   Scroll ===> HALF

Samples: 37      System: RMFB  Date: 05/13/02  Time: 13.09.00  Range: 60    Sec

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

RMFA   05/09/02 04.17.00          05/12/02 13.07.00
RMFB   05/09/02 04.16.00          05/12/02 13.10.00
RMFC   05/09/02 04.17.00          05/12/02 06.37.00
```

Figure 2-15. 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.

Field descriptions

For a description of the report header area of the Data Index, refer to "Header for single-system reports" on page 2-20, and "Header for sysplex reports" on page 2-20, 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 2-4. 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 will be blank. Also, note that the field will be blank if this line is representing the gatherer's in-storage buffer.
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" on page 2-30.) 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 MVS system name xxxxxxxx is defined in the SYS1.PARMLIB(IEASYSxxx) parameter SYSNAME.

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

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

***** Not Found *****
Uncataloged data set specified on the DATASET option of the Monitor III data gatherer session (the data set is unusable)

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

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

***** Open Error RC=xx reason=xxx *****
Error in opening the data set (the data set is unusable)

***** Close Error RC=xx reason=xxx *****
Error in closing the data set (the data set is unusable)

***** VSAM error RC=xx reason=xxx *****
Error in reading the VSAM data set (the data set is unusable)

***** DYNALLOC RC=xx IRC=xxxx ERC=xxxx *****
Dynamic allocation error (the data set is unusable)

***** UNALLOC RC=xx IRC=xxxx ERC=xxxx *****
Data set unallocated (the data set is unusable)

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

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

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

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

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

```
                        RMF Data Index Options
Command ===>

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

   DDNAMES/DSNAMES  ==> YES      Include DDNAMES / DSNAMES information (YES NO)
   Sort Order       ==> ASCEND   Sort data set names (ASCEND DESCEND)
```

Figure 2-16. 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 2-12 on page 2-27 and Figure 2-13 on page 2-28 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 2-14 on page 2-28 and Figure 2-15 on page 2-29 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 2-12 on page 2-27 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 2-17 for an example.

```
                  RMF V1R2  Data Index - RMFPLEX1             Line 1 of 14
Command ===>                                                 Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  Time: 10.31.40  Range: 100   Sec

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

 MVS1  11/29/02 09.55.00
               10.31.40                 * * *   In-storage buffer  * * *
 MVS1  11/29/02 05.53.20  SYS00003 SYS3.RMF.DS02
               10.11.40                 * * *   Currently active   * * *
 MVS1  11/29/02 01.31.40  SYS00002 SYS3.RMF.DS01
               05.53.20

 MVS2  11/29/02 09.55.00
               10.31.40                 * * *   In-storage buffer  * * *
 MVS2  11/29/02 01.31.40  SYS00003 SYS3.RMF.DS01
               10.11.40                 * * *   Currently active   * * *

 TEST
                                        * * *   No response        * * *
```

Figure 2-17. 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

How to request this report

To request the CACHDET report, select an **S** on the Primary menu, and then a **9** on the Sysplex Report menu (shown in Figure 2-6 on page 2-21), or you can enter the following command:

```
CACHDET [subsystem-id]
```

Purpose of the report

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

Contents of the report

```
                    RMF V1R2  Cache Detail     - SYSPLEX        Line 1 of 20
Command ===> _                                                 Scroll ===> HALF

Samples: 120     Systems: 5    Date: 11/29/02  Time: 16.11.30  Range: 120   Sec
                             CDate: 11/29/02 CTime: 16.11.25 CRange: 119   Sec

Volume /Num SSID   I/O  I/O    Hit - Cache Hit Rate - - DASD I/O -  Seq   Async
                    %   Rate    %   Read   DFW   CFW  Total Stage   Rate  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 2-18. CACHDET Report

There is no graphic version of this report available.

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

```
                    RMF V1R2  Cache Detail     - SYSPLEX        Line 1 of 20
Command ===> _                                                 Scroll ===> HALF

Samples: 120     Systems: 5    Date: 11/29/02  Time: 16.11.30  Range: 120   Sec
                              CDate: 11/29/02 CTime: 16.11.25 CRange: 119   Sec

Vol
                              RMF Cache Volume Detail

*AL  The following details are available for Volume SYSSM6 on SSID 0043
*NO  Press Enter to return to the Report panel.
*CA
SYS  Cache: Active        DFW: Active        Pinned: None
SYS
SYS         ------ Read ------   --------- Write ---------   Read   Tracks
SYS          Rate    Hit  Hit%    Rate   Fast    Hit  Hit%     %
SYS  Norm     3.7    3.6  98.8     0.8    0.8    0.8   100   82.2      0.1
SYS  Seq      0.0    0.0   100     0.1    0.1    0.1  93.3   21.1      0.0
SYS  CFW      0.0    0.0   0.0     0.0    0.0    0.0   0.0    0.0
SYS  Total    3.7    3.7  98.9     0.9    0.9    0.9  99.1   80.0
SYS
SYS  ------ Misc ------  - Non-Cache -  --- CKD ----  - Record Caching -
SYS  DFW Bypass :   0.0  ICL   :   0.0  Write:   0.0  Read Miss :   0.0
SYS  CFW Bypass :   0.0  Bypass:   0.0  Hits :   0.0  Write Prom:   0.0
SYS  DFW Inhibit:   0.0
SYS
SYS
SYS
MVS
```

Figure 2-19. CACHDET Report - Volume Details

```
                    RMF V1R2  Cache Detail     - SYSPLEX        Line 1 of 20
Command ===> _                                                 Scroll ===> HALF

Samples: 120     Systems: 5    Date: 11/29/02  Time: 16.11.30  Range: 120   Sec
                              CDate: 11/29/02 CTime: 16.11.25 CRange: 119   Sec

Vol
                               RMF Cache SSID Detail

*AL  The following details are available for SSID 0043
*NO  Press Enter to return to the Report panel.
*CA
SYS         ------ Read ------   --------- Write ---------   Read   Tracks
SYS          Rate    Hit  Hit%    Rate   Fast    Hit  Hit%     %
SYS  Norm    26.8   26.7  99.3    11.5   11.5   11.5  99.9   70.0      0.2
SYS  Seq      4.1    4.1  99.8     0.4    0.4    0.4  97.8   91.3      0.6
SYS  CFW      0.0    0.0   0.0     0.0    0.0    0.0   0.0    0.0
SYS  Total   31.0   30.8  99.4    11.9   11.9   11.9  99.9   72.2
SYS
SYS  ------ Misc ------  - Non-Cache -  --- CKD ----  - Record Caching -
SYS  DFW Bypass :   0.0  ICL   :   0.0  Write:   0.0  Read Miss :   0.0
SYS  CFW Bypass :   0.0  Bypass:   0.0  Hits :   0.0  Write Prom:   0.0
SYS  DFW Inhibit:   0.0
SYS
SYS
SYS
```

Figure 2-20. CACHDET Report - SSID Details

Field descriptions

Table 2-5. 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 where 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 2-6. Fields in the CACHDET Report - Volume and SSID Details

Field Heading	Meaning
Cache	Caching status of the device. **Active** Caching is active - requests to this device can be processed without DASD access. **Deact pending** Cache has been deactivated on request from host system or support facility, but transfer of modified data to DASD has failed. **Deactivated** Caching has been deactivated for this device.

Table 2-6. Fields in the CACHDET Report - Volume and SSID Details (continued)

Field Heading	Meaning
DFW	Status of the DASD FAST WRITE option. **Active** DASD FAST WRITE requests can be processed for this device. **Deact pending** 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. **Deactivated** DASD FAST WRITE requests are ignored for this device.
Pinned	A device has failed, and data that has not yet been written to DASD is pinned in cache or NVS for later recovery. **None** No data is pinned for this device. **DFW not suspend** Pinned data exists, but DASD FAST WRITE has not been suspended. **DFW suspend** 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 described above.
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.

Table 2-6. Fields in the CACHDET Report - Volume and SSID Details (continued)

Field Heading	Meaning
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.
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 not found in the cache, and access to DASD was 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 2-21. CACHDET Report Options Panel

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

How to request this report

To request the CACHSUM report, select an **S** on the Primary menu, and then an **8** on the Sysplex Report menu (shown in Figure 2-6 on page 2-21), or you can enter the following command:

```
CACHSUM
```

Purpose of the 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 2-33).

Contents of the report

```
                    RMF V1R2  Cache Summary    - SYSPLEX        Line 1 of 12
Command ===> _                                                 Scroll ===> HALF

Samples: 120     Systems: 5    Date: 11/29/02  Time: 16.11.30  Range: 120   Sec
                             CDate: 11/29/02 CTime: 16.11.25 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 2-22. CACHSUM Report

There is no graphic version of this report available.

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

```
                    RMF V1R2  Cache Summary    - SYSPLEX        Line 1 of 12
Command ===> _                                                 Scroll ===> HALF

Samples: 120     Systems: 5    Date: 11/29/02  Time: 16.11.30  Range: 120   Sec
                              CDate: 11/29/02 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 2-23. CACHSUM Report - SSID Details

Field description

Table 2-7. 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.	
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.

Table 2-7. Fields in the CACHSUM Report (continued)

Field Heading	Meaning
Hit Rate	I/O rate of all cache hits. This is the sum of: **READ** Rate of SEARCH/READ requests that completed without accessing the DASD **DFW** Rate of DASD FAST WRITE requests **CFW** 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 2-8. Fields in the CACHSUM Report - SSID Details

Field Heading	Meaning
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.
Cache	Caching status of the subsystem. **Active** Caching is active (online and usable). **Deact failed** A command requesting deactivation of cache has been received, but destaging to DASD has failed. **Deact in process** A command requesting deactivation of cache has been received, and destaging from cache to DASD is still in progress. **Deactivated** Cache has been deactivated by request from host system or support facility. **Error termination** An internal error stopped cache (cache is offline). **Maintenance** Cache has been disabled for maintenance. **Pending active** Caching is pending active, that is, cache is in the process of being brought online.
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.

Table 2-8. Fields in the CACHSUM Report - SSID Details (continued)

Field Heading	Meaning
NVS	Overall status of the non-volatile storage (NVS). **Active** NVS is online and usable. **Deact failed** A command requesting deactivation of NVS has been received but transfer from NVS to DASD has failed. **Deact in process** A command requesting deactivation of NVS has been received, and destaging to DASD is still in progress. **Deactivated** NVS has been deactivated by request from host system or support facility. **DFW inhibited** DASD FAST WRITE is inhibited because the battery is defective. **Error termination** An internal error caused termination of NVS. **Maintenance** NVS has been disabled for maintenance by the support facility.
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 2-6 on page 2-35 for a description of the other fields in the pop-up panel of the CACHSUM report.

CFACT - Coupling Facility Activity report

How to request this report

To request the Coupling Facility Activity report, select an **S** from the Primary menu and then a **7** on the Sysplex Report menu (shown in Figure 2-6 on page 2-21), or you can 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.

Purpose of the 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 panel 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.

Special considerations

For data gathering, you have specified (or used the default value in the ERBRMF04 Parmlib member) the gathering option CFDETAIL or NOCFDETAIL. With CFDETAIL, a large amount of data is being gathered that enables you to get many details about the usage of each structure in the coupling facility. Please, consider that this data gathering is done only on one member of the sysplex. This is called *sysplex master gathering* and has been implemented to reduce performance overhead on non-master members and to reduce the amount of data in SSHs and SMF records. The RMF Sysplex Data Server determines internally which member of the sysplex will be the master. This cannot be controlled externally by the operator or system administrator and is transparent for performance monitoring.

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

```
                  RMF V1R2  CF Activity      - UTCPLXHD      Line 13 of 76
Command ===>                                                Scroll ===> HALF

Samples: 120     Systems: 3    Date: 11/29/02  Time: 11.27.00  Range: 120   Sec

CF: ALL          Type  ST System      --- Sync ---     --------- Async --------
                                        Rate   Avg       Rate   Avg   Chng   Del
Structure Name                                Serv             Serv   %      %

IRRXCF00_B001    CACHE    *ALL          0.1    372       0.0      0   0.0   0.0
                          D0            0.0      0       0.0      0   0.0   0.0
                          D1            0.0      0       0.0      0   0.0   0.0
                          D2            0.1    372       0.0      0   0.0   0.0
IRRXCF00_P001    CACHE    *ALL          0.5    378       0.0      0   0.0   0.0
                          D0            0.2    353       0.0      0   0.0   0.0
                          D1            0.2    401       0.0      0   0.0   0.0
                          D2            0.2    378       0.0      0   0.0   0.0
ISGLOCK          LOCK     *ALL         14.5    266       0.0      0   0.0   0.0
                          D0            1.1    235       0.0      0   0.0   0.0
                          D1            7.2    293       0.0      0   0.0   0.0
                          D2            6.1    241       0.0      0   0.0   0.0
ISTGENERIC       LIST     *ALL          1.7    363      <0.1   8393  50.0   150
                          D0            0.8    382       0.0      0   0.0   0.0
                          D1            0.4    353      <0.1   8393  50.0   150
                          D2            0.5    341       0.0      0   0.0   0.0
IXCPLEX_PATH1    LIST     *ALL          0.0      0     248.6   2553   0.0  26.0
                          D0            0.0      0     129.4   2107   0.0  39.2
                          D1            0.0      0      26.6   3981   0.0  34.0
```

Figure 2-24. CFACT Report

There exists no graphic form of this report.

If the cursor is placed on any of the lines with coupling facility structure values, a pop-up panel appears showing details for this structure. Some of these detail values are available only, if the Monitor III Gatherer is running with the gathering option **CFDETAIL**.

```
                    RMF V1R2  CF Activity      - UTCPLXHD      Line 13 of 76
Command ===>                                                  Scroll ===> HALF

Samples: 120     Systems: 3    Date: 11/29/02  Time: 11.27.00  Range: 120   Sec

CF:  ┌──────────────────────────────────────────────────────────────────────┐
     │                   RMF Coupling Facility - Structure Details          │
Str  │                                                                      │
     │  List Structure          : ISTGENERIC                                │
IRR  │  Coupling Facility       : CF5B                                      │
     │  System                  : D1                                        │
     │                                                                      │
     │  Structure Size          :   10M   Connection Name : USIBMT6_M91HOST │
IRR  │  List entries    Total   : 55489   Jobname         : VTAM44          │
     │                  Current :     9   Status          : ACTIVE          │
     │  Data Elements   Total   :  1109   ASID            :   27            │
     │                  Current :     1   CF Level        :    5            │
ISG  │  Lock Entries    Total   :     4                                      │
     │                  Current :     0                                      │
     │  Contention      (%)     :   0.0                                     │
     │                                                                      │
IST  │  Press Enter to return to the Report panel.                          │
     │  If data is missing, see Help panel.                                 │
     │                                                                      │
     │                                                                      │
IXC  │                                                                      │
     └──────────────────────────────────────────────────────────────────────┘
                          D1            0.0      0      26.6   3981   0.0  34.0
```

Figure 2-25. CFACT Report - Details for a List Structure

```
                    RMF V1R2  CF Activity      - UTCPLXHD      Line 13 of 76
Command ===>                                                  Scroll ===> HALF

Samples: 120     Systems: 3    Date: 11/29/02  Time: 11.27.00  Range: 120   Sec

CF:  ┌──────────────────────────────────────────────────────────────────────┐
     │                   RMF Coupling Facility - Structure Details          │
Str  │                                                                      │
     │  Lock Structure          : ISGLOCK                                   │
IRR  │  Coupling Facility       : CF5B                                      │
     │  System                  : D1                                        │
     │                                                                      │
     │  Structure Size          :   10M   Connection Name: ISGLOCK#D1       │
IRR  │  List entries    Total   :     0   Jobname        : GRS              │
     │                  Current :     0   Status         : ACTIVE           │
     │  Lock Entries    Total   : 1049K   ASID           :    7             │
     │                  Current :  1205   CF Level       :    0             │
ISG  │  Contention          (%) :   3.7                                     │
     │  False Contention    (%) :   0.7                                     │
     │                                                                      │
     │  Press Enter to return to the Report panel.                          │
IST  │  If data is missing, see Help panel.                                 │
     │                                                                      │
     │                                                                      │
IXC  │                                                                      │
     └──────────────────────────────────────────────────────────────────────┘
                          D1            0.0      0      26.6   3981   0.0  34.0
```

Figure 2-26. CFACT Report - Details for a Lock Structure

```
                    RMF V1R2  CF Activity      - UTCPLXHD      Line 13 of 76
Command ===>                                                  Scroll ===> HALF

Samples: 120     Systems: 3    Date: 11/29/02  Time: 11.27.00  Range: 120   Sec

CF:
                      RMF Coupling Facility - Structure Details
Str
      Cache Structure         : IRRXCF00_P001
IRR   Coupling Facility       : CF5B
      System                  : D0

      Structure Size          : 1792K   Connection Name : IRRP001@D0
IRR   Direct. Entries Total   :   418   Jobname         : RACFDS
                    Current   :   230   Status          : ACTIVE
      Data Elements   Total   :   407   ASID            :   46
                    Current   :   230   CF Level        :    0
ISG   Request Rate            :   0.2
      Read Rate               :   0.4
      Write Rate              :   0.4
      Castout Rate            :   0.0
IST   XI Rate                 :   0.3
      Directory Reclaims      :     0

      Press Enter to return to the Report panel.
IXC   If data is missing, see Help panel.

                          D1            0.0      0      26.6   3981   0.0  34.0
```

Figure 2-27. CFACT Report - Details for a Cache Structure

Field descriptions

Table 2-9. Fields in the CFACT Report

Field Heading	Meaning
Note: 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	Status - can be one of three states in the SYSPLEX (*ALL) view data line for a structure: **blank** Structure is allocated and connected to at least one system during the entire mintime (normal case). **A** Active - structure became allocated and connected to at least one system during the mintime, or the structure has been added to the report because CFDETAIL has been switch on during this mintime. **I** Inactive - structure got disconnected from all systems during the mintime, or the structure has been removed from the report because CFDETAIL has been switch off during this mintime. **Note:** 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. Therefore, all values for these structures are reported as blank.

Table 2-9. Fields in the CFACT Report (continued)

Field Heading	Meaning
System	System name for the system connected to the structure (from IEASYSxx Parmlib member, SYSNAME parameter). 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.
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 all requests for this structure that changed from synchronous to asynchronous because the requests could not be serviced as synchronous request.
Async Del %	Percentage of all 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.
Note: The availability of the data in the pop-up panel depends on the **CFDETAIL** option of the Monitor III Gatherer session. If this option is not active, all values marked as follows have to be used carefully: **DET** Value will not be reported **MON I** 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. Each rate is reported as '<0.1' if the value is greater than 0 but would be rounded to 0.	
Structure Size (MON I)	Amount of storage that is currently allocated for this structure.
List Entries (MON I)	Number of list entries in use in a LIST or LOCK structure. **Total** Total number. **Current** Number of list entries in use.
Data Elements (MON I)	Number of data elements in use in a LIST or CACHE structure. **Total** Total number. **Current** Number of list data elements in use.
Lock Entries (MON I)	Number of lock table entries in use in a serialized LIST or a LOCK structure. **Total** Total number. **Current** Number of lock table entries in use.
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.

Table 2-9. Fields in the CFACT Report (continued)

Field Heading	Meaning
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. **Note:** 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 (DET)	Name of the last connection from the selected system.
Jobname (DET)	Name of the job that made the last connect from the selected system.
Status (DET)	The status of the last connection from the selected system. **Active** Connection established. **FailPers** Failed Persistent: Connection with CONDISP=KEEP has failed and all of the event exit responses have been received with RELEASECONN=NO. **Failing** Connection terminated abnormally and not all of the event exit responses have been received. **Disc** Disconnecting: Connection disconnected and not all of the event exit responses have been received. **NotKnown** None of the above.
ASID (DET)	ASID of the job that made the last connect from the selected system.
CF Level (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. **Total** Total number. **Current** Number of currently filled directory entries.
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.

Table 2-9. Fields in the CFACT Report (continued)

Field Heading	Meaning
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 2-28. 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

How to request this report

To request the Coupling Facility Overview report, select an **S** from the Primary menu and then a **5** on the Sysplex Report menu (shown in Figure 2-6 on page 2-21), or you can enter the following command:

```
CFOVER [cfname]
```

Purpose of the 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 as balanced as you want. This, of course, will not be the case if you have one coupling facility for production and the other one just as stand-by.

Contents of the report

```
                  RMF V1R2  CF Overview      - UTCPLXHD          Line 1 of 2
Command ===>                                                  Scroll ===> HALF

Samples: 120     Systems: 3    Date: 11/29/02  Time: 11.27.00  Range: 120   Sec

---- Coupling Facility -----    ---- Processor -----   Request    -- Storage --
Name      Type   Model Level    Util% Defined Effect   Rate       Size    Avail

CF5B      9674     RX4     5    12.4      3     3.0      232.3    1008M    679M
CF6B      9674     RX4     5    44.2      3     3.0       4572    1008M    651M
```

Figure 2-29. CFOVER Report

There exists no graphic form of this report.

Field descriptions

Table 2-10. Fields in the CFOVER Report

Field Heading	Meaning
Name	Coupling facility name.
Type	Coupling facility processor type.
Model	Coupling facility processor model.
Level	Coupling facility microcode level.

Table 2-10. Fields in the CFOVER Report (continued)

Field Heading	Meaning
Processor Util%	Average value of processor utilizations within the coupling facility. 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). If the average utilization is high, you can take the following actions: • In a PR/SM environment, you can dedicate the CP to the integrated coupling facility or assign additional CPs to the partition. • Move structures to a coupling facility with lower utilization. • Consider additional or larger coupling facilities.
Processor Defined	Number of logical processors defined for the coupling facility.
Processor Effect	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 between the LPAR dispatch time (CFCC execute and wait time) and the RMF Mintime length. 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.
Request Rate	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. The value is reported as '<0.1' if the rate is greater than 0 but would be rounded to 0.
Storage Size	The total amount of storage in the coupling facility, including both allocated and available space.
Storage Avail	The amount of coupling facility space that is not allocated to any structure and not allocated as dump space.

CFSYS - Coupling Facility Systems report

How to request this report

To request the Coupling Facility Systems report, select an **S** from the Primary menu and then a **6** on the Sysplex Report menu (shown in Figure 2-6 on page 2-21), or you can enter the following command:

```
CFSYS [cfname]
```

Purpose of the report

The Coupling Facility 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, you see for each coupling facility all systems connected to it, and their activities. High 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.

Contents of the report

```
                  RMF V1R2  CF Systems       - UTCPLXHD          Line 1 of 6
Command ===>                                                  Scroll ===> HALF

Samples: 120     Systems: 3    Date: 11/29/02  Time: 11.27.00  Range: 120   Sec

CF Name   System      Subch   -- Paths --   -- Sync ---   ------- Async -------
                      Delay   Avail Delay    Rate  Avg     Rate  Avg  Chng  Del
                        %             %           Serv           Serv   %    %

CF5B      D0            0.0     2     2.6     5.1   344    91.0  2890  0.1  7.6
          D1            0.2     2     7.3    10.6   315    33.8  4903  0.4 22.1
          D2            0.0     2     3.3     9.3   300    82.2  3042  0.0  7.2
CF6B      D0            5.1     1     2.2    2131   242   252.2  1974 31.0 39.3
          D1            1.1     1    34.1     0.9   548    30.9  3884  0.0 42.6
          D2            0.0     2     0.7    1895   267   144.8  2865  2.8  7.7
```

Figure 2-30. CFSYS Report

There exists no graphic form of this report.

If the cursor is placed on any of the lines with coupling facility systems values, a pop-up panel appears showing details for the subchannels and path IDs.

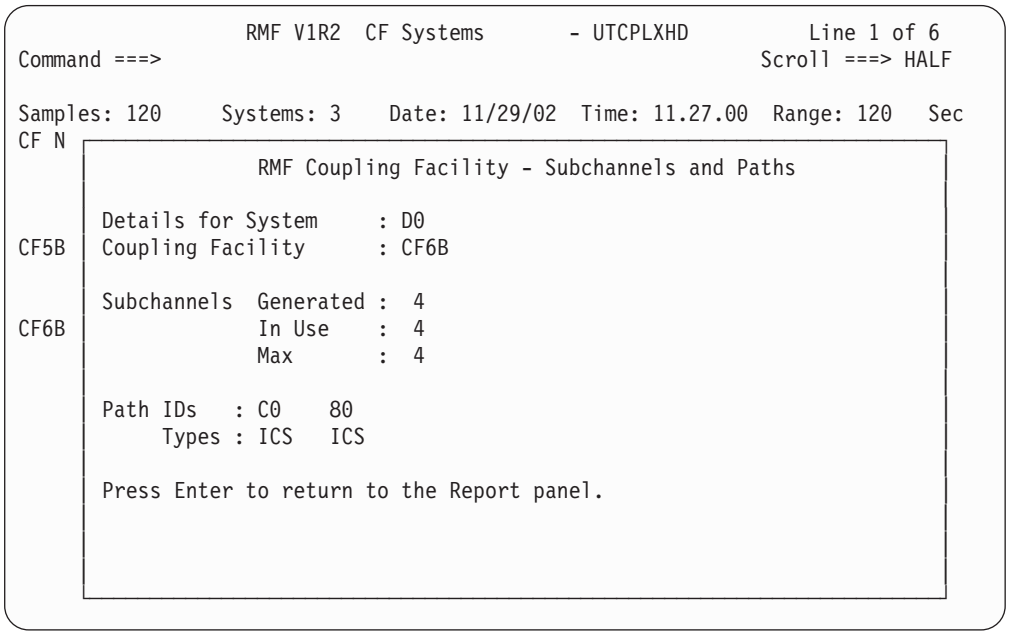

Figure 2-31. CFSYS Report - Subchannel Details

Field descriptions

Table 2-11. Fields in the CFSYS Report

Field Heading	Meaning
Note: 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).
Subch Delay %	The percentage of all coupling facility requests MVS 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.
Paths Avail	Number of physical paths (coupling facility channels) available to transfer coupling facility requests between this system and the coupling facility.

Table 2-11. Fields in the CFSYS Report (continued)

Field Heading	Meaning
Paths Delay %	The percentage of all coupling facility requests that were rejected because all paths to the coupling facility were busy. 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. **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, Subchannel Busy, not Path Busy, is the indicator for path contention. If Path Busy is low but Subchannel Busy 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 is defined in the I/O gen (see Subchannel Max). **PR/SM environment only:** If coupling facility channels are being shared among PR/SM partitions, Path Busy behaves differently. Potentially, you 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. **Possible actions:** Dedicate the coupling facility links on the sending processor or add additional links.
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 %	Percentage of all requests on behalf of connectors from this system that changed from synchronous to asynchronous because the requests could not be serviced as synchronous request.
Async Del %	Percentage of all hardware operations on behalf of connectors from this system being delayed by either subchannel contention or dump serialization.

Mon III - CFSYS report

Table 2-12. Fields in the CFSYS Report - Subchannel Details

Field Heading	Meaning
Subchannels	Subchannel configuration data. **Generated** Number of subchannels defined in the I/O gen. This could be more than the number MVS can optimally use for coupling facility requests. See the MAX field below. **In Use** Number of subchannels MVS is currently using for coupling facility requests. **Max** Maximum number of coupling facility subchannels MVS can optimally use for coupling facility requests. The limit is calculated by MVS 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. If this number is less than the Subchannel Gen, you should reduce the number of coupling facility subchannels in the I/O gen to match this number. Over-specifying subchannels causes unnecessary storage usage and, in a heavily loaded system, can cause a high number of rejected coupling facility requests due to path busy.
Path IDs	The IDs of the channels connecting a system with the coupling facility. The physical path utilization for these coupling facility links is shown in the CHANNEL report.
Types	Subchannel types for the channel paths with the displayed Path IDs.

CHANNEL - Channel Path Activity report

How to request this report

To request the Channel Path Activity report, select a **3** from the Primary menu and then a **12** on the Resource Report menu (shown in Figure 2-9 on page 2-24), or you can enter the following command:

```
CHANNEL
```

Purpose of the 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.

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.

Contents of the report

```
                    RMF V1R2  Channel Path Activity           Line 1 of 12
Command ===>                                                 Scroll ===> HALF

Samples: 100     System: SYS1  Date: 11/29/02  Time: 15.55.00  Range: 120   Sec


 Channel Path       Utilization(%)   Read(B/s) Write(B/s)    MSG  MSG Send Recv
ID No  G  Type  S   Part  Tot  Bus   Part  Tot  Part  Tot   Rate Size Fail Fail

    8    *FC_SM     21.7 38.9 40.2    10M  24M   11M  18M
   12    *FCV_M     13.6 26.8 33.3     9M  16M    7M  12M
    1  1 *FCV_M     12.4 31.9 43.3     8M  22M    7M  19M
    1    *CNC_M     17.2 34.4
80        FC    Y   22.0 43.9 44.0    11M  24M   11M  20M
81        FC_SM Y   21.7 38.9 40.2    10M  24M   11M  18M
82        FCV   Y   14.7 28.5 31.1     8M  17M    6M  11M
83        FCV_M Y   13.6 26.8 33.3     9M  16M    7M  12M
85        CTC_S Y    0.1  3.9
89        CNC_S      0.7  5.8
8C        CNC   Y    1.7 13.5
8D        CNC_S Y    1.8 13.7
DA        CNC_M Y   17.2 34.4
AE        IQD   Y                              705M   4G   923K  764    0    0
AF        IQD   Y                                1G   1G     1M  838    0    0
B0     1  FCV_M Y   12.4 31.9 43.3     8M  22M    7M  19M
```

Figure 2-32. CHANNEL Report

This is an example of the Channel Path Activity report in LPAR mode. The graphic form of this report shows the percent of total utilization for each channel.

Field descriptions

Table 2-13. Fields in the CHANNEL Report

Field Heading	Meaning
Channel Path ID	Hexadecimal number of the 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 * preceeding the channel path type, and the number of channels of the group is displayed in column No.
Channel Path G	Generation. 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, ...). For example, for a FICON channel, a number 1 indicates that the channel has an auto-negotiated throughput of 1Gbit/sec, or a number 2 indicates a throughput of 2Gbit/sec.
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, this field is blank. RMF continues to measure channel path activity. Check the operator console for messages.

Table 2-13. Fields in the CHANNEL Report (continued)

Field Heading	Meaning
Channel Path 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.
Note: 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).	

Utilization (%) Part	The channel path utilization percentage for an individual partition. RMF uses the values provided by CPMF (Channel Path Measurement Facility).
	`Part Channel Path Busy Time`
	`Utilization (%) = ------------------------- * 100`
	`Channel Path Elapsed Time`
	In BASIC mode or if EMIF is not installed, blanks are shown.
	For OSAEGbE, the value reflects the microprocessor utilization.

Utilization (%) Tot	The channel path utilization percentage for the entire system during an interval.
	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:
	`# SRM Observations of`
	`Total Channel Path Busy`
	`Utilization (%) = ---------- * 100`
	`# SRM samples`
	For unshared channels in LPAR mode, the value for total utilization is the same as partition utilization.
	For all channels in BASIC mode with CPMF available, the calculation is:
	`Total Channel Path Busy Time`
	`Utilization (%) = ------------------------- * 100`
	`Channel Path Elapsed Time`
	For OSAEGbE, the value reflects the microprocessor utilization.

Utilization (%) Bus	Percentage of bus cycles, the bus has been found busy for this channel in relation to the theoretical limit. For OSAEGbE, the value reflects the PCI bus utilization.
Read(B/s)	**Part** Data transfer rates from the control unit to the channel for this partition. **Total** Data transfer rates from the control unit to the channel for the entire system.
Write(B/s)	**Part** Data transfer rates from the channel to the control unit for this partition. **Total** Data transfer rates from the channel to the control unit for the entire system.
Note: The following fields are available only for HiperSockets.	
MSG Rate	Rate of messages sent by this partition.
MSG Size	Average size of messages sent by this partition (in bytes).
Send Fail	Rate of messages (sent by this partition) that failed, but not caused by an unavailable buffer in the receiving partition.

Table 2-13. Fields in the CHANNEL Report (continued)

Field Heading	Meaning
Recv Fail	Rate of messages (received by this partition) that failed due to an unavailable buffer. The value could indicate that more receive buffers are required.

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 2-14. Additional Fields in the CHANNEL Report

Field Name	Meaning
CHACTMVC	Total message sent rate
CHACTSVC	Total message sent size
CHACTFVC	Total receive fail rate

CPC - CPC Capacity report

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 2-7 on page 2-22), or you can enter the following command:

```
CPC
```

Purpose of the 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 capactity 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 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.

Note: The CPC report is available only if the system is running on a zSeries processor in a PR/SM environment.

Contents of the report

The report consists of two sections:

- **Header information** which offers MSU related values with the scope of that partition which requested the report as well as the processor type and model
- **Partition data** which displays the values for all partitions belonging to the CPC.

```
                    RMF V1R2 CPC Capacity                        Line 1 of 12
Command ===>                                                  Scroll ===> HALF
Samples: 59      System: Z201  Date: 11/29/02  Time: 13.16.00  Range: 60    sec

Partition: Z2           2064 Model 114
CPC Capacity:     410   Weight % of Max: 74.2    4h MSU Average:    41
Image Capacity:    50   WLM Capping %:    5.4    4h MSU Maximum:   185

Partition  --- MSU ---  Cap  Proc    Logical Util %   - Physical Util % -
             Def   Act  Def   Num    Effect   Total   LPAR  Effect  Total
 *CP                                                  1.0    10.1   11.1
  TZ1        130   122   NO   4.3      11.2    12.5    0.4     3.4    3.8
  Z1         150    89   NO   5.2       9.3     9.6    0.1     3.5    3.6
  Z2          50    58   NO   2.1      11.5    12.8    0.2     1.7    1.9
  Z3         N/A     0  YES   2.4       8.8    10.6    0.3     1.5    1.8
  PHYSICAL                                             0.1            0.1

 *ICF                                                  0.1     7.1    7.1
  CF1          0    29          1      99.9    99.9    0.0     7.1    7.1
  CF2          0     0          1       0.0     0.0    0.0     0.0    0.0
  PHYSICAL                                             0.1            0.1
```

Figure 2-33. CPC Report

Field descriptions

Table 2-15. Fields in the CPC Report

Field Heading	Meaning
All MSU-related values are measured in MSU/hour.	
Values for the partition which requested the report	
Partition	Partition name.
Processor/Model	Processor family and model number of the measured system.
CPC Capacity	Processor capacity available to the CPC.
Image Capacity	Processor capacity available to the MVS image measured in MSUs (millions of service units) per hour. The field is calculated as follows: • If the partition has a defined capacity limit, this value is that limit. • If the partition has no defined capacity limit, this value is calculated based on the partition's logical CP configuration and capping state.
Weight % of Max	Average weighting factor in relation to the maximum defined weighting factor for this partition. If the maximum weight value is changed on the Hardware Management Console (HMC), this field shows '****' during this interval.
WLM Capping %	Percentage of time when WLM capped the partition because the four-hours average MSU value exceeds the defined capacity limit.
4h MSU Average	Average value of consumed MSU/hour during the last four hours.
4h MSU Maximum	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.

Table 2-15. Fields in the CPC Report (continued)

Field Heading	Meaning
Values for all partitions belonging to the CPC	
Partition	Partition name. **Notes:** 1. The partition identified by the name PHYSICAL is not a configured partition. Data reported in these lines includes all of the uncaptured time which was used by LPAR but could not be attributed to a specific logical partition. 2. The values identified by the names *CP and *ICF show the total percentages.
MSU	**Def** Defined MSU capacity limit for the partition. **Act** Actual MSU consumption.
Cap Def	The capping option of the partition: YES or NO. This field indicates whether the operator has set 'Initial Capping ON' in the logical partition controls on the Hardware Management Console (HMC) for the partition.
Proc Num	Number of logical processors assigned to this partition.
Average Processor Utilization Percentages. • 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.	
Logical Util % (Effect)	The average partition effective dispatch time percentage. $\frac{\text{Effective Dispatch Time}}{\sum \text{Online Times}} * 100$
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. $\frac{\text{Total Dispatch Time} - \text{Effective Dispatch Time}}{\#\ \text{Physical Processors} * \text{Range Time}} * 100$
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. $\frac{\text{Total Dispatch Time}}{\#\ \text{Physical Processors} * \text{Range Time}} * 100$

DELAY - Delay report

How to request this report

To request the Delay report, select a **1** from the Primary menu, and then **4** on the Overview Report menu (shown in Figure 2-7 on page 2-22) or you can enter the following command using the format:

```
DELAY [class,s-class]
```

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

```
DELAY T, TSOPRIME
```

Purpose of the 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.

Contents of the report

```
                        RMF V1R2  Delay Report                Line 1 of 58
Command ===>                                                 Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  Time: 12.00.00  Range: 100 Sec

            Service     WFL USG DLY IDL UKN ----- % Delayed for -----  Primary
Name     CX Class    Cr  %   %   %   %   %  PRC  DEV STR SUB OPR  ENQ  Reason

*SYSTEM                  1   1   3  82  14   0    1   1   0   0    0
*TSO                    50   1   1  97   1   0    0   0   0   0    0
*BATCH                  10  27  75   0   6  13   15   5  16   9   19
*STC                     3   2   6  23  70   0    1   4   0   1    0
*ASCH                        0   0   0 100   0    0   0   0   0    0
*OMVS                        0   0   9  91   0    0   0   0   0    0
*ENCLAVE                 5   5  87   0   8  87  N/A   0 N/A N/A  N/A
MISTYDFS B  NRPRIME      0   0 100   0   0   0    0   0   0 100    0  Message
SWARRENA BO NRPRIME  S   0   0 100   0   0   0    0   0   0   0  100  SYSDSN
CAROLYN  T  TSOPRIME C   0   0 100   0   0   0    0 100   0   0    0  XMEM
MVSNFS   S  SYSSTC   SC  0   0 100   0   0   0    0 100   0   0    0  XMEM
SPEWAK2  B  NRPRIME      0   0 100   0   0   0    0   0   0   0  100  SYSDSN
STYMIE   TO TSOPRIME     0   0 100   0   0   0    0 100   0   0    0  XMEM
ENC00005 E  SYSSTC       4   2  43   0   0  43  N/A   0 N/A N/A  N/A  N/A
```

Figure 2-34. 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 2-16. 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: **S** Started task **T** TSO **B** Batch **A** ASCH **O** OMVS **?** Data is missing or not valid. Or it can indicate an enclave: **E** Enclave For summary entries, this field is blank. An **O** 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 (--------). **Note:** For systems running in compatibility mode, the fields `DMN` and `PG` are shown instead of `Service Class`, they display the domain and performance group for the specified job.
Cr	An indication whether WLM managed the address space as *storage critical* and/or *CPU critical* during the reporting interval. **C** CPU critical **S** Storage critical **SC** Both storage and CPU critical
WFL %	The workflow percentage of the job or job group. "Address space workflow (%)" on page 2-7 shows the formula used to calculate this value.
USG %	The using percentage for the job or job group. "Address space using (%)" on page 2-8 shows the formula used to calculate this value.
DLY %	The delay percentage for the job or job group. See "Address space delay (%)" on page 2-9 for more information.

Table 2-16. Fields in the DELAY Report (continued)

Field Heading	Meaning

IDL %	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 *idling* state if they are not using the processor or devices and are not delayed for any monitored reason.
	
	Jobs classified as in terminal wait meet all of the following conditions:
	• They are not found using any monitored resource
	• They are not found delayed for any monitored reason
	• They are swapped out
	• They are in terminal wait
	• They are waiting for a user ready indication before being swapped in.
	
	Jobs classified as in timer wait meet all of the following conditions:
	• They are not using or delayed for a monitored resource.
	• Their address space is waiting for a timer.
	
	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.
	
	The idling percentage for an address space during a refresh period is calculated as follows:
	`IDL % = (# Idle Samples / # Samples) * 100`
	
	Idle samples
	The number of samples that show the job in an idle state.
	
	The idling percentage for a group of address spaces during a range period is calculated as follows:
	`IDL % = (∑ Idle Samples / (# Samples * Avg # Address Spaces)) * 100`
	
	Note: The value reported might include some delay for a non-monitored resource.

UKN %	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. Examples of address spaces in an unknown state are: • 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. • Address spaces waiting for devices other than DASD or tape. 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.

Table 2-16. Fields in the DELAY Report (continued)

Field Heading	Meaning
% Delayed For	The percentage that each defined resource contributes to the overall delay of the job or job group. 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. If the percentages add up to more than DLY %, there is an overlap of delay states. The defined resources that can delay the job or job group are as follows: **PRC** The job or job group has ready work on the dispatching queue, but it is not being dispatched. **DEV** The job or job group is delayed for a DASD or tape. **STR** 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. **SUB** The job or job group is delayed for a JES, HSM, or XCF subsystem request. **OPR** 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: • A swappable address space will be swapped out, thus it can be OUTR and show storage delays. • 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. **ENQ** The job or job group is waiting to use an enqueued (reserved) resource.

Table 2-16. Fields in the DELAY Report (continued)

Field Heading	Meaning
Primary Reason	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. If the resource with the maximum delay is: **PRC** This field contains the name of the job that used the processor most often while the reported job was delayed. **STR** This field contains the storage delay category that had the most delay (COMM, LOCL, SWAP, VIO, XMEM, HIPR, OUTR). **DEV** This field contains the volume serial number of the device that the reported job was most often delayed for. **SUB** This field contains either JES, HSM, or XCF depending on which subsystem is causing the most delay. **OPR** This field contains *Message* if most of the delay was due to a message or *Mount* if most of the delay was due to a mount request. The field can contain *QUIESCE* if the operator quiesced the address space. A quiesced address space can show unexpected data: • A swappable address space will be swapped out, thus it can be OUTR and show storage delays. • 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. Cursor-sensitive control on this field gives you the Quiesce variation of the Job report. **ENQ** This field contains the major name of the resource most responsible for the delay. ***ENCLAVE** One or more enclaves are holding the processor. **RG-Cap** The job is delayed due to WLM resource capping. That means that • either the resource group for which the job is running, has used up its CPU service specified in the WLM policy, • 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 *z/OS MVS Planning: Workload Management*).

Monitor III Utility fields

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

Table 2-17. Additional Fields in the DELAY Report

Field Heading	Meaning
% Delayed For	• JES delay percentage • HSM delay percentage • XCF delay percentage • Operator mount delay percentage • Operator message delay percentage • Operator quiesce delay percentage • WLM resource group capping delay percentage

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 Figure 2-35 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 report.

```
                     RMF V1R2  Delay Report                  Line 1 of 58
Command ===>                                                Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  Time: 12.00.00  Range: 100   Sec


Name     CX  ┌──────────────────────────────────────────────────────────────┐
             │                 RMF Enclave Classification Data              │
*SYSTEM      │ The following details are available for enclave ENC00002 :   │
*TSO         │ Press Enter to return to the Report panel.                   │
*BATCH       │                                                              │
*STC         │   Transaction program name : WLMTOOL                         │
*ASCH        │   Userid                   : BHIM                            │
*OMVS        │   Transaction class        : JES2                            │
*ENCLAVE     │   Netword id               :                                 │
PRIMASID S   │   Logical unit name        :                                 │
ENC00007 E   │   Plan                     : TEST                            │
ENC00009 E   │   Package                  : UPS                             │
CATALOG  S   │   Connection               :                                 │
MAINASID S   │   Collection               : COLLECTION                      │
HOMEASID S   │   Correlation              : CTT                             │
VLF      S   │   Subsystem type           : DDF                             │
SYSBMAS  S   │   Function name            : FUNCTION                        │
ENC00004 E   │   Subsystem name           : DDF                             │
ENC00005 E   │                                                              │
ENC00001 E   └──────────────────────────────────────────────────────────────┘
MVSNFS   S  SYSSTC   SC   0   0 100   0   0   0    0 100   0   0    0  XMEM
SPEWAK2  B  NRPRIME       0   0 100   0   0   0    0   0   0   0  100  SYSDSN
STYMIE   TO TSOPRIME      0   0 100   0   0   0    0 100   0   0    0  XMEM
ENC00005 E  SYSSTC        4   2  43   0   0  43  N/A   0 N/A N/A  N/A  N/A
```

Figure 2-35. DELAY Report - Details for Enclave Classification Data

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, 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 DELAY report (YES NO)
  Criterion     ===> 0         Minimum delay to include job in DELAY report

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

                          Available Service classes
APPPRIME   NRPRIME    OMVS       TSOPRIME   SYSTEM     SYSSTC
```

Figure 2-36. 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.

Note: For systems running in compatibility mode, the field Group is shown instead of Service Class to let you specify the domain or performance group for which you want data reported. For Group, you can specify any of the available groups listed under Available Groups.

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.

Note: For systems running in compatibility mode, the list of available groups (domains and performance groups) is shown instead.

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 2-37.

```
                    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 2-37. 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

How to request this report

To request the Device Delays report, select a **3** from the Primary menu and then a **2** on the Resource Report menu (shown in Figure 2-9 on page 2-24), or you can enter the following command using the format:

```
DEV [class,s-class]
```

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

```
DEV T, TSOPRIME
```

Purpose of the 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.

Contents of the report

```
                        RMF V1R2  Device Delays                 Line 1 of 57
Command ===>                                                  Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  Time: 10.03.20  Range: 100   Sec

           Service  DLY USG CON  ------------ Main Delay Volume(s) ------------
Jobname  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
RSTSHYS0 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 TSO063    1 HSM004
```

Figure 2-38. DEV Report

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

Field descriptions

Table 2-18. 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: **S** Started task **T** TSO **B** Batch **A** ASCH **O** OMVS
Service Class	The name of the service class that a specified job has been running in. **Note:** For systems running in compatibility mode, the fields `DMN` and `PG` are shown instead of `Service Class`, they display the domain and performance group for the specified job.
DLY %	Delay the waiting job (address space) is experiencing because of contention for devices during the report interval, expressed as a percentage. **Note:** 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.) **Note:** To find all the using volumes for a jobname you must scan an entire resource-oriented device delay (DEVR) report, using the FIND command.
CON %	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. **Notes:** 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. 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.
Main Delay Volume(s)	Up to four DEV volumes contributing most to the delay of the job. The DEV volume having the largest delay percentage appears first. **VOLSER** The serial number of a DASD or tape contributing to the delay of a job. **%** The percentage of delay caused because the job was waiting to use the named volume.

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, STOR, STORC, STORF,
and XCF.

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

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

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

Figure 2-39. 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 2-68 (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

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
```

Purpose of the 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 page 2-75).

Contents of the report

```
                  RMF V1R2  Device Activity                    Line 1 of 118
Command ===>                                                  Scroll ===> HALF

Samples: 100     System: AQTS  Date: 11/29/02  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
TSO060 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
TSO024 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 2-40. 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 5-59).

  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 2-21 on page 2-78) except for the following:

Table 2-19. 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: `IosQ = IOS Queue Count / # Samples / 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	**USG** The average number of jobs using the device during the report interval. **DEL** 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 2-20. 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
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

How to request this report

To request the DEVR report, select a **3** from the Primary menu, and then a **3** on the Resource Report menu (shown in Figure 2-9 on page 2-24) or you can enter the following command using the format:

```
DEVR [volser]
```

Purpose of the 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

Contents of the report

```
                    RMF V1R2  Device Resource Delays        Line 1 of 374
Command ===>                                              Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  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     %   %

TSOL11 S    11.7 .085  100  30  69 PND   1 33903    MARYPATM B NRPRIME   30  70
  0194                                    3990-3
BPXLK1 S    29.0 .028   62  17  42 PND   3 33903    MICHAELL B NRPRIME   10  39
  0409                                 DB    1 3990-3   CHARLESR B NRPRIME    8  28
D24PK2 S    71.1 .008   52  15   6 PND  31 33903    MCPDUMP  S SYSSTC    12  36
  0307                                 DB   29 3990-3   SMF      S SYSTEM     1   4
HSML17 S    34.6 .014   18   9   8 PND   1 33903    DFHSM    S SYSSTC     7  10
  0743                                    3990-3   RSTSHYS0 B NRPRIME    1   5
SYSSMU S     1.6 .040    6   5   0 PND   1 33903    RHANSON  T TSOPRIME   0   4
  06AE                                 CMR   1 3990-3   CRISMAN  T TSOPRIME   0   4
                                                       GINNI    T TSOPRIME   0   3
                                                       JACKF    T TSOPRIME   1   2
                                                       KOCH     T TSOPRIME   0   2
```

Figure 2-41. 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 2-21. 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	**S** An S in the first line of this column indicates that the device was generated during system generation as a shared device. **PAV** 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 interval. If the number has been changed during the reporting interval, it is followed by an '*'.

Act Rate	The rate per second that I/O instructions (SSCH, RSCH, and HSCH) to a device completed successfully.
	The calculation is:
	`# I/O Instructions`
	`Act Rate = ------------------`
	`Range Time`

Resp Time	The average response time (in seconds) that the device required to complete an I/O request.
	The calculation is:
	`Active Time`
	`Resp Time = ------------------ + IOS Queue Time`
	`# I/O Instructions`

ACT %	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: `ACT % = PEND % + CON % + DSC %` **PEND %** 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). **CON %** Percentage of time the device was connected to a channel path to actually transfer data between the device and storage. **DSC %** 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.

CON %	The percent connect time. See the description under % ACT. RMF calculates the value as follows:
	`Accumulated Connect Time`
	`CON % = ----------------------- * 100`
	`Range Time`

Table 2-21. Fields in the DEVR Report (continued)

Field Heading	Meaning

DSC %	The percent disconnect time. See the description under %ACT. RMF calculates the value as follows:
	`Accumulated Disconnect Time`
	`DSC % = -------------------------- * 100`
	`Range Time`

Notes: 1. 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. 2. 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.	

PND % Reasons	The first entry is always the pending percentage (PND). See the description under % ACT. RMF calculates the value as follows:
	`Accumulated Pending Time`
	`PND % = ----------------------- * 100`
	`Range Time`
	
	DLY DB % and DLY CU % are included in pending time.
	
	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:
	
	DB 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.
	`Accumulated DB Delay Time`
	`DLY DB% = ------------------------  * 100`
	`Range Time`
	
	CMR 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.
	`Accumulated CMR Delay Time`
	`DLY CMR% = ------------------------- * 100`
	`Range Time`
	Note: If either hardware data or volume-related percentages are not available, this field is blank.

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.

Table 2-21. Fields in the DEVR Report (continued)

Field Heading	Meaning
C	A one-character abbreviation for the job class as follows: **S** Started task **T** TSO **B** Batch **A** ASCH **O** OMVS
Service Class	The name of the service class that a specified job has been running in. **Note:** For systems running in compatibility mode, the fields DMN and PG are shown instead of Service Class, they display the domain and performance group for the specified job.
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. This value is the amount of time the device was connected to a channel path to actually transfer data.
DLY %	Delay the waiting job (address space) is experiencing because of contention for a specific Volume during the report interval, expressed as a percentage. This value is the amount of time the job spent disconnected or pending.

Monitor III Utility fields

Table 2-22 shows additional fields for the Device Resource Delay report.

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

Table 2-22. 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
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.

```
                          RMF DEVR Report Options             Line 1 of 17
Command ===>                                                Scroll ===> HALF

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

      VOLSER ===> ALL         Volume to be reported or ALL

                              Available Volumes
BASECA     BASECB     BASEC0     BASEC1     BASEC2     BASEJM     BERDPK
BSS210     BSS999     CAT212     CAT84I     CB8480     CHKPTX     CHKPT1
CHKI01     CHKPT3     CHKPT4     CKSPL1     CKSPL2     CKSPL3     CLRPAK
C45I01     DASD03     DBLMN1     DBLMN2     DBLMN3     DBVL01     DBVL02
D60AF3     D60AF5     D60AF6     D60AF7     D60PAK     D602A2     D602A3
D602A4     D602B1     D602B2     D602B2     D602DC     D602D5     D602D6
D602D7     D6021B     D6021D     D6021D     D60215     D60217     D60411
D60412     D6044D     D60444     D60444     D60445     D60446     D606A8
D607AF     D607EC     D607E1     D60701     D60707     D608FO     D83CL3
D71CKP     D71CMN     D71SPP     D71WLD     D83CL1     D83CL2     D83D31
D83CMN     D83DB1     D83DB2     D83DMP     D83DRA     D83DRB     D83D31
D83D51     D83ILG     D83I80     D83I81     D83I90     D83I91     D83I92
D83I93     D83JCK     D83JC1     D83JC2     D83JD1     D83JD2     D83JFY
D83JS1     D83JS2     D83JVC     D83JV2     D83JV3     D83J22     D83J23
D83J31     D83J32     D83LOG     D83PAK     D83SPP     D83STA     D83STB
```

Figure 2-42. 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

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.

Purpose of the 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

Contents of the report

```
                  RMF V1R2  CLR010 Activity Trend              Line 1 of 20
Command ===>                                                 Scroll ===> HALF

Samples: 100     System: AQTS  Date: 11/29/02  Time: 11.58.20  Range: 100   Sec

    VolSer:   CLR010    Number: 0051  Type and CU-Type: 33903  3990-3
    Latest:   04/23/02 at 11.58.20    Range/Line:      100 Sec
    Earliest: 04/23/02 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 2-43. 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 data about 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 2-23. 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 2-21 on page 2-78 or in Table 2-19 on page 2-75.

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

How to request this report

To request the DSD report, select a **U** on the Primary menu, and then a **2** on the User menu. You can also enter the following command:

```
DSD
```

Purpose of the report

Figure 2-44 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.

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.

Under Working Set, the report lists the amount of central and expanded storage the user occupied while the job was swapped-in.

The COMM and LOCL fields include shared storage paging.

```
                    RMF V1R2  Storage Delays                Line 1 of 206
Command ===>                                                Scroll ===> HALF

Samples:    100   System: MVS1 Date: 11/29/02 Time: 10:31:40 Range:    100 Sec

               Service    DLY     ------------ % Delayed for ------------
Jobname    C   Class       %      COMM  LOCL  SWAP  OUTR   VIO  XMEM  HIPR

*MASTER*   S   STC_HIGH     0        0     0     0     0     0     0     0
PCAUTH     S   STC_HIGH     0        0     0     0     0     0     0     0
RASP       S   STC_LOW      0        0     0     0     0     0     0     0
TRACE      S   STC_LOW      0        0     0     0     0     0     0     0
XCFAS      S   STC_HIGH     0        0     0     0     0     0     0     0
GRS        S   STC_HIGH     0        0     0     0     0     0     0     0
SMXC       S   STC_HIGH     0        0     0     0     0     0     0     0
SYSBMAS    S   STC_HIGH     0        0     0     0     0     0     0     0
DUMPSRV    S   STC_HIGH     0        0     0     0     0     0     0     0
CONSOLE    S   STC_HIGH     0        0     0     0     0     0     0     0
ALLOCAS    S   STC_HIGH     0        0     0     0     0     0     0     0
TLCS       S   STC_LOW      0        0     0     0     0     0     0     0
GPDB       S   STC_LOW      0        0     0     0     0     0     0     0
NETVIEW1   S   STC_HIGH     0        0     0     0     0     0     0     0
TSO        S   STC_HIGH     0        0     0     0     0     0     0     0
APFTABLE   S   STC_HIGH     0        0     0     0     0     0     0     0
```

Figure 2-44. 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 2-55 on page 2-156) except for the XMEM and HIPR fields.

Table 2-24. 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

How to request this report

To request the Data Set Delays report, select a **3** from the Primary menu, and then a **3A** on the Resource Report menu (shown in Figure 2-9 on page 2-24). or you can enter the following command using the format:

```
DSND [dsname]
```

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

Purpose of the 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 reporting interval, or to change the list of data sets to be reported on, enter the command ROPTIONS to display the DSND Report Options panel.

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 V1R2  Data Set Delays                  Line 1 of 6
Command ===>                                                 Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  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 2-45. DSND Report

A graphic version of this report is not available.

Field descriptions

Table 2-25. 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 reporting 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	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. **Note:** See "Special considerations" on page 2-86.
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 reporting 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 reporting 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.IGX00022.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 2-46. 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 MVS 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 reporting interval. The data set names are sorted in alphabetical order.

Notes:

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

How to request this report

To request the Data Set Delays - Job report, select a **2** from the Primary menu, and then a **1A** on the Job Report menu (shown in Figure 2-8 on page 2-23). or you can enter the following command using the format:

```
DSNJ [jobname]
```

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

Purpose of the 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.

Contents of the report

```
                    RMF V1R2  Data Set Delays - Job            Line 1 of 5
Command ===>                                                 Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  Time: 10.03.20  Range: 100   Sec

Jobname: BOECL2              EXCP Rate:    123.5     Connect:  41%

ASID  -------------- Data Set Name ---------------  Volume  Num   DUSG% DDLY%
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.IGX00022.LST90514                       DATA38  0122      1     6
```

Figure 2-47. DSNJ Report

A graphic version of this report is not available.

Field descriptions

Table 2-26. 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. **Note:** 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 **Volume** and **Num**. This happens if only those I/O instructions have been detected for which no data set information is provided by the SMS subsystem, like • I/Os to system data sets (like paging or spooling) • I/Os to any data set which was opened prior to SMS subsystem initialization • I/Os like SENSE or RELEASE • System I/Os not done by an access method
Volume	Name of the volume on which the data set resides which was utilized during the current reporting 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. **Note:** See “Special considerations” on page 2-86.
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

How to request this report

To request the Data Set Delays - Volume report, select a **3** from the Primary menu, and then a **3B** on the Resource Report menu (shown in Figure 2-9 on page 2-24). or you can enter the following command using the format:

```
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.

Purpose of the report

The DSNV report presents information about the utilization of data sets which 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 which 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.

Contents of the report

```
                    RMF V1R2  Data Set Delays - Volume          Line 1 of 10
Command ===>                                                 Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  Time: 10.03.20  Range: 100   Sec

------------------------- Volume EDSS99 Device Data --------------------------
Number:    0B4A        Active:      84%        Pending:   22%     Average Users
Device:    3380E       Connect:     10%        Delay DB:  22%        Delayed
Shared:    Yes         Disconnect:  52%        Delay CM:   0%          0.4

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

Figure 2-48. DSNV Report

A graphic version of this report is not available.

Field descriptions

Table 2-27. 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 report (see "Device Delay variation" on page 2-125).
Data Set Name	Name of a data set which was utilized during the reporting 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. **Note:** 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 • I/Os to system data sets (like paging or spooling) • I/Os to any data set which was opened prior to SMS subsystem initialization • I/Os like SENSE or RELEASE • System I/Os not done by an access method
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. **Note:** See "Special considerations" on page 2-86.
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 2-49. DSNV Report Options Panel

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

Notes:

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

How to request this report

To request the ENCLAVE report, select an **O** on the Primary menu, and then a **6** on the Overview Report menu (shown in Figure 2-7 on page 2-22), or you can enter the following command:

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

Purpose of the 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.

Contents of the report

```
                    RMF V1R2  Enclave Report                  Line 1 of 16
Command ===>                                                 Scroll ===> HALF

Samples: 120     System: SYS5  Date: 11/29/02  Time: 13.13.30  Range: 120   Sec

Current options:   Subsystem Type: ALL                          -- CPU Util --
                   Enclave Owner:                               Appl%   EAppl%
                   Class/Group:                                  18.1     90.1

Enclave   Attribute  CLS/GRP  P Goal    % D X   EAppl%   TCPU    USG  DLY  IDL

*SUMMARY                                          70.04
ENC00003  CTT        PG004    1             O    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 2-50. ENCLAVE Report

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

If you place the cursor on any data field of an enclave, a pop-up panel appears showing performance statistics and classification attributes for the selected enclave.

```
                RMF Enclave Classification Data

The following details are available for enclave ENC00003 :
Press Enter to return to the Report panel.

 Detailed Performance Statistics:

 -- CPU Time --   ------------- Execution States -------------
 Total    26.78   #STS  -Using-  ------ Delay ------  IDL  UNK
 Delta    22.50         CPU I/O  CPU I/O STO CAP QUE
                  592   12 0.0   88 0.0 0.0 0.0 0.0  0.0  0.3

 Classification Attributes:

                                                        More:     +
  Subsystem Type: DDF      Owner: ENCLAVE1    System: RMF3
  Accounting Information     :
   Q123ERF7


  Collection Name            : COLLECTION
  Connection Type            :
  Correlation Information    : CTT
  LU Name                    :
  Netid                      :
  Package Name               : UPS
  Plan Name                  : TEST
  Procedure Name             :
  Process Name               :
  Transaction Class/Job Class: JES2
  Transaction Name/Job Name  : WLMTOOL
  Subsystem Collection Name  : SYS1
  Scheduling Environment     : SYS1
  Userid                     : JOEGEE
  Subsystem Instance         : DDF
  Subsystem Parameter        :
```

Figure 2-51. ENCLAVE Report - Enclave Details

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 at the end of the Monitor III mintime. Therefore, short-running enclaves will not appear in the report.

Field descriptions

Table 2-28. Fields in the ENCLAVE Report

Field Heading	Meaning
Subheader Section - You can define the setting of each field in the 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 or performance group.

Table 2-28. Fields in the ENCLAVE Report (continued)

Field Heading	Meaning
Appl%	The system CPU utilization, not considering enclave TCBs or SRBs.
EAppl%	The total system CPU utilization.
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 Report Options panel.
CLS/GRP	Service class or performance group the enclave is associated with.
P	Service class or performance group period, the enclave is currently running in.
Goal	Response time goal. The field is blank when running in compatibility mode or when another goal has been defined.
%	Response time percentile or velocity. The field is blank when running in compatibility mode or when another goal has been defined.
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: **O** The enclave originated on this system. **F** The enclave originated on another system in the sysplex but is participating on this system. **blank** The enclave is a single-system enclave.
Enclave Performance.	
EAppl%	Percentage of CPU time 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, queueing) based on #STS.
IDL	Percentage of idle samples based on #STS.

Table 2-29. Fields in the ENCLAVE Report - Enclave Details

Field Heading	Meaning
CPU Time Total	Total CPU time (in seconds) consumed by the enclave.
CPU Time Delta	CPU time (in seconds) consumed in the reported Monitor III range.

Table 2-29. Fields in the ENCLAVE Report - Enclave Details (continued)

Field Heading	Meaning
Execution States - In contrast to other Monitor III fields, the execution 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 #STS.	
#STS	Total number of state samples in the enclave.
Using CPU	Percentage of CPU using samples.
Using I/O	Percentage of I/O using samples.
Delay CPU	Percentage of CPU delay samples.
Delay I/O	Percentage of I/O delay samples.
Delay STO	Percentage of storage delay samples. This includes: • Waiting for paging I/O from common • Waiting for cross-memory page fault • Waiting for shared paging • Server private paging delay • Server VIO paging delay • Server hiperspace paging delay • Server MPL delay • Server swap-in delay
Delay CAP	Percentage of CPU capping samples.
Delay QUE	Percentage of queue delay samples.
IDL	Percentage of idle samples.
UNK	Percentage of unknown samples.

Report options

On the Report Options panel, you can select:

- An enclave filter by one of the following criteria:
  - Subsystem type, for example DDF, IWEB, or SOM
  - Owner job name, for example DB2MSTR
  - Class or group, for example PG005
- A list of classification attributes matching the selected subsystem.

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 ===>

Select one of the following options:

 1 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             _  Process Name
 _  Collection Name                    _  Scheduling Environment
 _  Connection Type                    _  Subsystem Collection Name
 S  Correlation Information            _  Subsystem Instance
 _  LU Name                            _  Subsystem Parameter
 _  Netid                              S  Subsystem Type
 _  Package Name                       _  Transaction Class/Job Class
 _  Plan Name                          _  Transaction Name/Job Name
 _  Procedure Name                     S  Userid
```

Figure 2-52. ENCLAVE Report Options Panel

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

ALL selects all active subsystems.

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

Service Class / Performance Group
Report on enclaves only that run in this service class or performance group.

Classification Attributes
You can select one or more classification attributes to be displayed in the report. There, the attributes (maximal eight characters) will be listed in the order at which they are specified in the Report Options panel.

ENQ - Enqueue Delays report

How to request this report

To request the Enqueue Delays report, select a **3** from the Primary menu, and then a **4** on the Resource Report menu (shown in Figure 2-9 on page 2-24), or you can enter the following command using the format:

```
ENQ [class,s-class]
```

Purpose of the 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.

Contents of the report

```
                        RMF V1R2  ENQ Delays                    Line 1 of 14
Command ===>                                                  Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  Time: 10.03.20  Range: 100   Sec

         DLY  -------------- Resource Waiting --------------  ---- Holding ----
Jobname   %    %  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
                                                             100 D71SJH1   SO
SWARRENA 100  100  EW  SYSDSN          (SYS)                  100 AMOLLOY   SO
                      SYS1.NUCLEUS                           100 SCHMATE   SO
                                                             100 DRAGON    SO
                                                             100 SCHUMAC   SO
                                                             100 D71SJH1   SO
LUCKYSM    1    1  EW  SYSZVVDS        (SYS)                    1 STEVEB    EO
+CATALOG                SYS1.MVS3.MCAT                           +CATALOG
*MASTER*   1    1  EW  SYSIKJBC        (SYSS)                   1 *MASTER*  EO
                      PETEG
```

Figure 2-53. ENQ Report

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

Field descriptions

Table 2-30. 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 +.

DLY %	Delay the waiting job is experiencing because of contention for any enqueued resource during the report interval. This value is calculated as follows:
	`Delay samples`
	`DLY % = ---------- * 100`
	`# Samples`
	Delay samples
	The number of samples when the job was delayed for one or more enqueued resources.
	Note: This DLY% value is also found in the ENQ field on the job delay report.

Resource Waiting %	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:
	`Delay samples`
	`Waiting % = ---------- * 100`
	`# Samples`
	Delay samples
	The number of samples when the job was delayed for the resource.
	Note: If there is no overlap in delay states, the WAITING% value(s) for a job add up to the DLY % value of the job.

Resource Waiting STAT	The status indicates whether the waiting job wants exclusive (EW) or shared (SW) use of the resource.
Major/Minor Names	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.

Table 2-30. Fields in the ENQ Report (continued)

Field Heading	Meaning
Major/Minor Names (cont'd)	The following are the most common enqueue major names and their associated resources: **Major Name** — **Resources** **MSFDC** — Service processor TP port **SPFDSN** — Data set name **SPFEDIT** — Data set name **SYSDSN** — System data sets **SYSIAT** — JES3 CHKPNT data set **SYSIEA01** — Dump data set **SYSIEFSD** — Serializes device allocations **SYSIEWLP** — SYSLMOD data set (Minor name is data set name **SYSIGGV1** — Master catalog **SYSIGGV2** — Catalogs (Minor name is catalog name) **SYSIKJBC** — TSO broadcast data set (Minor name is relative block address) **SYSIKJUA** — User attribute data set **SYSSMF01** — SMF SYS1.MANx data set **SYSVSAM** — VSAM data sets **SYSVTOC** — VTOC (Minor name is volser) **SYSAVM** — AVM queue or data areas **SYSZBDT** — MVS bulk data transfer (Minor name is node name **SYSZCAXW** — Catalog auxiliary work area **SYSZCMDS** — Master trace command or Message loss detection **SYSZCOMM** — Global Resource Serialization ring processing table **SYSZCSD** — CSD control block field **SYSZEC16** — Purge data set **SYSZIGGI** — TSB (Minor name is ASID) **SYSZIST0C** — Configuration restart data set (Minor name is ddname) **SYSZJES2** — JES2 buffer or data set **SYSZJWTP** — Job step messages **SYSZOPEN** — System data sets **SYSZPCCB** — Private catalog control block **SYSZPGAD** — PAGEADD command **SYSZPSWD** — Password data set **SYSZRPLW** — Catalog name (Minor name is catalog name) **SYSZSIPS** — SYSEVENT **SYSZSMF1** — SMF buffer **SYSZTIOT** — Device allocation **SYSZTRC** — System trace **SYSZUSRL** — User label tracks **SYSZVARY** — Reconfiguration commands **SYSZVMV** — Volume mount and verify **SYSZVOLS** — System volumes (Minor name is volser) **SYSZWTOR** — WTOR reply (Minor name is REPLYxx, where xx is the message id)

Table 2-30. Fields in the ENQ Report (continued)

Field Heading	Meaning
Holding %	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: `Holding % = Holding samples / # Samples * 100` **Holding samples** 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.
Holding Name/SYS	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

The ENQ Report Options panel is similar to the Device Report Options panel. See Figure 2-39 on page 2-73 for an example of the Report Options panel. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed. See Figure 2-37 on page 2-69 for an example.

ENQR - Enqueue Resource Delays report

How to request this report

To request the ENQR report, select a **3**, and then a **5** on the Resource Report menu (shown in Figure 2-9 on page 2-24), or or you can enter the following command using the format:

```
ENQR [resourcename]
```

Purpose of the 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.

Contents of the report

```
                        RMF V1R2  ENQ Resource Delays           Line 1 of 14
Command ===>                                                  Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  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 SCHMATE   SO
                                                              100 DRAGON    SO
                                                              100 SCHUMAC   SO
                                                              100 D71SJH1   SO
                                          100 SWARRENA  EW     100 AMOLLOY   SO
                                                              100 SCHMATE   SO
                                                              100 DRAGON    SO
                                                              100 SCHUMAC   SO
                                                              100 D71SJH1   SO
SYSIKJBC        (SYSS)                     1 *MASTER*  EW       1 *MASTER*  EO
 PETEG
SYSZVVDS        (SYS)                      1 LUCKYSM   EW       1 STEVEB    EO
 SYS1.MVS3.MCAT                             +CATALOG              +CATALOG
```

Figure 2-54. ENQR Report

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

Field descriptions

Table 2-31. 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.

Delayed %	The delay percentage of the job for a specific enqueued resource. This value is calculated as follows:
	`# Delay Samples`
	`Delayed % = --------------- * 100`
	`# Samples`
	Delay samples
	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.

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 %	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:
	`# Holding Samples`
	`Holding % = ---------------- * 100`
	`# Samples`
	Holding samples
	The number of samples when the holding job was holding the resource while the named job was delayed.

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 2-55. 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

How to request this report

To request the Group Response Time report, select a **1** on the Primary menu, and then a **5** on the Overview Report menu (shown in Figure 2-7 on page 2-22), or you can enter the following command using the format:

```
GROUP groupname           (in compatibility mode)
GROUP servclass,period    (in goal mode)
GROUP reportclass,period
```

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

```
GROUP HOTBATCH,1
```

Purpose of the report

The Group Response Time (GROUP) report presents information about using and delay values for a specific service or report class (or domain or performance group). The using and delay values are averages 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.

Contents of the report

```
                      RMF V1R2  Group Response Time
Command ===>                                                  Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  Time: 10.03.20  Range: 100   Sec

 Class: HOTBATCH     Period: 1    Description: Very Important Batch
 Primary Response Time Component: Using the processor

                                                  TRANS    --- Response Time ----
 WFL    Users    Frames   Vector   EXCP   PGIN   Ended    -- Ended TRANS-(Sec) -
  %    TOT  ACT   %ACT     UTIL    Rate   Rate   Rate      WAIT   EXECUT  ACTUAL
 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  OUTR    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 2-56. GROUP Report

A graphic report is not available for this report.

If the cursor is placed 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.

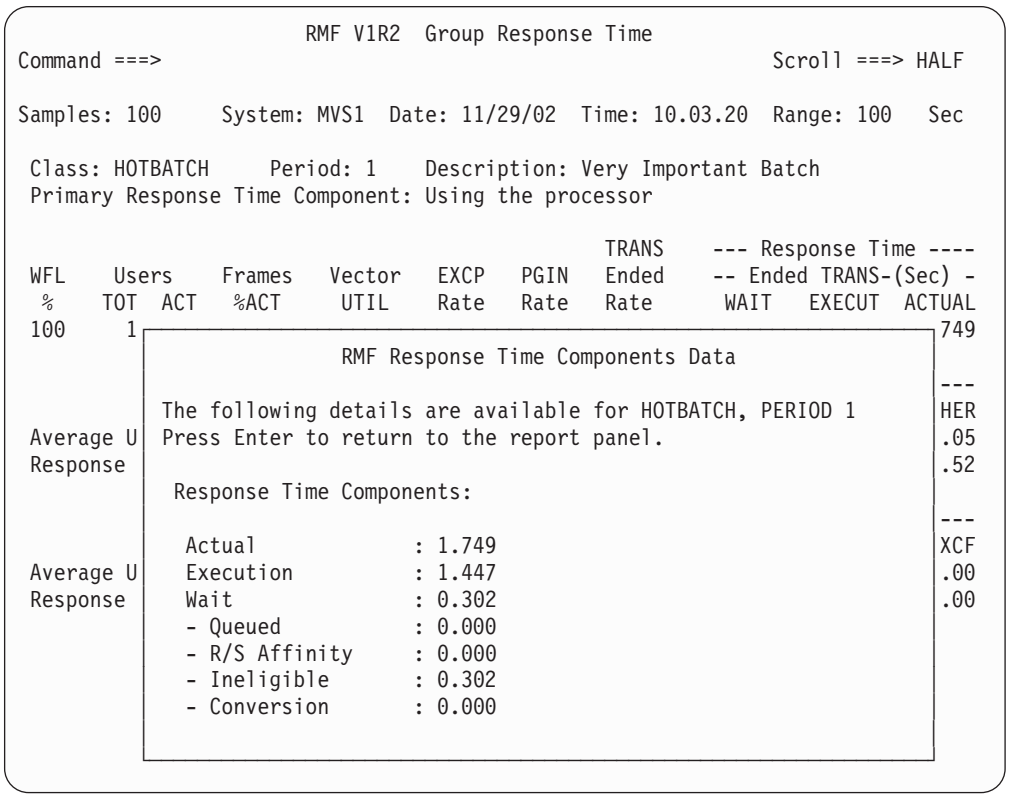

Figure 2-57. GROUP Report - Response Time Components

Field descriptions

Table 2-32. Fields in the GROUP Report

Field Heading	Meaning
Class	The name of the service or report class **Note:** For systems running in compatibility mode, the field `GROUP` is shown instead of `Class`, it displays the selected domain or performance group.
Period	The period number
Description	The description of the specified class, it is derived from the service policy.

Table 2-32. Fields in the GROUP Report (continued)

Field Heading	Meaning
Primary Response Time Component	A description of the component contributing to the largest percentage of total response time. The description can be: • Using the processor • Using I/O devices • Processor delay • Device delay • Storage delay for – common paging – local paging – virtual I/O – XMEM – HIPR – swap in – out and ready • Waiting for – JES – HSM – XCF • Waiting for volume mount • Waiting for operator reply • ENQ serialization delay • Delayed for unmonitored reasons
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	The number of users in the specified class. This category includes the following headings: **TOT** Total number of users equals the number of different users found in all address spaces for the specified class during the report interval. **ACT** Average number of active users found in all address spaces for the specified class during the report interval. An active user is either using a monitored resource, delayed for a monitored resource, or performing an activity that RMF does not measure. See the definition of Users/Active under "WFEX - Workflow/Exceptions report" on page 2-213 for more details.
Frames %ACT	The percentage of central and expanded 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 2-213. RMF accumulates the number of central storage frames for all active users during the report interval, then calculates the percentage as follows: `Frames %ACT = (ACSF + UESF) / (OCSF + IESF) * 100` **ACSF** Accumulated central storage frames **UESF** Expanded storage frames used by the active users **OCSF** Online central storage frames **IESF** Installed expanded storage frames
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.

Table 2-32. Fields in the GROUP Report (continued)

Field Heading	Meaning
EXCP Rate	The rate of EXCP requests per second for the specified class: $$\text{EXCP Rate} = \frac{\sum \text{all EXCP Requests}}{\text{Range Time}}$$
PgIn Rate	The rate at which pages are being swapped: $$\text{PgIn Rate} = \frac{\sum \text{Page-in Counts for Class}}{\text{Range Time}}$$
TRANS Ended Rate	The average number of ended transactions per second that occurred for the specified class during the report interval: $$\text{TRANS Ended Rate} = \frac{\text{\# Ended Transactions}}{\text{Range Time}}$$
Response Time	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. **Note:** 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. **WAIT** The average time (in seconds) that a transaction spent waiting because of one of these reasons: • **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. • **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. • **Ineligible - Operational or JES scheduling delay:** Average time a job was delayed due to operational delays or JES scheduling delays, examples are: – Job held by operator – Job class or job queue held – Duplicate jobname serialization – Job class execution limits • **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. **EXECUT** The average time that a transaction was active in the system. **ACTUAL** The sum of the execution time and the wait time, but does not include conversion time. If you place the cursor on one of these fields and press Enter, a **Response Time Components Data** pop-up panel will show a detailed breakdown of the different wait reasons and their average duration.

Table 2-32. Fields in the GROUP Report (continued)

Field Heading	Meaning
Average Users	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: **Total** 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. **Total** can be less than the sum of the individual categories if a user was found using or delayed in more than one category. **AVG USG** 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: **PROC** The average number of users using the processor during the report interval. **DEV** The average number of users using devices during the report interval. **Average Delay** 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. The average number delayed for is reported for the following categories: **PROC** Waiting for a processor **DEV** Waiting for a DASD or tape **STOR** Waiting for a COMM, LOCL, SWAP, XMEM, HIPR, or VIO page, or on the out/ready queue **SUBS** Waiting for services from JES, HSM, or XCF **OPER** Waiting for the operator to reply to a message or mount a tape **QUIESCE** The operator has quiesced the address space. A quiesced address space can show unexpected data: • A swappable address space will be swapped out, thus it can be OUTR and show storage delays. • 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. Cursor-sensitive control on this field gives you the Quiesce variation of the Job report. **ENQ** Waiting to use serially reusable resources that other jobs were using **OTHER** Unknown time **STOR Delay** A breakdown of the **Average Delay - STOR** field into the average number of users delayed for paging (Page), swapping (Swap), and swapped out and ready (OUTR). See the **% Delayed for** field description in Table 2-55 on page 2-156 for more information about the storage delays.

Table 2-32. Fields in the GROUP Report (continued)

Field Heading	Meaning
Average Users (continued)	**OUTR Swap Reason** A breakdown of the **STOR Delay OUTR** 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. The swap reasons can be: **TI** Terminal input wait **TO** Terminal output wait **LW** Long wait **XS** Auxiliary storage shortage **RS** Real storage shortage **DW** Detected long wait **RQ** Requested swap **NQ** Enqueue exchange swap **EX** Exchange swap **US** Unilateral swap **TS** Transition swap **IC** Improve central storage usage **IP** Improve system paging rate **MR** Make room for an out-too-long user **AW** APPC wait **IW** OMVS input wait **OW** OMVS output wait **SR** In-real swap See Table 5-51 on page 5-121 and Table 5-52 on page 5-121 for detailed information about the swap reasons. **SUBS Delay** A breakdown of the **Average Delay - SUBS** 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.

Table 2-32. Fields in the GROUP Report (continued)

Field Heading	Meaning
Response Time ACT	The average response time (in seconds) spent in each delay or using category for a transaction that was active during the report interval. The **Response Time ACT** line is divided into: **Total** Average time (in seconds) that an ended transaction was active in the system. The value is the same as the **Response Time, Ended TRANS (Sec), Active** value. **Note:** The value for Total and the sum of the individual using and delay values can be different. **Total** represents only ended transactions, while the breakdown of using and delay values represents all active transactions during the report interval. In addition, **Total** can be less than the sum of the individual categories if some of the users are delayed and using at the same time. **AVG USG** Average time (in seconds) that a transaction was using a processor (PROC) or device (DEV) during the report interval. **Average Delay** Average time (in seconds) that a transaction was delayed for the following reasons: **PROC** Waiting for a processor **DEV** Waiting for a DASD or tape **STOR** Waiting for a COMM, LOCL, SWAP, XMEM, HIPER, or VIO page, or on the out/ready queue **SUBS** Waiting for services from job-entry subsystem (JES), Hierarchical Storage Manager (HSM), or Cross-System Coupling Facility (XCF) **OPER** Waiting for the operator to reply to a message or mount a tape **QUIESCE** The operator has quiesced the address space. Cursor-sensitive control on this field gives you the Quiesce variation of the Job report. **ENQ** Waiting to use serially reusable resources that other jobs were using **OTHER** Unknown time **STOR Delay** Breakdown of the **Average Delay - STOR** 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 2-55 on page 2-156 for more information about the storage delays. **OUTR Swap Reason** Breakdown of the **STOR Delay OUTR** 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 **Average Users OUTR Swap Reason** field. See the field description on page 2-110 for a list of possible swap reasons. See the field descriptions for the Monitor I "PAGING - Paging Activity report" on page 5-107 for detailed information about the swap reasons. **SUBS Delay** Breakdown of the **Average Delay - SUBS** 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.

Table 2-32. Fields in the GROUP Report (continued)

Field Heading	Meaning
Statistical error	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. 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. To decrease statistical error, try increasing the report interval.

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
HIMMI    S   HOTBAT   S   HOTTSO   S   IBMUSER  S   IOHI     S   IOMED    S
IRLM     S   OE       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
STCSYS   S   STOCNSMR S   STODEX   S   STODIM   S   STOGIV   S   STOREC   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
DFHSM    R   DUMPSRV  R   GRS      R   HIMMI    R   IBM      R   LE       R
MASTER   R   MVSNFS   R   OE       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 2-58. 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 of the period you want reported (a number between 1 and 8).

Available Service and Report Classes
This list includes all service and report classes that had any activity during the current report interval.

Note: For systems running in compatibility mode, the Report Options panel displays the selected group and offers a selection list with all available domains and performance groups.

HSM - Hierarchical Storage Manager Delays report

How to request this report

To request the HSM report, select a **4** on the Primary menu, and then a **1** on the Subsystem Report menu (shown in Figure 2-10 on page 2-25), or or you can enter the following command using the format:

```
HSM [class,s-class]
```

Purpose of the 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.

Contents of the report

```
                        RMF V1R2  HSM Delays                      Line 1 of 3
Command ===>                                                   Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  Time: 10.03.20  Range: 100   Sec

          DLY  ---------------------- Main Delay Reason ------------------------
Jobname    %    %  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 2-59. HSM Report

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

Field descriptions

Table 2-33. 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 %	Delay the waiting job is experiencing because of contention for HSM during the report interval. This value is calculated as follows:
	`# Delay Samples`
	`DLY % = --------------- * 100`
	`# Samples`
	Delay samples
	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:
	• The MWE has a request from the user
	• The request is a "waited-on" request
	• The request has not completed processing
	• You receive one of the function codes listed under Main Delay Reason(s).
	Note: This DLY % value is also found in the HSM field on the job delay report.

Main Delay Reason(s)	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.
	% Indicates how much of the HSM delay of the job is caused by the reported subsystem function. This value is calculated as follows:
	`# Delay Samples`
	`% = --------------- * 100`
	`# Samples`
	Delay samples
	The number of samples when the job was delayed for HSM for a specific subsystem function.
	Note: The Main Delay Reason % values add up to the DEL % 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.
	The HSM F-codes (in decimal) and their explanations are as follows:
	F-Code Number — **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 2-39 on page 2-73 for an example. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed. See Figure 2-37 on page 2-69 for an example.

IOQUEUE - I/O Queuing Activity report

How to request this report

To request the I/O Queuing Activity report, select a **3** from the Primary menu, and then a **13** on the Resource Report menu (shown in Figure 2-9 on page 2-24), or you can enter the following command:

```
IOQUEUE
```

Purpose of the 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.

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 measurement 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.

Note: The report contains data for DASD control units only.

Data items that are not valid are marked by dashes (---) in the output display.

In contrast to the Monitor II IOQUEUE report, an LCU without any activity during the interval is not omitted from the report. If there are changes in the status or in the configuration during the interval, the report will be shown with the new configuration, but without measurement data.

Note: Starting with z/OS V1R2 RMF, the percentage about all channel paths being busy has been removed due to data inconsistencies for different channel types.

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.

Contents of the report

```
                          RMF V1R2  I/O Queuing Activity        Line 1 of 54
Command ===>                                                  Scroll ===> HALF

Samples:    100 System: MVS1  Date: 11/29/02  Time: 08.53.20     Range:  100 Sec

                   DCM Group       Cont Del Q   AVG  CHPID  %DP  %CU   AVG  AVG
Path DCM CTL Units MN MX DEF LCU   Rate Lngth   CSS  Taken  Busy Busy  CUB  CMR

11       400A 400B           0031                     2.86   0.0  1.1  0.0  0.0
         400C 400D
2E    Y  400A               0031                     5.88  10.7  0.2  0.0  0.0
20       400B               0031                     2.64  18.2  2.5  0.5  0.2
                  1  2  2   0031                     5.88  10.7  0.2  0.0  0.0
                              0031   1.1  4.78   1.7  11.39  10.2  1.0  0.0  0.0
11       440A               0032                     0.03   0.0 33.3  0.0  0.0
92    Y  440A               0032                     2.59  37.7  4.5  0.6  0.1
20       440B               0032                     0.00   0.0  0.0  0.0  0.0
                  0  1  2   0032                     2.59  37.7  4.5  0.6  0.1
                              0032   2.4  3.61   0.0   2.63  37.3  4.8  0.3  0.1
12       500A               0033                     0.29   5.5  0.0  0.0  0.0
93       500B               0033                     0.22  18.7  0.0  0.0  0.0
                              0033   0.0  0.00   0.0   0.51  11.7  0.0  0.0  0.0
12       700A               0034                     0.68  29.8  0.0  0.0  0.0
93       700B               0034                     0.46  43.7  0.0  0.0  0.0
                              0034   0.1  0.00   0.0   1.14  36.1  0.0  0.0  0.0
11       580A               0035                     0.03   0.0 33.3  0.0  0.0
20       580B               0035                     0.07   0.0 33.3  0.0  0.0
                  0  0  2   0035                     0.00   0.0  0.0  0.0  0.0
                              0035   0.0  0.00   0.0   0.10   0.0 33.3  0.0  0.0
11       5C0A               0036                     0.05   0.0  0.0  0.0  0.0
2E    Y  5C0A               0036                     1.39   2.3  0.0  0.0  0.0
```

Figure 2-60. IOQUEUE Report

The graphic form of this report shows the contention rate of each LCU.

Field descriptions

Table 2-34. Fields in the IOQUEUE Report

Field Heading	Meaning
Path	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.
DCM	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.
CTL Units	The hexadecimal identifier of each physical control unit associated with an online channel path in the LCU group.

Table 2-34. Fields in the IOQUEUE Report (continued)

Field Heading	Meaning
DCM Group	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. 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.
LCU	The hexadecimal number that identifies the logical control unit (LCU). 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. For each LCU, a summary line is reported in addition.
Cont Rate	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. `Cont Rate = # Enqueued Requests / Range Time`
Del Q Lngth	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. `Del Q Lngth = Accumulated Queue Length - # Enqueued Requests / # Enqueued Requests`
AVG CSS	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. `AVG CSS = Channel Subsystem Time / # I/O Operations Accepted`
CHPID Taken	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. `CHPID Taken = # I/O Operations Accepted on that Path / Range Time`
% DP Busy	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. `% DP Busy = DPB / DPB + CUB + SUC * 100` **DPB** Number of deferred I/O requests due to director port busy **CUB** Number of deferred I/O requests due to control unit busy **SUC** Number of successful I/O requests on that path

Table 2-34. Fields in the IOQUEUE Report (continued)

Field Heading	Meaning
% CU Busy	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. `% CU Busy = CUB / (DPB + CUB + SUC) * 100` **DPB** Number of deferred I/O requests due to director port busy **CUB** Number of deferred I/O requests due to control unit busy **SUC** Number of successful I/O requests on that path
AVG CUB	The average number of milliseconds of delay that an I/O request encountered for the channel path because the control unit was busy. `AVG CUB = Control Unit Busy Time / # I/O Operations Accepted on that Path`
AVG CMR	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). `AVG CMR = Initial Command Response Time / # I/O Operations Accepted on that Path`

JES - Job Entry Subsystem Delays report

How to request this report

To request the JES report, select a **4** on the Primary menu, and then a **2** on the Subsystem Report menu (shown in Figure 2-10 on page 2-25), or or you can enter the following command using the format:

```
JES [class,s-class]
```

Purpose of the 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.

Contents of the report

```
                        RMF V1R2  JES Delays                    Line 1 of 3

Command ===>                                                  Scroll ===> HALF

Samples:    100 System: MVS1  Date: 11/29/02  Time: 08.53.20     Range:  100 Sec

           DLY  ---------------------- Main Delay Reason(s) ----------------------
 Jobname    %    %  F-Code Explanation

 TCHANG      1    1   12   Waiting for JES2 job termination.
 MCKNSY1     1    1   12   Waiting for JES2 job termination.
 GROSSI      1    1   12   Waiting for JES2 job termination.
```

Figure 2-61. JES Report

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

Field descriptions

Table 2-35. 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 %	Delay the waiting job is experiencing because of JES during the report interval. This value is calculated as follows:
	`# Delay Samples`
	`DLY % = --------------- * 100`
	`# Samples`
	Delay samples
	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.
	For JES2 delay, RMF scans all subsystem job blocks (SJBs).
	See Main Delay Reason(s) in this report description for valid JES3 function codes.
	Note: This DLY% value is also found in the JES field on the Job Delay report.

Main Delay Reason(s)	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.
	% Indicates how much of the JES delay of the job is caused by the reported subsystem function. This value is calculated as follows:
	`# Delay Samples`
	`% = --------------- * 100`
	`# Samples`
	Delay samples
	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 "DEL %" field in this report description.
	Note: The Main Delay Reason % values add up to the DEL % 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.
	The JES2 F-codes (subsystem request number, in decimal) and their explanations are:
	F-Code Number **Explanation**
	01 Processing TSO OUTPUT command request.
	02 Waiting for JES2 to cancel a job.
	03 Waiting for job status info from JES2.
	12 Waiting for JES2 to purge a SYSOUT file.
	13 Waiting for JES2 to restart a job.

Table 2-35. Fields in the JES Report (continued)

Field Heading	Meaning
Main Delay Reason(s) (Cont'd)	The JES3 F-codes (subsystem interface service number, in decimal) and their explanations are: **F-Code Number** **Explanation** **023** Dynamically allocating data set to JES3. **026** Changing DDNAME of device or data set. **027** Changing data set use from SHR to OLD. **132** Allocate or deallocate spool data set. **138** Validating SYSOUT device name. **138** Waiting for JES3 to cancel a job. **138** Waiting for job status info from JES3. **144** Processing TSO OUTPUT command request. **144** Processing TSO OUTPUT command request. **147** Dynamically allocating data set to JES3. **149** Dynamically allocating data set to JES3. **151** Changing JES3 DD via dynamic allocation.

Report options

The JES Report Options panel is similar to the Device Report Options panel. See Figure 2-39 on page 2-73 for an example. Select YES for Jobs on the Report Options panel to display the Job Selection/Exclusion panel. See Figure 2-37 on page 2-69 for an example.

JOB - Job report

How to request this report

To request the Job report, select a **2** from the Primary menu and then a **5** on the Job Report menu (shown in Figure 2-8 on page 2-23), or you can enter the following command using the format:

```
JOB jobname
```

Purpose of the report

The Job report describes the reason why a specific job is delayed and provides possible causes leading to the delay.

Contents of the report

The Job report is available in ten variations. There is a Job report variation for each of the following types of delay:

- 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 report variation. The Job 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 selection menu, or cursor-sensitive control to select the type of delay to be reported. See “The Job Report Selection Menu” on page 2-23 for more information about requesting delay types.

Displaying the Job report

You can display the Job report in several ways:

- Place the cursor on a Jobname of any RMF report and press ENTER.
- Place the cursor on a % Delayed For field in a delay report or a Job 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 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 report

All variations of a Job report have the same top part:

```
                        RMF V1R2  Job Delays                      Line 1 of 1
Command ===>                                                   Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  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 2-62. Top Part of Job Report

Field descriptions

Table 2-36. Fields in the Job 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 2-63. Bottom Part of Job Report

The bottom part of the Job report has similar fields as the Delay report. Similar fields also appear on the Delay report. Table 2-16 on page 2-63 describes these fields. In the Job 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 report shows a job delayed by a device.

How to request this variation

To request the Device Delay variation, select a **3** from the Primary menu, and a **1** on the Job Report menu (shown in Figure 2-8 on page 2-23), or you can enter the following command using the format:

```
DEVJ jobname
```

```
------------------------- Volume BPXLK1 Device Data -------------------------
Number:    0409        Active:      62%        Pending:    3%     Average Users
Device:    33903       Connect:     17%        Delay DB:   1%        Delayed
Shared:    Yes         Disconnect:  42%        Delay CM:   0%          0.7
```

Figure 2-64. Device Delay Variation of the Job Report

The fields in the middle section of this report provide information about the device delaying the job.

Field descriptions

Table 2-37. Fields in the Device Delay Variation of the Job 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:	The number of parallel access volumes (base and alias) which were available at the end of the reporting interval. This field appears only for parallel access volumes. If the number has been changed during the reporting interval, it is followed by an '*'.
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 2-21 on page 2-78 for the calculation of this value.
Connect:	The percentage of time the device was connected to a channel path. Dashes in this field indicate hardware data is not available. See Table 2-21 on page 2-78 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: CON % > 40
Disconnect:	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 2-21 on page 2-78 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: DSC % > 40 and CON % < 30

Table 2-37. Fields in the Device Delay Variation of the Job Report (continued)

Field Heading	Meaning
Pending:	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 2-21 on page 2-78 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: PND % > 40 and CON % = 0 PND % > 40 and device is shared PND % > 30 and device is not shared.
Delay DB%:	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. $$\text{Delay DB\%} = \frac{\text{Accumulated DB Delay Time}}{\text{Range Time}} * 100$$
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 variation of the Job report shows a job delayed by a serially reusable resource.

How to request this variation

To request the Enqueue Delay variation, select a **3** from the Primary menu, and a **2** on the Job Report menu (shown in Figure 2-8 on page 2-23), or you can enter the following command using the format:

```
ENQJ jobname
```

```
---------------------------- Jobs Holding Resource -----------------------------
Job:      AMOLLOY             Job:      SCHMATE             Job:       DRAGON
Holding:     100%             Holding:     100%             Holding:     100%
Status:    Shared             Status:    Shared             Status:    Shared
```

Figure 2-65. Enqueue Variation of the Job 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 2-38. Fields in the Enqueue Variation of the Job 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 2-31 on page 2-104 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 report have the same format. They show a job delayed by either the HSM or JES address space.

How to Request this Variation

or you can enter the following command using the format:

```
HSMJ jobname
```

or

```
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 2-66. HSM Variation of Job Report

Field descriptions

The fields in the middle section of this report contain information about the HSM or JES address space.

Table 2-39. Fields in the HSM/JES Variation of the Job 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 2-7 for the calculation.
Using:	The using percentage of the JES/HSM address space. See "Address space using (%)" on page 2-8 for the calculation.
Delay:	The delay percentage of the JES/HSM address space. See Table 2-16 on page 2-63 for the calculation.
Primary delay category:	The category of delay contributing most to the overall delay. See the % Delayed For field in Table 2-16 on page 2-63 for details.
Primary delay reason:	The contents of this field depend on the primary delay category. See the Primary Reason field in Table 2-16 on page 2-63 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 variations of the Job report have the same format. They show a job delayed by either an operator message request or mount request.

How to request this variation

or you can enter one of the following commands using the format:

```
MSGJ jobname
```

or

```
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   0 100   0 Message
```

Figure 2-67. Operator Message Variation of the Job Report

Field descriptions

Table 2-40. Fields in the Operator Message Variation of the Job 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 variation of the Job report shows a job delayed by a processor.

How to request this variation

To request the Processor Delay variation, select a **3** from the Primary menu, and then an **8** on the Job Report menu (shown in Figure 2-8 on page 2-23) or you can enter the following command using the format:

```
PROCJ 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 2-68. Processor Delay Variation of the Job Report

The middle fields of this report provide information about the three main jobs holding the processor and causing delay.

Field descriptions

Table 2-41. Fields in the Processor Delay Variation of the Job 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 2-49 on page 2-144 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 2-18 on page 2-72 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 variation, select a **2** from the Primary menu, and then an **9** on the Job Report menu (shown in Figure 2-8 on page 2-23) or you can 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 -------------------------
        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   0 100   0 Quiesce
```

Figure 2-69. Quiesce Variation of the Job Report

Field descriptions

Table 2-42. Fields in the Quiesce Variation of the Job Report

Field Heading	Meaning
Primary Reason QUIESCE	The address space has been quiesced by the operator using the RESET command. A quiesced address space can show unexpected data: • A swappable address space will be swapped out, thus it can be OUTR and show storage delays. • 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.

Storage Delay variation

This variation of the Job report shows a job delayed by contention of storage.

How to request this variation

To request the Storage Delay variation, select a **3** from the Primary menu, and then an **10** on the Job Report menu (shown in Figure 2-8 on page 2-23) or you can 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 2-70. Storage Delay Variation of Job Report

The middle fields of this report provide information about the storage usage of the delayed job.

Field descriptions

Table 2-43. Fields in the Storage Delays Variation of the Job 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 2-55 on page 2-156 for the calculation of this value.
Idle Frames:	The average number of frames held by the job when it was idle. See Table 2-55 on page 2-156 for the calculation of this value.
Working Set:	The average amount of storage a user occupied while in storage. See Table 2-55 on page 2-156 for the calculation of this value.
Aux Slots:	The average number of auxiliary slots for each address space.

Table 2-43. Fields in the Storage Delays Variation of the Job Report (continued)

Field Heading	Meaning
Page In Rate	The rate at which pages are being read into central storage. $\text{Page In Rate} = \frac{\sum \text{all Page-in Counts for Group}}{\text{Resident Time}}$ The resident time is the total time the address space was swapped in. The page-in rate includes the shared storage page-ins.
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 2-61 on page 2-165 for the calculation of this value.
DIV Frames:	The number of central and expanded storage frames used by DIV. This count is accumulated only for jobs, not for domain or performance group. See Table 2-61 on page 2-165 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 report shows a job delayed by XCF.

How to request this variation

To request the XCF variation, select a **3** from the Primary menu, and then a **11** on the Job Report menu (shown in Figure 2-8 on page 2-23) or you can enter the following command using the format:

```
XCFJ jobname
```

```
------------------------------ XCF Path Summary -----------------------------
                   1st Path    2nd Path    3rd Path    4th Path
Delay:                4%          3%
Dev Number:          0E80        0CA0
```

Figure 2-71. XCF Variation of Job Report

Field descriptions

Table 2-44. Fields in the XCF Variation of the Job Report

Field Heading	Meaning
Delay:	The delay percentage of the XCF address space. See Table 2-16 on page 2-63 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 report. In addition to the delays shown above you can use the Utility to have the following delay percentages shown.

Table 2-45. Additional Fields in the Job 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 • it has used up its CPU service as specified in the WLM policy for the resource group to which the job belongs • 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 *z/OS MVS Planning: Workload Management*).

Report options

Each variation of the Job report uses the same Report Options panel. Use this panel to view and select an available jobname.

```
                          RMF Job Report Options                Line 1 of 66
Command ===>                                                  Scroll ===> HALF

Change or verify parameters for all job reports.

    Jobname ===> DCRPROCA    Name of job to be reported

                                Available Jobs
*MASTER*   ADAM       ADRIAN     ALAN1      ALEXIS     ALISONW    ALLOCAS
ALPERT     ALTQTS     ALTQTN66   ALTAN      AMELIA     AMSAQTS    AMYH
ANANIA     ANDREW     ANN        ARTHUR     ARVIN      ASCOTT     AWESOMX
AXLT       BART       BASS       BASS11     BATTER     BEERS      BENCH
BERNIEP    BETHP      BETHT      BILLIE     BILLR      BJXYZ      BOBJUD
BOYCOT     BRICK      BRUSH      BULL       BUTTON     BYRON      CANNON
CAROLL     CASTLE     CATALOG    CATHYM     CATHYX     CHRISD     CHRISTI
CHUCKG     CHUWU      CJWAXX     CLRUSER    CNPLFS     COFFEE     CONSOLE
COOK       CORNER     CORNY      COSTER     CRAIGJ     CRANE      CRIMP
```

Figure 2-72. 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 reports when it is available.

MSI - Migration SYSINFO report including execution velocity

How to request this report

To request the MSI report, select **U** on the Primary Menu, then **1** on the User menu, or you can enter the following command:

```
MSI
```

Purpose of the report

The Migration SYSINFO (MSI) report presents the same information as the System Information report except that the VEC UTL column is replaced with EXEC VEL.

Use this report in compatibility mode to estimate the execution velocity of special workloads that you want to plan execution velocity goals for in goal mode.

Contents of the report

```
                    RMF V1R2  System Information              Line 1 of 22
Command ===>                                                Scroll ===> HALF

Samples:    100   System: MVS1 Date: 11/29/02 Time: 10:31:40 Range:    100 Sec

   --------------------- 9672 Model R66 ------------------------ IPS = IEAIPST0
   Average CPU Util%:    73          Appl% / EAppl%:    63/ 65   OPT = IEAOPT00
   Avg. MVS CPU Util%:   84          Partition Name:  MVS1       ICS = IEAICS00

        WFL EXEC --Users--  RESP TRANS -AVG USG-  -Average Number Delayed For -
Name     %  VEL   TOT  ACT  Time  /SEC PROC  DEV  PROC  DEV STOR SUBS OPER  ENQ

*SYSTEM  85   91  146    3        4.39  0.5  3.0   0.0  0.0  0.0  0.0  0.0  0.0
*TSO     97   99   68    1        4.39  0.2  0.8   0.0  0.0  0.0  0.0  0.0  0.0
*BATCH   56   63    3    2        0.00  0.1  1.9   0.0  0.0  0.0  0.0  0.0  0.0
*STC     98   98   75    1        0.00  0.3  0.3   0.0  0.0  0.0  0.0  0.0  0.0
*ASCH             0    0  0.00  0.00  0.0  0.0   0.0  0.0  0.0  0.0  0.0  0.0
*OMVS             0    0  0.00  0.00  0.0  0.0   0.0  0.0  0.0  0.0  0.0  0.0
*ENCLAVE          0  N/A         N/A  0.0  N/A   0.0  N/A  0.0  N/A  N/A  N/A
DMN000   78   80   20    0  3.6K  0.00  0.1  0.1   0.0  0.0  0.0  0.0  0.0  0.0
DMN001   56   58    3    2  0.06  0.00  0.1  1.9   0.0  0.0  0.0  0.0  0.0  0.0
DMN002   71   79   67    0  0.83  3.96  0.1  0.1   0.0  0.0  0.0  0.0  0.0  0.0
DMN003   59   64    0    0  22.3  0.22  0.0  0.1   0.0  0.0  0.0  0.0  0.0  0.0
DMN009            1    0  0.00  0.00  0.0  0.0   0.0  0.0  0.0  0.0  0.0  0.0
PG000    53   59   20    0  0.00  0.00  0.1  0.1   0.0  0.0  0.0  0.0  0.0  0.0
PG001    72   81    3    2  0.52  0.00  0.1  1.9   0.0  0.0  0.0  0.0  0.0  0.0
```

Figure 2-73. Migration SYSINFO Report

For each domain and performance group in compatibility mode, this report shows an **execution velocity** that you can use to define velocity goals.

Execution velocity is also shown on the SYSSUM report. In a sysplex with only one processor the value for service classes and workload groups with low utilization may be different on the SYSSUM report and the SYSINFO report. A different sample technology is used in the SYSSUM report, giving the more exact value.

Figure 2-73 on page 2-136 shows a modified SYSINFO report that replaces the **Vector Utilization (VEC UTL)** column with the **Execution Velocity (EXEC VEL)** column. The value shown in this column is calculated in the same way as it is for a service class or service class period in goal mode.

Field descriptions

The fields in the MSI report are identical to the fields in the SYSINFO report, except for the Execution Velocity field. See Table 2-69 on page 2-180 for a detailed description of all fields.

Table 2-46. Fields in the MSI Report

Field Heading	Meaning
EXEC VEL	The value for execution velocity 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.

OPD - OMVS Process Data report

How to request this report

To request the OMVS Process Data report, select an **1** from the Primary menu and then a **7** on the Overview Report menu (shown in Figure 2-7 on page 2-22), or you can enter the following command:

```
OPD
```

Purpose of the report

UNIX System Services (USS) 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 shortly 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.

Contents of the report

```
                  RMF V1R2 OMVS Process Data                   Line 1 of 24

Command ===>                                                 Scroll ===> HALF

Samples: 18      System: SYS4  Date: 11/29/02  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 2-74. OPD Report

There exists no graphic form of this report.

Table 2-47. 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 that the process consumed during the report range.
Total	Total computing time since the process has been started.

Table 2-47. Fields in the OPD Report (continued)

Field Heading	Meaning
Server	If the process represents a server, one of the following is shown: **FILE** Network file server **LOCK** Network lock server **FEXP** Network file exporter **SFDS** Shared file server For non-server processes, 'N/A' is shown.

The following two pop-up panels show examples of process details for a server process and a non-server process:

```
                    RMF OMVS Process Data - Details

Press Enter to return to the Report panel.

Start Time/Date : 12.08.57 03/09/2002
Command         : GFSAMAIN
Process-ID :       25    Parent Process-ID :        1
Jobname    :  MVSNFSS    User Name         :   MVSNFS
ASID       :     0049    Hexadecimal ASID  :     0031

Appl%  :  0.0   Total CT     : 0.485   LW-PID     :        0

Server Information:
  Name : MVSNFS
  Type : FILE   Active Files :     0   Max. Files :     200K

Process State : MF
M: Multiple threads, no pthread_create used
F: File system kernel wait
```

Figure 2-75. OPD Report - Details for Server Process

```
                    RMF OMVS Process Data - Details

Press Enter to return to the Report panel.

Start Time/Date : 12.08.52 03/09/2002
Command         : PORTMAP
Process-ID :       20    Parent Process-ID :        1
Jobname    :  PORTMAP    User Name         :  PORTMAP
ASID       :     0046    Hexadecimal ASID  :     002E

Appl%  :  0.0   Total CT     : 0.099   LW-PID     :        0

Server Information:
  Name : N/A
  Type : N/A    Active Files :   N/A   Max. Files :      N/A

Process State : 1MF
 1: Single Thread
 M: Multiple threads, no pthread_create used
 F: File system kernel wait
```

Figure 2-76. OPD Report - Details for Non-Server Process

Table 2-48. 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 ID of the address space.
Appl%	Percentage of TCB and local/global SRB time that the process consumed during the report range.
Total CT	Total computing time since the process has been started.
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: **FILE** Network file server **LOCK** Network lock server **FEXP** Network file exporter **SFDS** 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: **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. Canceled, 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 preceeding X) format
   3. Jobname     ===>           Jobname associated with a process
   4. User        ===>           User name associated with a process
```

Figure 2-77. 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.

PROC - Processor Delays report

How to request this report

To request the Processor Delay report, select a **4** from the Primary menu, then an **1** from the Resource Report menu (shown in Figure 2-8 on page 2-23) or you can enter the following command using the format:

```
PROC [class,s-class]
```

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

```
PROC T, TSOPRIME
```

Purpose of the 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.

Contents of the report

```
                        RMF V1R2 Processor Delays              Line 1 of 138
Command ===>                                                  Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  Time: 10.03.20  Range: 100   Sec

             Service  DLY USG  Appl EAppl ----------- Holding Job(s) -----------
Jobname   CX Class     %   %      %     %  %  Name      %  Name      %  Name

#STONE1   B  NRPRIME   41  60   55.1  62.1  18 D10PAV1   15 SAYLEKR   12 DAVEP
SAYLEKR   B  NRPRIME   24  21   20.7  20.7  13 DB2MDIST  12 #STONE1    9 CATALOG
SAYLEKQ   B  NRPRIME   22  14   14.8  14.8  16 D10PAV1   11 SAYLEKR    8 TREVORJ
RSTSHYS0  BO NRPRIME   15  24   18.6  18.6   9 DAVEP      8 #STONE1    7 GRS
CHARLESR  B  NRPRIME   15   7    2.9   4.2   8 D10PAV1    6 SAYLEKR    5 DAVEP
LUCKYSMP  B  NRPRIME    6   8    4.1   4.1   5 DAVEP      4 STEVEB     2 DSEAGER
DFHSM     S  SYSSTC     5  17   13.0  13.0   3 D10PAV1    3 #STONE1    2 *MASTER*
APETER    B  NRPRIME    5   5    3.7   3.7   3 DB2MDIST   3 SAYLEK     2 #STONE1
MARYPATM  B  NRPRIME    5   2    2.2   3.1   2 GRS        2 DAVEP      2 #STONE1
D10PAV1   T  TSOPRIME   4  25   23.5  23.5   2 BKELLER    2 DFHSM      2 RHANSON
DAVEP     T  TSOPRIME   4  25   14.9  14.9   3 DB2MDIST   2 DSEAGER    2 LUCKYSMP
HJPB88E   B  NRPRIME    4   3    1.1   1.1   3 #STONE1    2 GRS        2 DB2MDIST
```

Figure 2-78. PROC Report

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

Field descriptions

Table 2-49. 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: **S** Started task **T** TSO **B** Batch **A** ASCH **O** OMVS An **O** 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. **Note:** For systems running in compatibility mode, the fields `DMN` and `PG` are shown instead of `Service Class`, they display the domain and performance group for the specified job.

DLY %	Delay the waiting job (address space) is experiencing because of contention for the processor during the report interval, expressed as a percentage.
	`# Delay Samples`
	`DLY % = --------------- * 100`
	`# Samples`
	
	Delay samples
	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.
	Note: This DLY% value is also found in the PROC field on the Job Delay report.

USG%	The percentage of time when the job is receiving service from the processor.
	`# Using Samples`
	`USG % = --------------- * 100`
	`# Samples`
	
	Using samples
	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.

Appl%	Percentage of CPU time as sum of TCB time, global and local SRB time, and preemtable or client SRB time consumed on behalf of this address space.
EAppl%	Percentage of CPU time as sum of TCB time, global and local SRB time, preemtable or client SRB time, and enclave CPU time consumed within this address space.

Table 2-49. Fields in the PROC Report (continued)

Field Heading	Meaning
Holding Job(s)	Up to three jobs that, by their use of the processor, contributed most to the delay of the job listed under Jobname. **%** The percentage of delay caused by the named job to the job waiting to use the processor. `% = # Holding Samples / # Samples * 100` **Holding samples** The number of samples when the job was using the processor and delaying the other job (indicated in the Jobname field). **Note:** 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.
Name	The name of a job contributing to the delay of the job waiting to use the processor. The name ***ENCLAVE** 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.

Monitor III Utility fields

Table 2-50 shows the additional fields you can select for this report.

Table 2-50. Additional Fields in the PROC Report

Field Heading	Meaning
PRCPCAPD	This column contains the actual delay caused by WLM • due to a resource group maximum • or due to disretionary goal management. This means that the work in question may be overachieving its goal. It may be capped in order to divert its resources to run discretionary work (see also section 'Using Discretionary Goals' in *z/OS MVS Planning: Workload Management*). 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.
PRCPVEC	The ratio of vector usage time to total processor usage time for this job.
PRCPTCBT	TCB time consumed in this address space.
PRCPSRBT	SRB time consumed in this address space by local or global SRBs.
PRCPPCST	Preemptable or client SRB time consumed on behalf of this address space.
PRCPEPST	Preemtable or client SRB time and enclave CPU time consumed within this address space.

Report options

The PROC Report Options panel is similar to the DEV Report Options panel. See Figure 2-39 on page 2-73 for an example. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed. See Figure 2-37 on page 2-69 for an example.

RG - Resource Group Data report

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
```

Purpose of the 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.

Contents of the report

Figure 2-79 shows a modified version of the Sysplex Summary report that replaces all columns containing **actual values** and the **Performance Index** column with columns for **Duration**, **Resource Group** and **Minimum**, **Maximum** and **Actual Capacity**.

```
                    RMF V1R2 Sysplex Summary - RMFPLEX1       Line  1 of 35
Command ===>                                                 Scroll ===> HALF

WLM Samples: 100   Systems:  8 Date: 11/29/02 Time: 10:31:40 Range: 100   Sec

Service Definition: SLA_2002              Installed at: 07/01/02, 06.02.00
     Active Policy: WORKDAYS              Activated at: 07/01/02, 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                                     PRODUCT    2000   4500   1471
         1  3             1.00M AVG    1000
         2  D
BATCH.2  S  D                                  PRODUCT    2000   4500    927
CICSTOR  S  1     90                           CICS        N/A   2100    752
CICSAOR  S  1     90                           CICS        N/A   2100    846
CICSDOR  S  1     90                           CICS        N/A   2100    492
DEVELOP  S                                     DEVELOP     500   1000    803
         1  1             1.00  AVG     500
         2  3             3.00  90%    2000
         3  5     40
^POSMULTI S  1             0.50  90%
SYSSTC   S       N/A       N/A
```

Figure 2-79. RG Report

Field descriptions

Fields particular to the RG report are described here. For a detailed description of other fields, see Table 2-72 on page 2-194.

Table 2-51. 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	Minimum service rate, in unweighted CPU service units per second, as defined for that resource group. **Note:** The minimum and maximum values must not be specified by the user.
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.

RLSLRU - VSAM LRU Overview report

How to request this report

To request the VSAM LRU Overview report, select an **S** from the Primary menu and then a **12** on the Sysplex Report menu (shown in Figure 2-6 on page 2-21), or you can enter the following command:

```
RLSLRU
```

In addition, you can navigate to this report with cursor-sensitive control from the VSAM RLS Activity report.

Purpose of the 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.

Contents of the report

```
                  RMF V1R2 VSAM LRU Overview   - SYSPLEX         Line 1 of 2
Command ===>                                                  Scroll == => HALF

Samples: 120     Systems: 2    Date: 11/29/02  Time: 13.25.00  Range: 120   Sec

MVS       Avg CPU  - Buffer Size -  Accel  Reclaim  ------ Read -----
System     Time      Goal   High      %       %     BMF%   CF%  DASD%

SYS4      390.3       10M    10M     0.0     0.0    83.0   0.0   17.0
SYS5      454.2       10M    10M     0.0     0.0    83.0   0.0   17.0
```

Figure 2-80. 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: SYS5
Press Enter to return to the Report panel.

  Size      Low       High      Average
   2K     1,000      3,000        1,500
   4K     8,000     10,000        9,000
   6K         0          0            0
   8K         0          0            0
  10K         0          0            0
  12K         0          0            0
  14K         0          0            0
  16K         0          0            0
  18K         0          0            0
  20K         0          0            0
  22K         0          0            0
  24K         0          0            0
  26K         0          0            0
  28K         0          0            0
  30K         0          0            0
  32K         0          0            0
```

Figure 2-81. VSAM LRU Overview - Buffer Counts by Pool

Field descriptions

Table 2-52. Fields in the RLSLRU Report

Field Heading	Meaning
MVS System	MVS system name.
Avg CPU Time	Average CPU time spent by BMF LRU processing during each reporting interval (milliseconds).
Buffer Size: Goal	Buffer size goal (bytes). If no valid goal has been defined, MAX will be shown.
Buffer Size: High	Buffer size actual high value (bytes).
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.
BMF Read %	Percentage of READ requests that could be satisfied from local buffers being managed by SMSVSAM.
Cache Read %	Percentage of CF cache structure READ requests.
DASD Read %	Percentage of READ requests to DASD.

Table 2-53. Fields in the Buffer Counts by Pool Panel

Field Heading	Meaning
Size	Size of the storage pool.
Low High Average	Low, high and average value of the number of BMF buffers for the pool during this interval.

RLSSC/RLSDS - VSAM RLS Activity report

How to request this report

The VSAM RLS Activity report has two different flavours:

- 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.

To request the VSAM RLS Activity report, select an **S** from the Primary menu and then a **10** (for storage class) or an **11** (for data set) on the Sysplex Report menu (shown in Figure 2-6 on page 2-21), or you can enter the following command:

```
RLSSC [storage_class]

RLSDS [data_set_name]
```

In addition, you can navigate between both reports with cursor-sensitive control.

Purpose of the report

With VSAM RLS, GETs and PUTs are executed by SMSVSAM on behalf of the application. When the application's data request can be satisifed 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?

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 V1R2   VSAM RLS Activity  - SYSPLEX        Line 1 of 12
Command ===>                                                   Scroll == => HALF

Samples: 59      Systems: 2    Date: 11/29/02  Time: 13.16.00  Range: 60    sec

LRU Status   : Good
Contention % :  0.0
False Cont % :  0.0

Stor Class  Access  Resp   -------- Read ----------  ------ BMF -------  Write
                    Time    Rate  BMF%   CF%  DASD%  Valid%  False Inv%  Rate

RLS
             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
             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
             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
```

Figure 2-82. VSAM RLS Activity by Storage Class - Sysplex Totals View

Cursor-sensitive control on the LRU STATUS field displays the VSAM LRU Overview report (see Figure 2-80 on page 2-148) with the Local Buffer Manager LRU statistics for each system.

Cursor-sensitive control on a storage class name redisplays the report with a system and CF cache structure breakdown for the selected storage class.

```
                  RMF V1R2 VSAM RLS Activity   - SYSPLEX        Line 1 of 23
Command ===>                                                   Scroll == => HALF

Samples: 120     Systems: 2    Date: 11/29/02  Time: 13.25.00  Range: 120   Sec

LRU Status   : Good        Storage Class : RLS
Contention % :  0.0        Cache Set     : PUBLIC1
False Cont % :  0.0

System/CF   Access  Resp   -------- Read ----------  ------ BMF -------  Write
                    Time    Rate  BMF%   CF%  DASD%  Valid%  False Inv%  Rate

*ALL
             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
             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
             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 2-83. VSAM RLS Activity by Storage Class - System/CF Structure View

At 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 2-24 on page 2-43) 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 V1R2 VSAM RLS Activity   - SYSPLEX        Line 1 of 20
Command ===>                                                  Scroll == => HALF

Samples: 120     Systems: 2    Date: 11/29/02  Time: 13.25.00  Range: 120  Sec

LRU Status   : Good
Contention % :  0.0
False Cont % :  0.0

Sphere/DS   Access  Resp   -------- Read ----------  ------ BMF -------  Write
                    Time    Rate  BMF%   CF%  DASD%  Valid%  False Inv%  Rate

BMAI.VSAMIN.MEGA
 BMAI.VSAMIN.MEGA.AIX.DATA
             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
             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
             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
             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 2-84. VSAM RLS Activity by Data Set - Sysplex Totals View

Cursor-sensitive control on the LRU STATUS field displays the VSAM LRU Overview report (see Figure 2-80 on page 2-148) with the Local Buffer Manager LRU statistics for each system.

Cursor-sensitive control on a data set name redisplays the report with a system breakdown for the selected VSAM data set. At top of the report, the sysplex wide totals for the data set will be displayed (indicated by *ALL) followed by report lines per system.

```
                RMF V1R2 VSAM RLS Activity   - SYSPLEX         Line 1 of 9
Command ===>                                                 Scroll == => HALF

Samples: 120     Systems: 2    Date: 11/29/02  Time: 13.25.00  Range: 120  Sec

LRU Status   : Good        Storage Class : RLS01
Contention % :  0.0        Cache Set     : RLSCSET
False Cont % :  0.1        Data Set      : RLSADSW.VF01D.INVENTOR.INDEX

System/CF   Access  Resp   -------- Read ----------  ------ BMF -------  Write
                    Time    Rate  BMF%   CF%  DASD%  Valid%  False Inv%  Rate

*ALL
             DIR    0.000  24.27  97.7   2.2    0.1    100      0.00      0.00
             SEQ    0.000   0.00   0.0   0.0    0.0    0.0      0.00      0.00
N64
 RLSCACHE01
             DIR    0.000   0.20   100   0.0    0.0    100      0.00      0.00
             SEQ    0.000   0.00   0.0   0.0    0.0    0.0      0.00      0.00
N65
 RLSCACHE01
             DIR    0.000  24.07  97.7   2.2    0.1    100      0.00      0.00
             SEQ    0.000   0.00   0.0   0.0    0.0    0.0      0.00      0.00
```

Figure 2-85. VSAM RLS Activity by Data Set - System View

Cursor-sensitive control on a CF structure name displays the CF Activity report (see Figure 2-24 on page 2-43) for this structure.

Field descriptions

Table 2-54. Fields in the VSAM RLS Activity Report

Field Heading	Meaning
LRU Status	LRU status of local buffers under control of BMF (Buffer Management Facility). **Good** BMF is at or below its goal on all systems. **Accelerated** BMF is over the goal on at least one system, and the buffer aging algorithms were accelerated. **Reclaimed** 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.
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.
Storage Class	Storage class name.
Cache Set	DFSMS cache set name.
System	MVS system name (*ALL indicates a sysplex-wide view).
CF	CF cache structure name.
Sphere	VSAM sphere name.
DS Data Set	VSAM data set name.
Access	Indicates whether the values in this row are shown for direct access (DIR) or sequential access (SEQ).
Resp Time	Average response time of all requests (seconds).

Table 2-54. Fields in the VSAM RLS Activity Report (continued)

Field Heading	Meaning
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. **Note:** 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: • 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). • The coupling facility has lost track of the integrity status of the buffer (BMF false invalid READ hits). To make use of BMF Valid%, following formulas are helpful: • BMF Invalid Read Hits = True + False Invalid Read Hits • BMF Read Hit% = BMF Read% / BMF Valid% * 100 • BMF Invalid Read Hit% = BMF Read Hit% - BMF Read%
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.

STOR - Storage Delays report

How to request this report

To request the Storage Delays report, select a **3**, from the Primary menu, then an **6** from the Resource Report menu (shown in Figure 2-9 on page 2-24) or you can enter the following command using the format:

```
STOR [class,s-class]
```

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

```
STOR T, TSOPRIME
```

Purpose of the report

The Storage Delays report is job-oriented. It displays storage delay information for all jobs.

Contents of the report

```
                        RMF V1R2  Storage Delays               Line 1 of 678
Command ===>                                                  Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  Time: 10.03.20  Range: 100   Sec

             Service  DLY    ------- % Delayed for ------   -- Working Set --
Jobname   C  Class     %     COMM  LOCL  SWAP  OUTR  OTHR   Central  Expanded
JES2      S  SYSSTC   100       0     0     0     0   100      1931      1311
CAROLYN   T  TSOPRIME 100       0     0     0     0   100        41       917
MVSNFS    S  SYSSTC   100       0     0     0     0   100       166       262
DB2MDIST  S  SYSSTC   100       0     0     0     0   100       156        83
STYMIE    T  TSOPRIME 100       0     0     0     0   100       319       410
INIT      S  SYSSTC    73       0     0     0     0    73       102        29
INIT      S  SYSSTC    53       0     0     0     0    53         0        90
HJPB88E   B  ********  47       0     0     0     0    47       206        35
INIT      S  SYSSTC    26       0     0     0     0    26       103        29
IFBOBR00  B  NRPRIME    1       0     0     0     0     1       401        56
JHWANGX   B  NRPRIME    1       0     0     0     0     1       225        29
*MASTER*  S  SYSTEM     0       0     0     0     0     0       287       285
PCAUTH    S  SYSSTC     0       0     0     0     0     0        61         6
RASP      S  SYSTEM     0       0     0     0     0     0        99        58
TRACE     S  SYSSTC     0       0     0     0     0     0       138         3
DUMPSRV   S  SYSTEM     0       0     0     0     0     0        39        89
XCFAS     S  SYSTEM     0       0     0     0     0     0       505       159
GRS       S  SYSSTC     0       0     0     0     0     0      4387      2031
```

Figure 2-86. STOR Report

The graphic form of this report shows the percentage of each user's time that COMM, LOCL, SWAP, OUTR, and OTHR contributed to the delay of the job for storage.

Field descriptions

Table 2-55. 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: **S** Started task **T** TSO **B** Batch **A** ASCH **O** OMVS
Service Class	The name of the service class that a specified job has been running in. **Note:** For systems running in compatibility mode, the fields `DMN` and `PG` are shown instead of `Service Class`, they display the domain and performance group for the specified job.

DLY %	Delay the waiting job (address space) is experiencing because of contention for storage during the range period, expressed as a percentage.
	`# Delay Samples`
	`DLY % = --------------- * 100`
	`# Samples`
	Delay samples
	The number of samples where the job was delayed for storage one or more times due to COMM, LOCL, OTHR, SWAP, and OUTR
	Note: This value is also found in the STOR field on the Delay report.

% Delayed for	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. The meaning of each category follows: **COMM** 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. **LOCL** 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. **SWAP** 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. **OUTR** The percentage that being swapped out and ready contributes to the delay of the job. **OTHR** The sum of: • VIO (virtual I/O) • XMEM (paging delays from cross-memory address spaces) • 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. Cursor-sensitive control on this field gives you the STORR report. 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.

Table 2-55. Fields in the STOR Report (continued)

Field Heading	Meaning
Working Set	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.

Monitor III Utility fields

You can use the Monitor III Utility to customize the Storage Delays report. In addition to the delays shown above, you can use the Utility to have the delay percentages in Table 2-56 shown in the Storage Delays report.

Table 2-56. 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 2-39 on page 2-73 for an example of the Report Options panel. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed. See Figure 2-37 on page 2-69 for an example.

STORC - Common Storage report

How to request this report

To request the Common Storage report, select a **3**, from the Primary menu, then a **10** from the Resource Report menu (shown in Figure 2-9 on page 2-24) or you can enter the following command using the format:

```
STORC [class,s-class]
```

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

```
STORC T, TSOPRIME
```

Purpose of the 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 2-57 on page 2-159.

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 2-58 on page 2-160.

Notes:

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.

Contents of the report

```
                        RMF V1R2  Common Storage               Line 1 of 680
Command ===>                                                  Scroll ===> HALF

Samples: 100     System: MVS3  Date: 11/29/02  Time: 10.03.20  Range: 100 Sec

                                    ---- Percent ----   ------- Amount --------
System Information                  CSA ECSA SQA ESQA     CSA  ECSA   SQA  ESQA
 IPL Definitions                                        3380K   28M 1296K   97M
 Peak Allocation Values              30   73  90   43   1007K   20M 1171K   42M
 Average CSA to SQA Conversion        1    0           28672     0
 Average Use Summary                 30   73  90   43    999K   20M 1162K   42M
 Available at End of Range           70   27  12   57   2379K 7856K  158K   55M

               Service        ELAP  -- Percent Used -   ----- Amount Used -----
Jobname  Act C Class    ASID  Time  CSA ECSA SQA ESQA     CSA  ECSA   SQA  ESQA
%MVS                                 6    7  52   38    196K 2016K  676K   37M
%REMAIN                              2    2   0    0   78920  523K  3944 54216
JES2         S SYSSTC     33   5.2D   1   30   0    0   20712 8748K    48   544
*MASTER*     S SYSTEM      1   5.2D   2    6  15    2   65944 1736K  198K 2349K
VTAM         S SYSSTC     67   3.3D   1   10   0    0   20618 2959K     0   352
TSO          S SYSSTC     86   3.3D   4    0   0    0    138K  2632     0   224
RMF33        S SYSSTC     72   5.2D   0    0   2    1       0   424 32792  948K
```

Figure 2-87. STORC Report

Graphic support for the report is not available.

Field descriptions

Table 2-57. 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 *z/OS MVS Initialization and Tuning Reference*.
Peak Allocation Values	The peak common storage (CSA, ECSA, SQA, and ESQA) values detected by Monitor III during the specified report interval. The peak allocation values include common storage used by: • Active jobs • System activity not related to a specific job • Ended jobs that did not release all common storage 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.

Table 2-57. Fields in the STORC Report - System Information Section (continued)

Field Heading	Meaning
Average Use Summary	The average common storage (CSA, ECSA, SQA, and ESQA) usage during the specified report interval. The average use values include common storage used by: • Active jobs • System activity not related to a specific job • Ended job that did not release all common storage
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 report interval.

Table 2-58. 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: **%MVS** Summary information about common storage being getmained with the GETMAIN parameter OWNER(SYSTEM). **%REMAIN** Summary information about common storage that was not released by ended jobs. ***srvcls** 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. **Note:** For systems running in compatibility mode, the names ***DMNxxx** and ***PGxxx** show summary information about address spaces in a specific domain or performance group.
Act	Job status at end of the specified report interval. **N** A job ended during the report interval and released all common storage **H** 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 **H,** the STORCR (Common Storage Remaining) report will be shown with the selected job reported first. **Note:** 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: **S** Started task **T** TSO **B** Batch **A** ASCH **O** OMVS
Service Class	The name of the service class that a specified job has been running in. **Note:** For systems running in compatibility mode, the fields `DMN` and `PG` are shown instead of `Service Class`, they display the domain and performance group for the specified job.
ASID	The unique system-assigned identifier for the address space in which the job is running.

Table 2-58. Fields in the STORC Report - Job Information Section (continued)

Field Heading	Meaning
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. **Note:** 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). **Note:** 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 shown above you can use the Utility to have the following values shown.

Table 2-59. 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 2-88. 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 2-37 on page 2-69 for an example.

STORCR - Common Storage Remaining report

How to request this report

To request the Common Storage Remaining report, select a **3**, from the Primary menu, then an **11** from the Resource Report menu (shown in Figure 2-9 on page 2-24) or you can enter following command:

```
STORCR
```

Purpose of the 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.

Contents of the report

```
                     RMF V1R2  Common Storage Remaining       Line 1 of 752
Command ===>                                                 Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  Time: 10.03.20  Range: 100   Sec

                                        Amount of Common Storage
                       Job Ended       Not Released at End of Job
Jobname     ID       Date     Time        CSA  ECSA   SQA  ESQA

%REMAIN                                 78920  523K  3944 54216
CATALOG            11/15/02  05.54.59   62736 12192     0  3808
AUTONET            11/17/02  03.21.53       0  100K     0     0
DCOLCMQS JOB09987  11/20/02  05.47.29       0     0  2648     0
NPM      STC23711  11/17/02  00.09.42    4608     0     0     0
PPCDLYR  JOB09019  11/19/02  05.12.24       0 36912     0     0
DYNPARS  STC23607  11/15/02  05.56.10       0 34024     0     0
DRAGONJ  JOB22261  11/19/02  10.37.07       0 27392     0     0
DRAGONJ  JOB20441  11/19/02  09.26.19       0 27384     0     0
RLW      TSU27070  11/17/02  18.55.46       0 24768     0     0
AUTONET            11/18/02  17.08.02       0 13280     0   128
SARSIM   TSU04303  11/18/02  15.10.51       0 12288     0     0
DFHSM    STC23692  11/17/02  11.59.44       0     0   464     0
```

Figure 2-89. STORCR Report

A graphic report is not available for this report.

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 2-60. 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

How to request this report

To request the Storage Frames report, select a **3**, from the Primary menu, then an **7** from the Resource Report menu (shown in Figure 2-9 on page 2-24) or you can enter the following command using the format:

```
STORF [class,s-class]
```

Purpose of the report

The Storage Frames report contains detailed frame counts, auxiliary slot count, and page-in rate for each address space.

Contents of the report

```
                        RMF V1R2  Storage Frames               Line 1 of 678
Command ===>                                                 Scroll ===> HALF

Samples: 100     System: MVS3  Date: 11/29/02  Time: 10.03.20  Range: 100   Sec

           Service     -Frame Occupancy-  -- Active Frames --   AUX   PGIN ES
Jobname  C Class    Cr TOTAL  ACTV  IDLE  WSET   FIXED    DIV   SLOTS RATE RATE

DB2MDBM1 S SYSSTC   S  64897 64897     0 64897     432      0    3212    0   68
STAWICK  T TSOPRIME    31462     0 31462     0     182      0       0    0    0
ENGD     T TSOPRIME    25719     0 25719     0      40      0       0    0    0
BCOVEY   T TSOPRIME    23697     0 23697     0       0      0       0    0    0
HENNET   T TSOPRIME    21191     0 21191     0     155      0       0    0    0
HOSE     T TSOPRIME    17559     0 17559     0       0      0       0    0    0
DAVEB    T TSOPRIME    15505     0 15505     0     125      0       0    0    0
JHWANG   T TSOPRIME    13327  1334 11993 13339     116      0       0    0    4
SHIH2    T TSOPRIME    10117     0 10117     0       0      0       0    0    0
SYSBMAS  S SYSSTC       9929  9929     0  9929      65      0      24    0    0
SHIH     T TSOPRIME S   7426     0  7426     0       0      0       0    0    0
GRS      S SYSSTC       6419  6419     0  6419      95      0      43    0    6
```

Figure 2-90. STORF Report

A graphic report is not available for this report.

Field descriptions

Table 2-61. 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: **S** Started task **T** TSO **B** Batch **A** ASCH **O** OMVS

Table 2-61. Fields in the STORF Report (continued)

Field Heading	Meaning
Service Class	The name of the service class that a specified job has been running in. **Note:** For systems running in compatibility mode, the fields DMN and PG are shown instead of Service Class, they display the domain and performance group for the specified job.
Cr	An **S** in this column indicates that WLM managed the address space as *storage critical* during the reporting interval.
Frame Occupancy	This field is divided into three categories. **TOTAL** The sum of the ACTV and IDLE frames. The shared page counts are not included in TOTAL. **ACTV** The average number of frames held by the job while it was active. This value represents the average number of active central and expanded storage frames the job used during the report interval. $$ACTV = \frac{\sum \text{Central Storage Frames} + \sum \text{Expanded Storage Frames}}{\text{\# Samples}}$$ **IDLE** The average number of frames held by the job while it was idle. This value represents the average number of central and expanded storage frames the jobs used when it was idle during the report interval. $$IDLE = \frac{\sum \text{Central Storage Idle Frames} + \sum \text{Expanded Storage Idle Frames}}{\text{\# Samples}}$$
Active Frames	This field is broken into three categories. **WSET** See Table 2-55 on page 2-156 for a description of this value. The shared page counts are not included in WSET. **FIXED** 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. $$FIXED = \frac{\sum \text{Fixed Frames}}{\text{\# Samples}}$$ **DIV** The number of central and expanded 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. $$DIV = \frac{\sum \text{DIV Frames}}{\text{\# DIV Samples}}$$
AUX SLOTS	The number of auxiliary slots for each address space.

Table 2-61. Fields in the STORF Report (continued)

Field Heading	Meaning
PGIN RATE	The average number of page-ins per second for an address space. $$\text{PGIN RATE} = \frac{\sum \text{Page-in Counts for Group}}{\text{Resident Time}}$$ 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.
ES RATE	The rate at which pages are being read into central storage from expanded storage. $$\text{ES RATE} = \frac{\sum \text{Expanded Storage Page-in Counts}}{\text{Resident Time}}$$ The ES RATE includes single and blocked pages, but does not include hiperspace or VIO pages. For systems running in 64-bit mode, this field will be blank.

Monitor III Utility fields

You can use the Monitor III Utility to customize the STORF report. The additional fields show the shared storage page measurements.

Table 2-62. Additional Fields in the STORF Report

Field Heading	Meaning
Shared Page-in rate	Shared storage page-in rate from auxiliary storage. The page-in rate is calculated by dividing the total number of shared pages brought in from auxiliary storage by the resident time. (The resident time is the total time an address space was swapped-in in seconds).
Shared page validations	Shared page validations.
Total valid shared pages	Total number of shared pages in central storage that are valid in this address space. "Valid" means that the page has been brought into central storage and has been made addressable for the address space.
Total number of shared pages	Total number of shared pages for the requested range. This is the number of pages owned by the address space that have been shared (whether they were referenced or not).

Report options

The STORF Report Options panel is similar to the DEV Report Options panel. See Figure 2-39 on page 2-73 for an example of the Report Options panel. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed. See Figure 2-37 on page 2-69 for an example.

STORR - Storage Resource Delays report

How to request this report

To request the Storage Resource Delays report, select a **3**, from the Primary menu, then an **8** from the Resource Report menu (shown in Figure 2-9 on page 2-24) or you can enter the following command:

```
STORR
```

Purpose of the report

The Storage Resource Delays report (STORR) provides information about storage problems and paging space delay by volume serial.

Contents of the report

```
                      RMF V1R2  Storage Resource Delays         Line 1 of 7
Command ===>                                                  Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  Time: 10.03.20  Range: 100   Sec

------------------------------ Storage Summary -------------------------------
               AVG HI UIC/    Frames      ------------ % Frames --------------
Storage         MIGR AGE      Online      NUC  SQA  CSA  LPA  ACTV  IDLE  AVAIL
Central           224         212990        1    5    1    1    22    68      0
Expanded         103K         826368             0    0    0    12    69     18
                                                Total SQA + ESQA Overflow:   28K
---------------------------- Page/Swap Activity ------------------------------
Volume DEV      CU       PAV ACT CON DSC PND Pend    SPACE  - AVG Active Users-
Serial Type     Type          %   %   %   %  Reasons TYPE   TOTL LOCL SWAP COMM

PGT859 33903    3990-3         0   0   0   0 None    LOCL    0.0  0.0  0.0  0.0
PGT80C 33903    3990-3         0   0   0   0 None    LOCL    0.0  0.0  0.0  0.0
PGT1B5 33903    3990-3         0   0   0   0 None    LOCL    0.0  0.0  0.0  0.0
                                                    COMM    0.0  0.0  0.0  0.0
                                                    PLPA    0.0  0.0  0.0  0.0
PGT80A 33903    3990-3         0   0   0   0 None    LOCL    0.0  0.0  0.0  0.0
PGT857 33903    3990-3         0   0   0   0 None    LOCL    0.0  0.0  0.0  0.0
```

Figure 2-91. STORR Report

The report has two sections.

Storage Summary
This section includes general information about the use of central and expanded 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.

For systems running in 64-bit mode, all fields that are related to expanded storage will be blank.

Field descriptions

Table 2-63. Fields in the STORR Report - Storage Summary Section

Field Heading	Meaning
AVG HI UIC/ MIGR AGE	*Central storage* - The average high unreferenced interval count (AVG HI UIC) indicates storage contention. *Expanded storage* - The time (in seconds) that a page resides in expanded storage before it migrates to auxiliary storage (MIGR AGE).
Frames Online	*Central storage* - The number of central storage frames, excluding read-only frames. Nucleus frames are included in this value. *Expanded storage* - The number of usable expanded storage frames. Nucleus frames are not included in this value.
% Frames	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 or expanded storage frames during the report interval. The categories are as follows: **NUC** *Central storage* - Percentage of central storage frames allocated to the nucleus (NUC). *Expanded storage* - Nucleus frames do not exist. **SQA** *Central storage* - Percentage of central storage frames allocated to the system queue area (SQA). *Expanded storage* - Percentage of expanded storage frames allocated to the system queue area. **CSA** *Central storage* - Percentage of central storage frames allocated to the common storage area (CSA)l. *Expanded storage* - Percentage of expanded storage frames allocated to the common storage. The number of CSA expanded storage frames is the difference between the total CSA frames and the central storage CSA frames. **LPA** *Central storage* - Percentage of central storage frames allocated to the link pack area (LPA). *Expanded storage* - Percentage of expanded storage frames allocated to the link pack area. The number of LPA expanded storage frames is the difference between the total LPA frames and the central storage LPA frames.

Table 2-63. Fields in the STORR Report - Storage Summary Section (continued)

Field Heading	Meaning
% Frames (continued)	**ACTV** *Central storage* - 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. *Expanded storage* - Percentage of frames allocated to jobs that were active. This value represents the number of expanded storage frames allocated to all active address spaces. It includes idle, using, and unknown time. **IDLE** *Central Storage* - 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. *Expanded storage* - Percentage of expanded frames allocated to jobs that were idle. This value represents the number of expanded storage frames allocated to idle address spaces. See definition of Users in the Workflow/Exceptions report for a description of idle users. **AVAIL** *Central storage* - Percentage of available central frames. *Expanded storage* - Percentage of available expanded storage frames.
Total SQA + ESQA Overflow	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. `Total SQA + ESQA Overflow = Total Overflow / # Samples` **Total overflow** Total overflow above and below 16M

Table 2-64. 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	The number of parallel access volumes (base and alias) which were available at the end of the reporting interval. If the number has been changed during the reporting interval, it is followed by an '*'.
ACT %	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` **PND Time** 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). **CON Time** The time the device was connected to a channel path to actually transfer data between the device and storage. **DSC Time** 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.

Table 2-64. Fields in the STORR Report - Page/Swap Activity Section (continued)

Field Heading	Meaning

CON %	The percent connect time:
	`Connect Time`
	`CON % = ------------ * 100`
	`Range Time`

DSC %	The percent disconnect time:
	`Disconnect Time`
	`DSC % = --------------- * 100`
	`Range Time`

PND %	The percent pending time:
	`Pending Time`
	`PND % = ------------ * 100`
	`Range Time`
	Notes:
	1. 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.
	2. 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.

Pend Reasons	The reason for the delay and the percentage of delay.
	DB 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.
	`Accumulated DB Delay Time`
	`DLY DB% = ------------------------ * 100`
	`Range Time`
	CMR 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.
	`Accumulated Command Response Delay Time`
	`DLY CMR% = ------------------------------------ * 100`
	`Range Time`
	Note: If either hardware data or volume related percentages are not available, this field is blank.

SPACE TYPE	The space types for which the percentage of the volume's delay is reported. The types appear in the following order: **LOCL** User private area **COMM** Common area **PLPA** Pageable link pack area

Table 2-64. Fields in the STORR Report - Page/Swap Activity Section (continued)

Field Heading	Meaning
AVG Active Users	The average number of jobs waiting for the data set. $$\text{AVG Active Users} = \frac{\sum \text{Waiting Jobs}}{\text{\# Samples}}$$ **Sum of all waiting jobs** Sum of all delay samples for all jobs waiting for the data set. This category is divided into: **TOTL** 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 %. **LOCL** The percentage that local (private) storage paging contributes to the delay from the time of the page fault until I/O is completed. **SWAP** The percentage that swapping contributes to the delay from the time of swap initiation until the last swap page I/O is completed. **COMM** 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. 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.

Monitor III Utility fields

You can use the Monitor III Utility to customize the STORR report. In addition to the delays shown above you can use the Utility to have the following delay percentages shown.

Table 2-65. 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

How to request this report

To request the Storage Delay Summary report, select a **3**, from the Primary menu, then an **9** from the Resource Report menu (shown in Figure 2-9 on page 2-24) or you can enter the following command using the format:

```
STORS wlm
```

Purpose of the report

This Storage Delay Summary (STORS) report provides you with an overview of storage usage by service classes, report classes, and workload groups (or by domain and performance group).

Contents of the report

```
                        RMF V1R2  Storage Delay Summary          Line 1 of 22
Command ===>                                                   Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  Time: 10.03.20  Range: 100   Sec

------------------------------ Storage Summary -------------------------------
               AVG HI UIC/    Frames      ------------ % Frames --------------
Storage         MIGR AGE      Online      NUC  SQA  CSA  LPA  ACTV  IDLE  AVAIL
Central           224         212990        1    5    1    1    22    68      0
Expanded         103K         826368             0    0    0    12    69     18
                                                 Total SQA + ESQA Overflow:   28K
Group    T  -- Users --  - Average Number Delayed For-  - Average Frames-  PGIN
             TOTL  ACTV   ANY COMM LOCL SWAP OUTR OTHR   ACTV  IDLE FIXED  RATE

PRIMEAPP W      3     0     0    0    0    0    0    0   1442     0     0   0.0
APPPRIME S      3     0     0    0    0    0    0    0   1442     0     0   0.0
         1      2     0     0    0    0    0    0    0    736     0     0   0.0
         2      1     0     0    0    0    0    0    0    706     0     0   0.0
PRIMEBAT W     11    10   0.5    0    0    0    0  0.5   6016     0   330   2.4
NRPRIME  S     11    10   0.5    0    0    0    0  0.5   6016     0   330   2.4
         1      9     9   0.5    0    0    0    0  0.5   4121     0   261   2.4
         2      0     0     0    0    0    0    0    0    509     0    18   0.0
         3      1     1     0    0    0    0    0    0   1385     0    51   0.0
PRIMETSO W    527     8   2.0    0    0    0    0  2.0  17932  717K  9653   0.0
TSOPRIME S    527     8   2.0    0    0    0    0  2.0  17932  717K  9653   0.0
         1    526     8   2.0    0    0    0    0  2.0  15780  717K  9577   0.0
         2      1     1     0    0    0    0    0    0   1955     0    67   0.0
         3      0     0     0    0    0    0    0    0    197     0     9   0.0
PRIMOMVS W      1     0     0    0    0    0    0    0    768   168   183   0.0
OMVS     S      1     0     0    0    0    0    0    0    768   168   183   0.0
```

Figure 2-92. STORS Report

The top section on the report provides overall system information and is the same as the 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, OUTR, and OTHR.

For systems running in 64-bit mode, all fields that are related to expanded storage will be blank.

Field descriptions

Table 2-66. Fields in the STORS Report

Field Heading	Meaning
Storage Summary	Fields in this section are described in Table 2-63 on page 2-169.
Group	The name of the group, including: • Workload group names • Service class names • Report class names **Note:** For systems running in compatibility mode, the fields DMN and PG are shown instead of Group, they display the domain and performance group.
T	A one-character abbreviation for the type of workload manager group as follows: **W** Workload group name **S** Service class name **R** Report class name **n** Service/report class period
Users	The number of users within the group. This category includes the following headings: **TOTL** The total number of users equals the number of different users found in all address spaces for the group listed during the report interval. **ACTV** The average number of active users is a measure of system workload. See "WFEX - Workflow/Exceptions report" on page 2-213 for the definition of User/Active.
Average Number Delayed For	The average number of delayed users is summarized for the following categories: **ANY** Delay the group experienced because of contention for any of the following measured storage reasons during the report interval. **COMM — LOCL — SWAP — OUTR — OTHR** For descriptions of these delays, see the corresponding field in the STOR report (Table 2-55 on page 2-156).
Average Frames	The average number of storage frames the group held during the report interval. This field reports on the following frame categories: **ACTV — IDLE** See Table 2-61 on page 2-165 for a description of these counts. **FIXED** The average number of fixed frames the job was using during the report interval including frames both above and below the 16 megabyte line. $$\text{Avg Fixed Frames} = \frac{\sum \text{Fixed Frames}}{\text{\# Samples}}$$
PGIN RATE	The rate at which pages are being read into central storage. $$\text{PGIN RATE} = \frac{\sum \text{All Page-in Counts for Group}}{\text{Resident Time}}$$ The address-space related shared storage page-ins are included in the PGIN RATE.

Monitor III Utility fields

You can use the Monitor III Utility to customize the Storage Delay Summary report. In addition to the delays shown above, you can use the Utility to have the delays in Table 2-67 shown in the Storage Delay Summary report.

Table 2-67. 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 2-93. 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

```
Sel  Group    T   Sel  Group   T
 S   CICS*___ S    S   IMS*___ S
     ________ _        ________
 X   *ALL
```

You can use the wild card to select by type, for example, to list service classes only, specify:

```
Sel  Group    T   Sel  Group   T
 S   *_______ S        _______ _
     ________ _        _______ _
 X   *ALL
```

T - type Type can be:

W Workload group name
S Service class name
R Report class name

Note: For systems running in compatibility mode, the Report Options panel contains a selection list with the available domains and performance groups.

SYSENQ - Sysplex Enqueue Delays report

How to request this report

To request the SYSENQ report, select an **S** on the Primary menu, and then a **4** on the Sysplex Report menu (shown in Figure 2-6 on page 2-21), or you can enter the following command:

```
SYSENQ
```

Purpose of the report

The SYSENQ report is similar to the ENQR report (see "ENQR - Enqueue Resource Delays report" on page 2-103), 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.

Contents of the report

```
                    RMF V1R2  Sysplex ENQ Delays - RMFPLEX1     Line 1 of 10
Command ===>                                                  Scroll ===> HALF

Samples: 114     Systems: 3    Date: 11/29/02  Time: 12.58.30  Range: 100   Sec

------ Resource Name ------ ------- Delayed --------  ------- Holding --------
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
 SYSMCS#CL2
SYSZMCS                      20 CONSOLE  SYS1     EW   18 CATALOG  SYS1     EO
 SYSMCS#CL1                                             2 CATALOG  SYS1     EO
DSNJBSDS                     99 S412MSTR SYS1     EW   99 S411MSTR SYS1     EO
 MODIFY
```

Figure 2-94. SYSENQ Report

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

Mon III - SYSENQ report

Field descriptions

Table 2-68. Fields in the SYSENQ 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.
	`# Delay Samples`
	`Delayed % = --------------- * 100`
	`# Samples`
	Delay samples
	The number of samples when the job was delayed for a specific enqueued resource.

-- 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 MVS 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.

-- Holding -- %	The percent of the range that a specific job was holding the resource while the named job was delayed.
	`# Holding Samples`
	`Holding % = ---------------- * 100`
	`# Samples`
	Holding samples
	The number of samples when the holding job was holding the resource while the named job was delayed.
	Because more than one job can hold the resource at a time, these values can add up to more than 100%.

-- Holding -- Jobname	The name of the job that is holding the resource that the delayed job is waiting for. 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 +.
-- Holding -- Sys-Name	The MVS system name where the job is running on.
-- Holding -- ST	The status indicates whether the holding job has exclusive (EO) or shared (SO) use of the resource.

SYSINFO - System Information report

How to request this report

To request the System Information report, select a **1** from the Primary menu, then a **2** from the Overview Report menu (shown in Figure 2-7 on page 2-22) or you can enter the following command using the format:

```
SYSINFO wlm
```

Purpose of the 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 (or domain, performance group, or report performance group), and the total number of jobs using resources or delayed for resources.

Contents of the report

```
                        RMF V1R2  System Information            Line 1 of 28
Command ===>                                                  Scroll ===> HALF

Samples: 100     System: MVS3  Date: 11/29/02  Time: 10.03.20  Range: 100   Sec

   ---------------------- 9672 Model RX4 ---------------------- Policy: STANDARD
   Average CPU Util%:    73          Appl% / EAppl%:    63/ 65 Date:   11/29/02
   Avg. MVS CPU Util%:   84          Partition Name:  MVS1     Time:   08.00.11

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      34  664   26       13.95  5.1  5.0    1.9  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        0.00  0.0  0.0    0.0  0.0  0.0  0.0  0.0  0.0
*OMVS             2    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
PRIMETSO  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 2-95. 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 number of processors and vector processors online, 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 MVS 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 5-41.

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 (or domains (DMN), performance groups (PG), and report performance groups (RPG)).

Note: It might be possible that there is enclave activity in your system (for example, indicated by EAppl% > Appl%), and summary line *ENCLAVE is showing 0 users. The reason is that only those enclaves are shown in the report that have been sampled at least twice and that are active at the end of the Monitor III mintime. Therefore, short-running enclaves will not appear in the report.

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.

Notes:

1. RPG and 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 RPG and report class lines.

Field descriptions

Table 2-69. Fields in the SYSINFO Report

Field Heading	Meaning
Processor	Processor familiy and model (or N/A — if model information is not available). If the processor does not support the Concurrent Processor Upgrade function, the version number (in hexadecimal) will be displayed.
Average CPU Util%	The average CPU utilization percentage for all processors is also displayed on the WFEX report. See Table 2-78 on page 2-217 for the calculation. If data is missing or not valid, this field contains ’***’.
Appl%	Percentage of CPU time used by all address spaces during the report interval. This value is divided by the number of processors that have been active during this interval.
EAppl%	Percentage of CPU time used by all address spaces and enclaves during the report interval. This value is divided by the number of processors that have been active during this interval.
Avg. MVS CPU Util%	MVS view of processor utilization.
Partition Name	Partition name.
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. **Note:** For systems running in compatibility mode, this block of data displays the names of the Parmlib members IEAIPSxx, IEAOPTxx, and IEAICSxx.

Table 2-69. Fields in the SYSINFO Report (continued)

Field Heading	Meaning
Group	The name of a class (*SYSTEM, *TSO, *BATCH, *STC, *ASCH, or *OMVS), an enclave (*ENCLAVE), or a group, including: • Workload group names • Service class names • Report class names **Note:** In compatibility mode, domains, performance groups, or report performance groups are shown under the header `Name`.
T	Type of workload manager group: **W** Workload group name **S** Service class name **R** Report class name **n** 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: **TOT** The total number of users. **ACT** The average number of active users. See the definition of Users/Active under "WFEX - Workflow/Exceptions report" on page 2-213 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 • K - times one thousand seconds • M - times one million seconds **Note:** 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. `TRANS /SEC = Completed Transaction Count / 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: **PROC** Average number of users using the processor during the report interval. **DEV** Average number of users using devices during the report interval.

Table 2-69. Fields in the SYSINFO Report (continued)

Field Heading	Meaning
Average Number Delayed For	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. The average number delayed for is reported for the following categories: **PROC** Number of users experiencing delay because of contention for the processor during the report interval. **DEV** Number of users experiencing delay because of contention for the devices during the report interval. **STOR** Number of users experiencing delay because of contention for storage during the report interval. **SUBS** Number of users experiencing delay because of contention for JES, HSM, or XCF during the report interval. **OPER** 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. **ENQ** Number of users experiencing delay because of contention for an enqueued resource during the report interval.

Monitor III Utility fields

You can use the Monitor III Utility to customize the SYSINFO report. In addition to the information shown above you can use the Utility to have the following values shown.

Table 2-70. Additional Fields in the SYSINFO Report

Field Heading	Meaning
Execution Velocity	The value for execution velocity 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.
Active frames percentage	The average number of frames held by the job while it was active. This value represents the average number of active central and expanded storage frames the job used during the report interval.
Average number delayed JES	The average number of users experiencing delay when requesting service from JES.
Average number delayed HSM	The average number of users experiencing delay when requesting service from HSM.
Average number delayed XCF	The average number of users experiencing delay when requesting service from XCF.
Average number delayed Mount	The average number of users experiencing delay because of an operator mount request.
Average number delayed Message	The average number of users experiencing delay because of an operator message request.
Percentage of CPU time (TCB+SRB) used	Percentage of processor time used by a class or group.
Percentage of SRB time used	Percentage of SRB time used by a class or group.
Percentage of TCB time used	Percentage of TCB time used by a class or group.

Table 2-70. Additional Fields in the SYSINFO Report (continued)

Field Heading	Meaning
Vector utilization	Contains the vector utilization percentage for a processor with a vector processor online. This value is shown directly in the report when the system is running in compatibility mode.

Report options

The Report Options panel is exactly the same as for the SYSSUM report and STORS report, shown in Figure 2-93 on page 2-175. Selections made on either options panel affect all three reports.

SYSRTD - Response Time Distribution report

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 2-6 on page 2-21) or you can enter the following command using the format:

```
SYSRTD serviceclass, period
SYSRTD reportclass, period
```

For example, to get a Response Time Distribution report for the service class POSMULTI and service class period 1, enter:

```
SYSRTD POSMULTI, 1
```

Purpose of the 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.

Contents of the report

```
                      RMF V1R2  Response Time - RMFPLEX1        Line 1 of  5
Command ===>                                                  Scroll ===> HALF

WLM Samples: 100     Systems: 8  Date: 11/29/02 Time: 10.31.40 Range: 100   sec


Class : POSMULTI  %  40|
Period: 1         of  |
                  TRX |
Goal:                 |
0.500 sec for 90%    20|
                       |
                       |
                       |                    XXXX
                       |                    XXXX
                     0|-----//---+---+---+---+---+---+---+---+---+//---- Time
                       <0.30   0.350       0.500           0.700   >0.75(sec)

                --Avg. Resp. Time--   Trx     --Subsystem Data--  --Exec Data--
System   Data   WAIT  EXECUT ACTUAL   Rate    Actv  Ready Delay   Ex Vel  Delay

*ALL            0.208  0.379  0.587   124.7      9     4    54
MVS1     all    0.061  0.311  0.372    84.3      7     8    74
MVS2     all    0.149  0.984  1.134    40.4      2     0    42
MVS3     part   0.0    0.0    0.0       0.0      0     0     0
MVS5     none
```

Figure 2-96. 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.
- The details of 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 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 V1R2  Response Time - RMFPLEX1         Line  1 of  5
Command ===>                                                  Scroll ===> HALF

WLM Samples: 100     Systems: 7  Date: 11/29/02 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.




                --Avg. Resp. Time--   Trx     --Subsystem Data--  --Exec Data--
System   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 2-97. 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 2-191 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 that 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 area that makes up the chart represent 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 chart 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 goal mode 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 2-57 on page 2-107) 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 2-71. 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 2-185.
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. **all** Data could be retrieved that covers the report interval shown in the report header. **part** Data was retrieved that contains at least one time gap within the report interval. **none** 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 ***ALL**). Refer to "SYSSUM - Sysplex Summary report" on page 2-189 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 ***ALL**).
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 ***ALL**). Refer to "SYSWKM - Work Manager Delays report" on page 2-204 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 *Exec Vel - Actual* on the Sysplex Summary report with the only difference that the value here is only calculated for a single system (except ***ALL**). Refer to "SYSSUM - Sysplex Summary report" on page 2-189 for a description.
Execution Data Delay	This is the general execution delay used for the execution velocity calculation. Note that in a service group, more than one transaction can be delayed at the same point of time. So if for example, two transactions in average are delayed each time RMF takes a measurement sample, then 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 2-58 on page 2-114, 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

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 2-6 on page 2-21), or you can enter the following command using the format:

```
SYSSUM wlm
```

Purpose of the 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

The report provides an overview of workload groups, service classes, service class periods, report classes, and report class periods. It allows ″summarizing″ of actual goal mode 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.

Contents of the report

```
                    RMF V1R2 Sysplex Summary - RMFPLEX1       Line  1 of 35
Command ===>                                                 Scroll ===> HALF

WLM Samples: 100   Systems:  8 Date: 11/29/02 Time: 10:31:40 Range: 100   Sec
                                                            Refresh: 100   Sec

 ----|||||--|----X--XX-XXXX-XX-XX--||---XXXXX---|--||||||||--XXXX--X--XX----XX

Service Definition: SLA_2002              Installed at: 07/01/02, 06.02.00
     Active Policy: WORKDAYS              Activated at: 07/01/02, 23.17.54

               ------- Goals versus Actuals --------  Trans --Avg. Resp. Time-
               Exec Vel  --- Response Time ---  Perf  Ended  WAIT EXECUT ACTUAL
Name     T  I  Goal Act  ---Goal--- --Actual--  Indx  Rate   Time  Time  Time

ATM&POS  W           N/A                              168.5  0.152 0.459 0.611
ATMLABOR S  1        N/A   0.80  90%        88%   1.1   44.4  0.000 0.700 0.700
POSMULTI S  1        N/A   0.50  90%        85%   1.2  124.7  0.208 0.379 0.587
BATCH    W           75                               18.3  3.11H 24.8M 3.52H
BATCHP1  S           61                               11.2  4.08M 9.26M 13.3M
         1  3        65   10.0M AVG  24.5M AVG   2.5    8.5  5.93M 18.6M 24.5M
         2  D        40                               2.7  12.3M 9.83M 22.1M
BATCHP2  S  D        89                               7.1  7.41H 48.2M 8.21H
CICSSTC  W           68                               0.0  0.000 0.000 0.000
CICSTOR  S  1    90  12                         7.5    0.0  0.000 0.000 0.000
CICSAOR  S  1    90  57                         1.6    0.0  0.000 0.000 0.000
CICSDOR  S  1    90  99                         0.9    0.0  0.000 0.000 0.000
SYSTEM   W           84                               0.0  0.000 0.000 0.000
SYSSTC   S      N/A  84    N/A                        0.0  0.000 0.000 0.000
TSO      W           63                               6.4  0.608 1.933 2.541
DEVELOP  S           63                               6.4  0.073 1.360 1.433
         1  1        65   1.00  AVG  0.80  AVG   0.8    3.9  0.000 0.797 0.797
         2  3        60   3.00  90%        92%   1.0    2.3  0.000 2.991 2.991
```

Figure 2-98. 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.

```
        Report Interval

Time = 12:00:00    :03:20    :06:40    :10:00    :13:20
            |---------|-x-------|-x-------|---------|
                      ↑         ↑
                      |---------|  (first report interval)

                           ↑         ↑
                           |---------|  (second report interval)

                                ↑         ↑
                                |---------|  (third report interval)
```

If you end GO mode, and do not change the Refresh or range values, his torical 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 2-98 on page 2-190 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:

```
                    RMF V1R2 Sysplex Summary - RMFPLEX1       Line  1 of 35
WLM Samples: 100   Systems:  8 Date: 11/29/02 Time: 10:26:40 Range: 100   Sec

              >>>>>>> ------------------ <<<<<<
```

Figure 2-99. 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 TSODEPTA, TSODEPTB, and TSODEPTC. 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 TSODEPTA and TSODEPTB. 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 2-72 on page 2-194 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.

For each report class (indicated by the type R), one line is shown containing the same data as for a service class, that means *Execution Velocity - Actual*, *Trans Ended Rate*, and *Avg. Resp. Time*.

In the *Goal* columns, **N/A** will be shown, and the *Importance* column will remain blank. Furthermore, a period breakdown does not exist, because periods cannot be defined for a report class.

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 2-100 on page 2-194.

```
                    RMF V1R2 Sysplex Summary - RMFPLEX1       Line  1 of  5
Command ===>                                                  Scroll ===> HALF

WLM Samples: 100   Systems:  8 Date: 11/29/02 Time: 10:31:40 Range: 100   Sec
                                                            Refresh: 100   Sec

----|||||--|----X--XX-XXXX-XX-XX--|---XXXXX---|--|||||||--XXXX--X--XX----XX


Service Definition: SLA_2002              Installed at: 07/01/02, 06.02.00
     Active Policy: WORKDAYS              Activated at: 07/01/02, 23.17.54

               ------- Goals versus Actuals --------  Trans --Avg. Resp. Time-
               Exec Vel  --- Response Time ---  Perf  Ended  WAIT EXECUT ACTUAL
Name     T  I  Goal Act  ---Goal--- --Actual--  Indx  Rate   Time  Time  Time

ATM&POS  W          N/A                                168.5  0.152 0.459 0.611
BATCH    W           75                                 18.3  3.11H 24.8M 3.52H
CICSSTC  W           68                                  0.0  0.000 0.000 0.000
SYSTEM   W           84                                  0.0  0.000 0.000 0.000
TSO      W           63                                  6.4  0.608 1.933 2.541
```

Figure 2-100. SYSSUM Report for Workload Groups

Field descriptions

Table 2-72. 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 goal mode 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 goal mode 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. **W** Workload group **S** Service class **R** Report class **n** 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 **D** in this column, since an importance cannot be defined for a discretionary goal. **1** Highest - describes highest priority service class period for most important work **2** High **3** Medium **4** Low **5** Lowest **D** Discretionary

Table 2-72. Fields in the SYSSUM Report (continued)

Field Heading	Meaning
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 *Response Time* goal and actual columns remain empty. If a response time goal with a percentile was specified, the *Actual* 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 *Avg. Resp. Time*. 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.
Execution velocity Actual	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. A high value indicates little workload contention while a low value indicates that the requests for system resources are delayed. See "Execution velocity" on page 2-10 for details about the execution velocity. Whenever subsystem delays are available for that service class period, **N/A** is shown in this field. The Work Manager Delays report provides more information for these service class periods. 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. In the *Execution Velocity - Actual* column, two exceptions may occur: • N/A appears for "transaction" service classes, that means, for classes served by "server" service classes. For these classes, the execution velocity cannot be defined meaningfully. • 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.
Time	The time units shown in the columns *Goal versus Actuals - Response Time* and *Avg. Resp. Time* are all seconds, if no unit is marked. Time values bigger than one minute are marked with **M**, and time values bigger than one hour are marked with **H**.
Response Time Goal	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: • The average target response time for all ended transactions. • The percentage of transactions that should terminate within the time specified in the goal. This percentage is specified together with the response time value. For a goal without percentage, **AVG** is shown in this field. This field shows **N/A** for system service classes, since they do not have a user defined goal, and it is empty for heterogeneous report class periods.
Response Time Actual	• Average response time goal: The value represents the average response time for all ended transactions, followed by **AVG**. • Response time goal with percentile: The percentage of transactions that actually ended within the time specified in the goal.

Table 2-72. Fields in the SYSSUM Report (continued)

Field Heading	Meaning
Performance Index	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. If RMF cannot calculate the performance index, this field contains **N/A**: • For a period with a response time goal: The sum of completed transactions is zero • 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) This field also shows **N/A** for report classes and system service classes, since they do not have a user defined goal. The *Performance Index* field can also be shown in dark blue. This has the same reason as for the *Execution Velocity - Actual* column, described above. RMF calculates the performance index depending on the type of goal: **Execution velocity goal** $$\text{Perf Indx} = \frac{\text{Goal \%}}{\text{Actual \%}}$$ **Average or percentile response time goal** $$\text{Perf Indx} = \frac{\text{Actual (sec)}}{\text{Goal (sec)}}$$ "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: 1. Calculate the number of transactions N that correspond to the goal: $$N = \frac{\sum \text{Transactions} * \text{Goal Percentage}}{100}$$ 2. Add up all transactions until a bucket M is reached where the sum is greater than N. 3. The "actual" response time in the formula for the performance index shown above is the response time value belonging to the bucket M. **Note:** 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 '****'.

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 below 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:

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 2-73. Fields in the SYSSUM Report - Continuation

Field Heading	Meaning
Trans Ended Rate	The number of transactions ended per second.
Avg. Resp. Time WAIT Time	The average time (in seconds) that a transaction spent waiting because of one of these reasons: • 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. • 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. • Ineligible - Operational or JES scheduling delay Average time a job was delayed due to operational delays or JES scheduling delays, examples are: – Job held by operator – Job class or job queue held – Duplicate jobname serialization – Job class execution limits • 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. The time a job was delayed due to TYPRUN=HOLD or TYPRUN=JCLHOLD is NOT included in any of the transaction times. 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. In the *Avg. Resp. Time* 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.
Avg. Resp. Time EXECUT	For CICS transactions, this includes execution time in AOR and following regions. For IMS transactions, this includes execution time within the MPR. For Batch, TSO, etc., this is the average time that transactions spent in execution. In the *Avg. Resp. Time* 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.

Table 2-73. Fields in the SYSSUM Report - Continuation (continued)

Field Heading	Meaning
Avg. Resp. Time ACTUAL Time	In general, this is the sum of the wait and execution times described above, but does not include ineligible time. For CISC 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 tend to be longer than the local transaction, EXECUT could be greater than ACTUAL. 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. In the *Avg. Resp. Time* 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.

Monitor III Utility fields

You can use the Monitor III Utility to customize the SYSSUM report. In addition to the values shown above you can use the Utility to have the following information shown.

Table 2-74. Additional Fields in the SYSSUM Report

Field Heading	Meaning
Goal type	This goal type indication includes a list of all different types of goals. **1** Percentile response time goal **2** Average response time goal **3** Velocity goal **4** Discretionary goal
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	Minimum service rate, in unweighted CPU service units per second, as defined for that resource group. **Note:** The minimum and maximum values must not be specified by the user.
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 2-93 on page 2-175. 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

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.

Purpose of the 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.

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 reporting 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.

Notes:

1. 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.
2. The SRB% and TCB% columns are empty if a domain has been selected for reporting.

```
                RMF V1R2  *SYSTEM   Trend                     Line 1 of 20
Command ===>                                                 Scroll ===> HALF

Samples: 100     System: AQTS  Date: 03/30/02  Time: 14.30.00  Range: 100 Sec

  Latest:    03/30/02 at 14.30.00   Range/Line:      100 Sec
  Earliest:  03/30/02 at 13.58.20   Total Range:    2000 Sec   00.33.20

          CPU  SRB   TCB   WFL -Users-- -AVG USG-  -Avg. Act. Users Delayed For-
Time       %    %     %     %  TOT  ACT PROC DEV   PROC DEV  STOR SUBS OPER ENQ

14.30.00  39   2.5  31.0  45  623   11  2.5  2.8   1.5  2.4  0.0  1.4  1.3  0.0
14.28.20  53   2.5  41.9  52  623   13  3.6  4.4   1.0  4.2  0.0  0.9  1.2  0.0
14.26.40  50   2.5  39.8  52  623   14  3.8  4.6   1.8  4.8  0.0  0.1  0.7  0.3
14.25.00  35   2.2  26.5  38  621   15  2.0  4.2   1.1  4.5  0.0  1.1  0.6  3.0
14.23.20  35   2.2  26.3  35  620   16  2.2  4.2   0.8  4.4  0.0  1.5  1.4  4.0
14.21.40  39   2.2  30.7  39  620   16  2.6  4.5   0.8  3.3  0.0  1.6  0.2  5.0
14.20.00  40   2.2  32.1  51  620   10  2.9  3.6   0.8  3.0  0.0  0.9  1.5  0.0
14.18.20  41   1.9  33.8  52  620    9  2.5  3.0   0.6  1.7  0.0  1.7  1.1  0.0
14.16.40  49   2.2  39.9  53  620   11  3.2  3.4   1.0  2.9  0.0  0.8  0.7  0.5
14.15.00  54   2.2  45.4  54  623   12  3.4  3.6   0.8  2.5  0.0  1.2  0.4  0.9
14.13.20  61   2.0  52.2  48  623   12  3.8  2.7   1.0  2.1  0.0  1.6  1.5  1.0
14.11.40  69   2.7  57.2  59  623   12  4.3  3.9   1.8  2.6  0.0  0.7  0.3  0.2
14.10.00  59   2.2  50.7  57  623   11  3.8  3.4   1.1  2.7  0.0  0.7  0.9  0.0
14.08.20  64   2.5  54.1  62  623   10  4.0  3.4   1.1  1.7  0.0  0.4  1.1  0.2
14.06.40  52   1.9  44.2  62  622    8  3.4  2.5   1.0  1.2  0.0  0.1  1.3  0.0
14.05.00  48   2.0  39.7  67  619    9  3.2  3.6   0.8  1.1  0.0  0.2  1.2  0.0
14.03.20  58   1.9  50.6  67  618    9  3.6  3.3   0.7  1.5  0.0  0.7  0.5  0.0
14.01.40  54   2.1  44.4  57  620   12  3.5  4.4   0.7  2.0  0.0  2.1  1.2  0.0
14.00.00  63   2.6  50.8  50  619   14  4.0  4.5   2.3  2.9  0.0  2.0  1.2  0.0
13.58.20  63   2.6  50.8  50  619   14  4.0  4.5   2.3  2.9  0.0  2.0  1.2  0.0
```

Figure 2-101. SYSTREND Report

Field descriptions

Table 2-75. 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 WFEX report. See Table 2-78 on page 2-217 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 2-69 on page 2-180).

Cursor-sensitive control

Cursor-sensitive control allows you to navigate to the SYSINFO or a delay report for a selected reporting range.

Table 2-76. SYSTREND Report - Cursor-sensitive Control for Navigation

Report Column where Cursor-sensitive Control is Used	*Displayed Report*
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

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 2-6 on page 2-21) or you can enter the following command using the format:

```
SYSWKM serviceclass,period
SYSWKM reportclass,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
```

Purpose of the 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 2-102 on page 2-205 and Figure 2-103 on page 2-205 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 a possible.

Contents of the report

```
                RMF V1R2  Work Manager Delays - RMFPLEX1         Line 1 of 3
Command ===>                                                  Scroll ===> HALF

WLM Samples:   100   Systems: 2 Date: 03/25/02 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-Usg I/O-Usg  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   SYSSTC       48        7       1      73     0      0
```

Figure 2-102. 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 MVS 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 V1R2  Work Manager Delays - RMFPLEX1         Line 1 of 6
Command ===>                                                  Scroll ===> HALF

WLM Samples:     100 Systems:  2 Date: 03/25/02 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 2-103. 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 MVS 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 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 2-77. 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 *nnnnn* TRX.	*nnnnn* 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 *nnnnn* TRX.	*nnnnn* is the total number of transactions that completed their execution phase during this report interval.
Abnormally ended	The number of abnormally ended transactions from all reported systems. This value is not included in the number of total completed transactions.

Table 2-77. Fields in the SYSWKM Report (continued)

Field Heading	Meaning
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** This line represents states of the begin-to-end phase of the transactions. **X** This line represents states of the execution phase of the transactions.
Response time breakdown	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`). 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. **Tot** 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. The value is a sum of all the figures shown in this row in the other ″Response time breakdown″ columns. **Note:** Because one transaction can be counted in more than one state during a report interval, this number can be larger than 100. **Act** Time spent in an active state. 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). 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. **Rdy** Time spent in a ready state. 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. **Idle** Time spent idle means that no work request (or transaction) is available to be run by the work manager.

Table 2-77. Fields in the SYSWKM Report (continued)

Field Heading	Meaning
Response time breakdown Delayed by	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. **LOCK** Time spent waiting for a lock. **I/O** Time spent of waiting for I/O. 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. **CONV** Time spent waiting for conversation. **DIST** Time spent waiting for distributed request state samples. 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. **SESS** Time spent waiting for a session to be established. This is a sum of the time spent waiting for sessions to be established locally (for example, on the current MVS image), somewhere in the network, or somewhere in the sysplex. **TIME** Time spent waiting for a timer. **PROD** Time spent waiting for another product. **LTCH** Time spent waiting for a latch. **MISC** Time spent waiting for an unidentified resource. **SSLT** Time spent waiting for an SSL thread. **REGT** Time spent waiting for a regular thread. **WORK** Time spent waiting for registration to a work table.
Switched Time (%)	Percentage of time that transactions spent routed to another region for processing. This percentage also refers to the *Average Response Time* shown in the sub-header. For a begin-to-end phase, the sum of these percentages should approximately equal the value shown in the *CONV* column. For an execution phase, these percentages, as well as the figure in the *CONV* column, are expected to be zero. **LOC** Percentage of time that transactions spent switched on this MVS image. Subsystems might set this state when they function ship a transaction to another component within the same MVS image. **SYS** Percentage of time that transactions spent switched to another MVS image in the sysplex. Subsystems might set this state when they function ship a transaction to another component on another MVS image within the sysplex. **REM** 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.

Table 2-77. Fields in the SYSWKM Report (continued)

Field Heading	Meaning
Adress Spaces Serving this Service Class *srvcls*	*Srvcls* is the name of the *Class* at the top of the report. 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. 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.
Jobname	Jobname of the server.
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	Capping delay is defined as
	`# Capping Samples`
	`Capp = ----------------`
	`# Samples`
	
	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 *z/OS MVS Planning: Workload Management*).

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     CICSPS     CICSREGS   CICSSC     CICSTRX    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 2-104. 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.

WFEX - Workflow/Exceptions report

How to request this report

To request the Workflow/Exceptions, select a **1** from the Primary menu, then an **1** from the Overview Report menu (shown in Figure 2-7 on page 2-22) or you can enter the following command:

```
WFEX
```

Purpose of the 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.

Using the information given in the report

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 derived from the IPS (if you have customization set to YES) or 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.

Contents of the report

```
                    RMF V1R2 Workflow/Exceptions                Line 1 of 12
Command ===>                                                   Scroll ==> HALF

Samples:    100 System: MVS1  Date: 11/29/02  Time: 08.53.20     Range:  100 Sec

 --------------------------- Speed (Workflow) ---------------------------------
            Speed of 100 = Maximum, 0 = Stopped         Average CPU Util:  81 %
 Name       Users Active      Speed          Name       Users Active      Speed
 *SYSTEM      505     13         54          *DEV         117     29        100
 ALL TSO      433     10         55          *MASTER*       1      0          0
 ALL BATCH      2      0         42          TSOPROD      420      8         57
 ALL STC       70      2         55          BTCHPROD       2      0         42
 ALL ASCH                  Not avail
 ALL OMVS       3      0        100
 *PROC        129     14         78

 ----------------------------- Exceptions -----------------------------------
 Name       Reason        Critical val. Possible cause or action
 *SLIP*     SLIP PER TRAP               SLIP ID=SR01 is active.
 BEVK       Rate < 2.0    1.220 /sec    Tx rate is 1.220 /s.
 BSHR       STOR-COMM     23.1 % delay
 CSAHOG     JCSA% > 15    18.3 %        Job CSA usage 18 %, system 57 %.
 POK063     DAR > 20      23.22 /sec    I/O rate is 23.22 /s on volume POK063.
 DATAPK     Not avail                   Volume DATAPK is not mounted.
 BBUSER02   PINRT > 2     3.410 /sec    Page-in rate is 3.410 /s.
```

Figure 2-105. 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 2-105 is an example of the Workflow/Exceptions report. For information about the WFEX Report Options panels, see “Report options” on page 2-221. 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 2-231.

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 2-6.

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.

Workflow/Exceptions graphic report

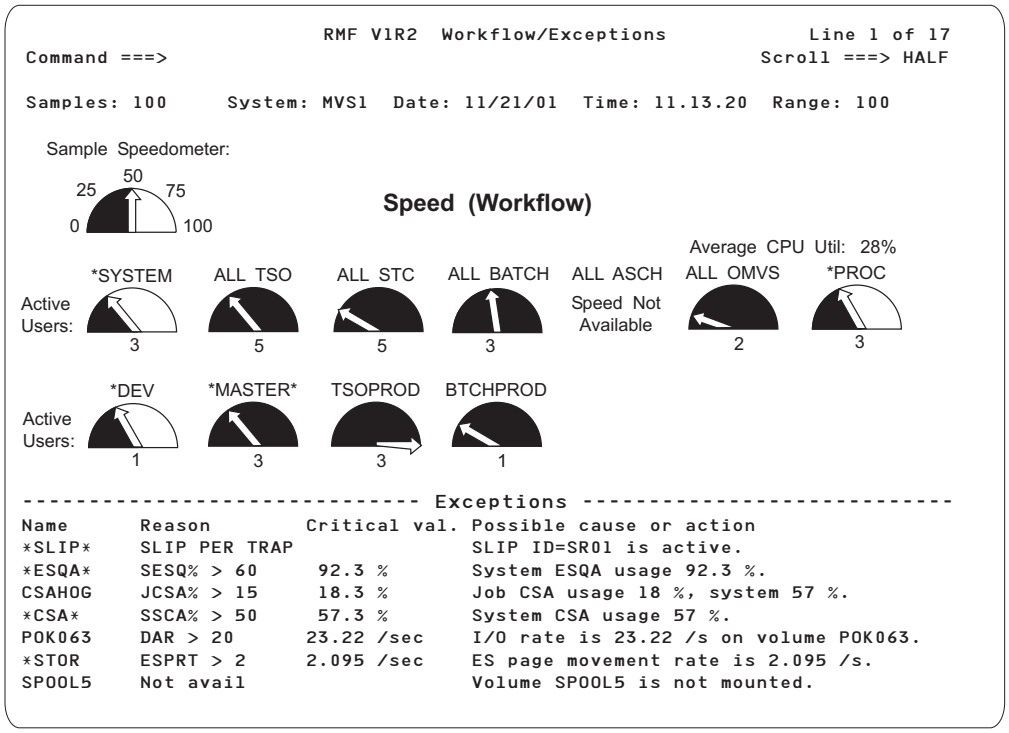

Figure 2-106. WFEX Graphic Report

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

The Workflow/Exceptions report fields are described in Table 2-78 and Table 2-79.

Table 2-78. Fields in the Speed Section of the WFEX Report

Field Heading	Meaning
Average CPU Util	The average utilization percentage for all processors. RMF obtains the accumulated processor wait time at the beginning and end of the report interval. The time of a processor is computed as the range time minus the processor wait time. RMF calculates the average CPU busy time as the accumulated processor time for all processors divided by the number of online processors. $$\text{AVG CPU UTL} = \frac{\text{Average CPU Busy Time}}{\text{Range Time}} * 100$$ In a PR/SM environment in LPAR mode, the following calculations are used: • For a dedicated partition, the above calculation applies. • For a non-dedicated partition with Wait Completion YES: $$\text{CPU UTL} = \frac{\text{Processor Dispatch Time} - \sum \text{Wait Times}}{\sum \text{Online Times}} * 100$$ • For a non-dedicated partition with Wait Completion NO: $$\text{CPU UTL} = \frac{\text{Processor Dispatch Time}}{\sum \text{Online Times}} * 100$$ Processor dispatch time is the time a processor is dispatched by PR/SM to this partition. Wait time is the time a processor was not used. The average CPU utilization percentage for all processors is also displayed on the SYSINFO report.
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.

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. The fields in the exceptions section of the WFEX report are described in Table 2-79.

Field descriptions

Table 2-79. Fields in the Exceptions Section of the WFEX Report

Field Heading	Meaning
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 a threshold from the **Definition and Criteria** 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.
Reason	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: • A resource contributing most to the delay of a job or job group (main delay) 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%). • A reason meeting a specified exception criterion A specific reason is displayed when a single delay exception is met, even if another resource contributed more to the overall delay. For example - If you specify ENQ% > 10 for jobname xxx, then the Reason field displays **ENQ-majorname** where majorname is the name of the enqueue resource that is causing the delay. If you specify COMM% > 20 for jobname xxx, then the Reason field displays **STOR-COMM.** • An exception The exception statement is displayed when a STOR class exception or another type of single exception that is not delay-oriented is met. For example - If you specify CPUS% > 60, then the Reason field displays **CPUS% > 60.** • SLIP PER TRAP 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. If the **Reason** field displays an exception statement or the SLIP PER TRAP exception, the field is not split by a hyphen. However, the **Reason** column is split into two at the hyphen when a main delay or a specific reason for the exception is displayed. The left part of the column depicts the resource contributing to the exception condition. The resource displayed is either the resource contributing most to the overall delay, or the resource that has been specifically defined (single exception).

Table 2-79. Fields in the Exceptions Section of the WFEX Report (continued)

Field Heading	Meaning
Reason (continued)	The right part of the column depicts the reason for the exception condition. The reason depends on which resource caused it. If the resource under **Reason** is: **PROC** Name of the job using the processor most often when the delayed job wanted to use it appears to the right of the hyphen. **DEV** Volume serial number of the device that the job was most often delayed for appears to the right of the hyphen. **ENQ** Major name of the serially reusable resource causing the greatest percentage of delay appears to the right of the hyphen. **STOR** Either COMM, LOCL, VIO, SWAP, or OUTR appears to the right of the hyphen. If you requested a COMM%, LOCL%, VIO%, SWAP%, or OUTR% single exception, the possible causes are: **COMM** Common storage paging **LOCL** Local storage paging **VIO** Virtual I/O paging **SWAP** Swap-in delay **OUTR** Swapped out and ready. **XMEM** Cross memory address space paging **HIPR** Standard hiperspace paging If STOR is the main delay, or a single STOR% exception is met, only **STOR** appears under **Reason.** The storage delay reason appears as an informational message under **Possible cause or action.** Note that the informational message can be overwritten by a user specified message. For single jobs, 'Main reason **SSSS** causes **xxx %** delay' appears under **Possible cause or action.** For job groups, 'Main reason **SSSS** delays **xxx** users' appears under **Possible cause or action.** **Note:** In both cases, **SSSS** is either COMM, LOCL, VIO, SWAP, XMEM, HIPR, or OUTR. **OPER** 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. **SUBS** Either JES, HSM, or XCF appears to the right of the hyphen. Cursor-sensitive control is split at the hyphen of the **Reason** 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.

Table 2-79. Fields in the Exceptions Section of the WFEX Report (continued)

Field Heading	Meaning
Critical val.	The critical value that caused the exception condition. The following values can appear under **Critical val**: **% delay** Percentage of delay caused by the delay category in the **Reason** field. **users** Average number of delayed users for the group or resource, or maximum number of users. **%** Using percentage in the case of single STOR class exceptions, single common storage exceptions, device class exceptions, CPU% exceptions, or CPUS% exceptions. **frames** Number of frames online in the case of ONLF and ONLXF exceptions, or working set size in expanded storage. **bytes** Number of bytes of storage used in the case of TSQAO exceptions. **/sec** Rate per second in the case of DAR, RATE, PINRT, ESMRT, or ESPRT exceptions. **sec** Amount of time in seconds in the case of time-related exceptions (i.e. AT, DRT, QT, RT, TET, and ESMAG). **replies** Number of outstanding replies. If the exception resource or user is unavailable on the system, the Critical val. field remains blank. For the formula used to calculate the delay of an address space or group of address spaces see "Address space delay (%)" on page 2-9.
Possible Cause or Action	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. 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.

Field descriptions

Table 2-80. 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.

Table 2-80. Fields in the Graphic WFEX Report (continued)

Field Heading	Meaning
Speedometer	How to read a speedometer: 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. Each speedometer is made up of 2 parts, separated by the speedometer needle: 1. A **left** part, which is always solid. The solid left part represents the proportion of a user's time spent doing useful work. 2. A **right** 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. 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. A line in the Exceptions section of the report corresponds to the speedometers with solid right parts and gives details about the exception. The speedometer needle points to the relative speed of the job in the system, from 0 to 100.

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 2-107. 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 2-231.

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 those in the IPS, if customization is set to YES in your Session Options, or if not, to the RMF default values, enter the RESET command.

Table 2-81. Fields in the WFEX Action Panel

Field Heading	Meaning
Action	In the Action column, you can enter the following commands: **S** Access the Definition and Criteria panel **AD** Add a workflow or exception indicator **C** Copy a workflow or exception indicator **D** Delete a workflow or exception indicator **M** Move a line **MM** Move a block **B** Place the moved block or line or copied line on the preceding line **A** 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 (or a domain, performance group, or report performance group number), 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 **Name** on the Workflow/Exceptions (WFEX) report. For an EX-ANY option, the **Name** 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 **Not Displayed** appears under **Row** and **Position,** 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  _____  __   ____  ____
_____  __   ____  ____     _____  __   ____  ____     _____  __   ____  ____
_____  __   ____  ____     _____  __   ____  ____     _____  __   ____  ____
_____  __   ____  ____     _____  __   ____  ____     _____  __   ____  ____
_____  __   ____  ____     _____  __   ____  ____     _____  __   ____  ____
_____  __   ____  ____     _____  __   ____  ____     _____  __   ____  ____
```

Figure 2-108. WFEX Definition and Criteria Panel

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.

Mon III - WFEX report

Table 2-82. Fields in the WFEX Definition and Criteria Panel

Field Heading	Meaning
Class	You specify the class of resources for exception in the Class field. Class can be: **SYSTEM** All jobs in the system **TSO** All TSO/E users **BATCH** All batch jobs **STC** All started tasks **DMN** Jobs grouped by domain **PG** Jobs grouped by performance group **RPG** Jobs grouped by report performance group **JOB** Single job by name **PROC** Processor **DEV** Device **STOR** Storage **ASCH** ASCH address space **OMVS** OMVS address space **SRVCLS** Service class
Qualifier	In the Qualifier field, you specify qualifiers for the following exception classes: **BATCH** Batch jobs - specify the job class (JES initiator class). **BATCH** does not require a qualifier, but if you leave **Qualifier** blank, RMF fills in the default ALL. **JOB** Single job by name - specify a jobname. If indicator EX-UNAVAIL is specified, no qualifier is allowed. **DMN** Jobs grouped by domain - specify a domain number. **DEV** Device - specify a device number. **DEV** does not require a qualifier, but if you leave Qualifier blank, RMF fills in the default ALL. If indicator EX-UNAVAIL is specified, no qualifier is allowed. **PG** Jobs grouped by performance group - specify a performance group number. **RPG** Jobs grouped by report performance group - specify a report performance group number. **SRVCLS** Jobs grouped by service class - specify a valid service class name. **Wildcard Support** You can specify a wildcard character at the end of the input string for the qualifier of the classes JOB and DEV. 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.

Table 2-82. Fields in the WFEX Definition and Criteria Panel (continued)

Field Heading	Meaning
Indicator	In the Indicator field, you specify the workflow or exception indicator type. For Indicator, you can define the following: **WF** 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. **EX-ANY** Indicates that RMF reports one line per address space that meets the condition you specify under Criteria Set. **EX-AVG** Indicates that RMF reports an exception when the average value of the specified group fulfills the conditions you specify under Criteria Set. **EX-GROUP** 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. 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. **EX-UNAVAIL** 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 2-109 on page 2-227 for a description of the panel.
Label	In the Label field, you specify a label for the workflow and exception indicators. 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. 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.
Alert	In the Alert field, you specify the warning signal for the workflow indicator or the exception line. When the threshold values you specified in the **<>, Yel,** and **Red** columns are met, RMF informs you through the alert signal. For Alert, you can specify: **BLINK** The workflow indicator on the tabular report or the exception line on both the tabular and graphic reports blinks **BEEP** Your workstation beeps **BOTH** Your workstation beeps, and the workflow indicator on the tabular report or exception line on both the graphic and tabular reports blinks **NONE** No alerting signal.
Text	In the Text field, you can specify the text for the **Possible Cause or Action** field of the Exceptions section of the Workflow/Exceptions report. **Text** is optional. If you leave it blank, RMF either: • Dynamically fills it in with additional information or with a suggestion of what to do • Leaves it blank when no information is available.

Table 2-82. Fields in the WFEX Definition and Criteria Panel (continued)

Field Heading	Meaning
Criteria Set	Use the **Criteria Set** to specify the exception conditions and color highlighting that RMF should check. If you are defining a workflow indicator, the **Criteria Sets** are optional. If you are defining an exception (EX-ANY, EX-AVG, EX-GROUP), you must enter at least one **Name**, one comparison operand (**<>**), and one threshold value (**Yel** or **Red**) in one criteria set.
Name	You enter the conditions that RMF is to check for under **Name**. 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 **Name** 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 2-227 for more information. You can specify or select up to seven criteria names for each criteria set. 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. To delete a criteria name, on the selection panel, you can either: • Blank out the **S** next to the criteria 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 criteria line.
<>	Specify a comparative operator in the **<>** field. You can specify one of the following: `< > <= >= = ¬= ¬> ¬< GT LT GE LE EQ NE NG NL`
Yel Red	Specify a numeric threshold value in the **Yel** and **Red** 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. 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.

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                    Not allowed on this panel
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
_  ______    _  ______    T  DRV021    T  D13IOG    T  D13MCP    T  D94RM1
Y  PAGE08    T  SPOOL1    T  410PRM    T  410SRO
```

Figure 2-109. 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 (or domains, performance groups), ASCH, and OMVS

  Criteria names AT, QT, and RT are available only for domains and performance groups.
- 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     _ SQAO% 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 %                 _ TSQAO 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 2-110. 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     _ SQAO% 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 %                 _ TSQAO 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 2-111. 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 %                  _ SQAO% SQA overflow %
_ ESMAG Exp storage migration age      _ SSQA% System SQA use %
_ ESMRT Exp storage migration rate     _ TSQAO Total SQA overflow
_ ESPRT Exp storage page movement rate _ UIC   Unreferenced interval count
_ ESQO% ESQA overflow %                _ WFL%  Workflow/Speed %
_ ISTO% Idle storage %
```

Figure 2-112. WFEX Criteria Names for Class: Processor

```
                  RMF WFEX Criteria Names for Class: RPG
Command ===>

Select (S) a maximum of 7 items to use in a criteria set. Press END.


_ ASTO% Active storage %               _ RATE  RPG transaction rate
_ CPUS% CPU utilization %              _ RT    RPG response time
_ 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 %                _ SQAO% SQA overflow %
_ ISTO% Idle storage %                 _ SSQA% System SQA use %
_ LPA%  LPA storage %                  _ TSQAO Total SQA overflow
_ ONLF  Online real storage frames     _ UIC   Unreferenced interval count
_ ONLXF Online expanded storage frames
```

Figure 2-113. WFEX Criteria Names for Class: Report Performance Group

```
                  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 %                _ SQAO% SQA overflow %
_ ISTO% Idle storage %                 _ SSQA% System SQA use %
_ LPA%  LPA storage %                  _ TSQAO Total SQA overflow
_ ONLF  Online real storage frames     _ UIC   Unreferenced interval count
```

Figure 2-114. 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   Avgerage 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           _ SQAO% SQA overflow %
_ DRT   Device response time           _ SSQA% System SQA use %
_ DSC%  Device disconnect time %       _ TSQAO 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 2-115. 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 IPS member name (in compatibility mode) or service policy name (in goal mode) and system ID that was in effect when the data was gathered. If the IPS member, 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 IPS member or 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.

In addition, for compatibility mode reports, RMF sets up the workflow and exception indicators for the WFEX report based on the specifications for workload in the IPS member associated with the data.

You may specify the data set containing the IPS information in the Input Data Set field. If nothing is specified in this field, the default SYS1.PARMLIB is used as the name of the parmlib data set.

Note: You must be authorized for whichever data set is used.

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 IPS member or service policy name and system ID of the data you are looking at.

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.

XCF - Cross-System Coupling Facility Delays report

How to request this report

To request the XCF Delays report, select a **4** from the Primary menu, then a **3** from the Subsystem Report menu (shown in Figure 2-10 on page 2-25) or you can enter the following command:

```
XCF [class,s-class]
```

Purpose of the report

The XCF Delays report lets you investigate situations where executing jobs are delayed when requesting service from XCF.

Contents of the report

```
                        RMF V1R2  XCF Delays                      Line 1 of 1
Command ===>                                                   Scroll ===> HALF

Samples: 100     System: MVS1  Date: 11/29/02  Time: 10.03.20  Range: 100   Sec

             Service    DLY     ------------ Main Delay Path(s) -----------
Jobname   C  Class       %        %  Path    %  Path    %  Path    %  Path

GRS       S  SYSSTC       8       3  0CA0    2  0C80    2  0EA0
```

Figure 2-116. 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 tying up 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 truely 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 2-83. 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: **S** Started task **T** TSO **B** Batch **A** ASCH **O** OMVS
Service Class	The name of the service class that a specified job has been running in. **Note:** For systems running in compatibility mode, the fields DMN and PG are shown instead of Service Class, they display the domain and performance group for the specified job.

DLY %	Delay the waiting job (address space) is experiencing because of delay for XCF during the report interval, expressed as a percentage.
	`# Delay Samples`
	`DLY % = --------------- * 100`
	`# Samples`
	Delay samples
	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.
	Note: This DLY% value is also found in the XCF field on the Job Delay report.

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 2-39 on page 2-73 for an example. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed. See Figure 2-37 on page 2-69 for an example.

Chapter 3. Snapshot reporting with Monitor II

All about Monitor II

This chapter 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
DDMN	Domain Activity Report
DEV	Device Activity Report
HFS	HFS Statistics Report
ILOCK	IRLM Long Lock Detection Report
IOQUEUE	I/O Queuing Activity Report
LLI	Library List Report
PGSP	Page/Swap 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
TRX	Transaction Activity Report

© Copyright IBM Corp. 1990, 2003

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,3`
- Background Session

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

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

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.

Note: Starting with z/OS V1R2 RMF, there is no longer a local 3270 display session. However, you can have access to Monitor II reports without an active TSO/TCAS subsystem by means of the RMF Client/Server Enabling (RMFCS), you find a description in the *z/OS RMF User's Guide*.

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

If you are using the ISPF interface

Each report consists of

- A header line identifying the report.
- A line for commands and scroll amount field.
- A status line for MIG, CPU, UIC, and PR. This line also contains the SMF identifier associated with this system and the current setting of the DELTA mode.
- A variable number of data lines.

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

                      MIG= nnnK CPU=  n UIC=nnnn PR=   n    System= syst  Total
```

Figure 3-1. Header of a Monitor II ISPF Session Report

If you are using the TSO/E interface

Each report consists of

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

Figure 3-2 shows the different report areas for a TSO display session. For a description of each area, see Table 3-1 on page 3-4.

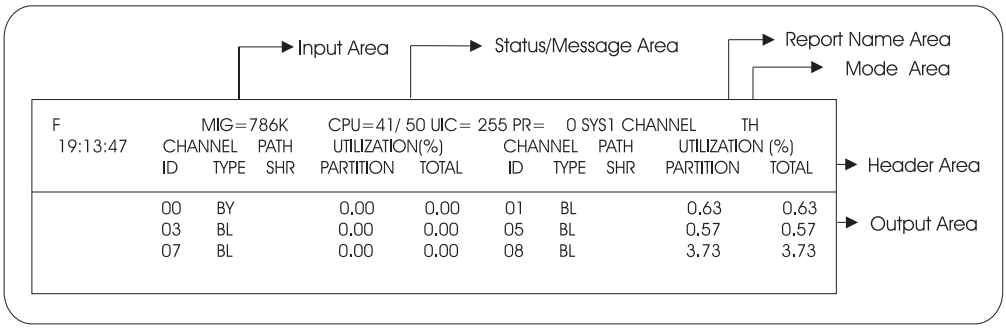

Figure 3-2. Header of a Monitor II TSO Session Report

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 3-1.

Monitor II Reports

Table 3-1. Monitor II Display Session Areas

Area	Function
Report title	The type of measurement data
F	Indicates more pages
Input Area	Issue all commands from here. Separate commands from the MIG=xxx by at least one blank space, or use all 32 spaces, otherwise you get a syntax error.
Status/Message Area	This area contains: **MIG** Migration age (in 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. This value will not be shown in z/Architecture machines (64-bit mode). **CPU** Current average processor utilization. This information depends on the activity of Monitor I. If Monitor I CPU gathering is **active**: • In **LPAR mode,** the header line shows two views separated by a slash (/): – The MVS view of the CPU utilization For the different aspects of CPU utilization refer to "CPU - CPU Activity report" on page 5-41. – The LPAR view of the CPU utilization • In **Basic mode** (non-LPAR mode), the header line shows the MVS view. If Monitor I CPU gathering is **not active**: • In **LPAR mode,** the header line shows: – The SRM view (in compatibility mode, it includes queue length and can have a value of up to 128) of the CPU utilization. – '***' due to missing CPU measurement data for the LPAR view • In **Basic mode** (non-LPAR mode), the header line shows the SRM view. **UIC** The highest unreferenced interval count. **PR** The rate of page-ins per second excluding swap-ins, VIO (virtual input/output), reclaims, and hiperspaces. **System** The SMF identifier associated with this system.
Report Name Area	The report name.
Mode Area	The current setting for the delta mode (either D for delta or T for total) and hardcopy mode (either H for hardcopy, or blank)
Header Area	Consists of two lines of column headings that identify the data fields included in the report.
Output Area	Contains the report data.

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, they can be displayed in:

- **Total** mode

  A total mode report shows the cumulative total since the beginning of the Monitor I interval.

- **Delta** mode

  A delta mode report shows the change in the activity since the previous request for the report.

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(RMFDMTSO) 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 3-3 shows an example of printed output from a display session.

```
                                          RMF MONITOR II HARDCOPY LOG

                                               SESSION NAME TSO

           z/OS V1R4               SYSTEM ID SYS1            DATE 06/14/2002
                                   RPT VERSION V1R2          TIME 17.12.25

++++++++++                                 +++++++ TOTAL MODE ++++++++++++++++++++++++++ ARD      +++++++++++++++++
++ MIG= 729K CPU= 41 UIC= 254 PR=   0 SYSTEM=SYS1 ++++++++++++++++++++++++++++++++++++++++ 17.12.28 +++++++++++++++++
17:12:28 DEV    FF PRV LSQA LSQA X    SRM  TCB    CPU   EXCP SWAP LPA CSA NVI V&H
JOBNAME  CONN  16M  FF  CSF  ESF M CR ABS  TIME   TIME  RATE RATE  RT  RT  RT  RT
*MASTER* 11432   0  12  116   18      0.0 7973.9  17944 ---- ---- --- --- --- ---
PCAUTH   0.000   0   2   30    3 X    0.0   0.04   0.05 ---- ---- --- --- --- ---
RASP     0.000 --- --- ---- ---- X    0.0   0.02  27.48 ---- ---- --- --- --- ---
TRACE    0.000   0   3   49    2 X    0.0   0.02   0.03 ---- ---- --- --- --- ---
XCFAS    0.091   0  22  154   70 X    0.0 203.12 224.59 ---- ---- --- --- --- ---
GRS      29514  17  38   33    5 X    0.0 2135.2  10346 ---- ---- --- --- --- ---
SMXC     0.000   0   2   20   10      0.0  44.81  51.38 ---- ---- --- --- --- ---
SYSBMAS  0.000   0   5   19    8      0.0   0.01   0.02 ---- ---- --- --- --- ---
DUMPSRV  278.1   0   2   37   13      0.0  19.30  61.10 ---- ---- --- --- --- ---
CONSOLE   1148   0   2   31    5 X    0.0 1189.7 1365.2 ---- ---- --- --- --- ---
ANTMAIN  0.085   0   2   23   10 X    0.0   0.02   0.05 ---- ---- --- --- --- ---
ALLOCAS  0.000   0   2   21    3 X    0.0   0.02   0.03 ---- ---- --- --- --- ---
SMF      324.0   0   2   36   10 X    0.0   1.32  43.32 ---- ---- --- --- --- ---
VLF      0.956   1  32   34   10 X    0.0 231.96 235.34 ---- ---- --- --- --- ---
LLA      215.2  10  12   56   10 X    0.0  80.48  96.14 ---- ---- --- --- --- ---
DFRMM    265.3   0   2   61   23 X    0.0  80.95  98.33 ---- ---- --- --- --- ---
FTPSRV01 0.783   1   1   32   10   S  0.0   7.22   7.64 ---- ---- --- --- --- ---
```

Figure 3-3. 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 session 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 3-4 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
PCAUTH   SYSSTC   1 NS    76   33    0   33     0   11    0    0   +0   +0   +0
TRACE    SYSSTC   1 NS    77  176    0  176     0   11    0    0   +0   +0   +0
GRS      SYSTEM   1 NS    FF  561    0  561    33K  11    0    0   +0   +0   +0
CONSOLE  SYSTEM   1 NS    FF  174    0  174     0   11    1    0   +0   +0   +0
ALLOCAS  SYSTEM   1 NS    71  986    0  986     0   11    1    0   +0   +0   +0
LLA      SYSSTC   1 NS    71  170    0  170     0   30    0  150   +0   +0   +0
RMF33    SYSSTC   1 NS    71  159    0  159     0   34    0  150   +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
BOYLEMMM BATCH    2 IN    78   85   48  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  491  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   99  597     0    0    3    0   +0   +0   +0
ZAPPERD  TSO      2 IN    72  110   11  121     0    0    1  150   +0   +0   +0
PTRACYB  TSO      2 IN    78  174   55  229     0    0    3    0   +0   +0   +0
IRLMPROC SYSSTC   1 NS    FC   61    0   61     0   33    0    0   +0   +0   +0
```

Figure 3-4. Format of Printed Reports from a Monitor II Background Session

ARD/ARDJ - Address Space Resource Data report

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)
  ```

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

Purpose of the report

The report gives 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. Such the report gives for example information on processor time, paging, and central storage.

Using the information given in the report

The 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.

Contents of the report

In the ARD report of Figure 3-5, the number of data lines in the report depends on the number of address space identifiers in the system that meet your selection criteria.

In the ARDJ report of Figure 3-7 the number of rows depends on your requests to build a table of information for a particular job.

```
                RMF - ARD Address Space Resource Data           Line 1 of 85
Command ===>                                                  Scroll ===> HALF

                  MIG=49.6K CPU=  9     UIC= 255 PR=  13    System= SYS1 Total

11:08:41 DEV    FF PRV LSQA LSQA X CR SRM TCB   CPU   EXCP SWAP LPA CSA NVI V&H
JOBNAME  CONN  16M  FF  CSF  ESF M    ABS TIME  TIME  RATE RATE  RT  RT  RT  RT

*MASTER* 11539   0 680   95   20      0.0  1069 13187 0.00 0.00 0.0 0.0 0.0 0.0
PCAUTH   0.000   0   2   30    3 X    0.0  0.00  0.01 0.00 0.00 0.0 0.0 0.0 0.0
RASP     0.000 --- --- ---- ---- X    0.0  0.00 63.93 0.00 0.00 0.0 0.0 0.0 0.0
TRACE    0.000   0   3  172    2 X    0.0  0.00  0.01 0.00 0.00 0.0 0.0 0.0 0.0
DUMPSRV  352.7   0   2   63    7      0.0 21.26 31.32 0.00 0.00 0.0 0.0 0.0 0.0
XCFAS    16315   0  2K 1435    0 X    0.0  1237  3053 7.00 0.00 0.0 0.0 0.0 0.0
GRS      1.796   0  51 1134  103 X S  0.0  9261 10583 0.00 0.00 0.0 0.0 0.0 0.0
SMXC     0.000   0   2   41   26      0.0  2529  2862 0.00 0.00 0.0 0.0 0.0 0.0
SYSBMAS  0.000   0 289 1272  672      0.0 826.6 827.4 0.00 0.00 0.0 0.0 0.0 0.0
CONSOLE  8.261   0  23  172    0 X    0.0 340.4 357.3 0.00 0.00 0.0 0.0 0.0 0.0
WLM      1.075   0  35   97   60 X    0.0  1045  1238 0.00 0.00 0.0 0.0 0.0 0.0
ANTMAIN  28267   0  18  158   13 X SC 0.0  2054  2285 0.00 0.00 0.0 0.0 0.0 0.0
ANTAS000 0.742   0   4   68   15 X    0.0  0.18  0.22 0.00 0.00 0.0 0.0 0.0 0.0
OMVS      2438   0 201 2655   91 X    0.0 772.2 856.4 0.00 0.00 0.0 0.0 0.0 0.0
```

Figure 3-5. ARD Report

If the system is running in z/Architecture, the report has the following structure:

```
                RMF - ARD Address Space Resource Data           Line 1 of 85
Command ===>                                                  Scroll ===> HALF

                           CPU=  9     UIC= 255 PR=  13    System= SYS1 Total

11:08:41 DEV    FF   FF PRV LSQA X CR SRM  TCB   CPU  EXCP SWAP LPA CSA NVI V&H
JOBNAME  CONN  16M   2G FF   CSF M    ABS  TIME  TIME RATE RATE  RT  RT  RT  RT

*MASTER* 11539   0    0 680   95      0.0  1069 13187 0.00 0.00 0.0 0.0 0.0 0.0
PCAUTH   0.000   0    0   2   30 X    0.0  0.00  0.01 0.00 0.00 0.0 0.0 0.0 0.0
RASP     0.000 ---  --- --- ---- X    0.0  0.00 63.93 0.00 0.00 0.0 0.0 0.0 0.0
TRACE    0.000   0    0   3  172 X    0.0  0.00  0.01 0.00 0.00 0.0 0.0 0.0 0.0
DUMPSRV  352.7   0    0   2   63      0.0 21.26 31.32 0.00 0.00 0.0 0.0 0.0 0.0
XCFAS    16315   0    0  2K 1435 X    0.0  1237  3053 7.00 0.00 0.0 0.0 0.0 0.0
GRS      1.796   0    0  51 1134 X S  0.0  9261 10583 0.00 0.00 0.0 0.0 0.0 0.0
SMXC     0.000   0    0   2   41      0.0  2529  2862 0.00 0.00 0.0 0.0 0.0 0.0
SYSBMAS  0.000   0    0 289 1272      0.0 826.6 827.4 0.00 0.00 0.0 0.0 0.0 0.0
CONSOLE  8.261   0    0  23  172 X    0.0 340.4 357.3 0.00 0.00 0.0 0.0 0.0 0.0
WLM      1.075   0    0  35   97 X    0.0  1045  1238 0.00 0.00 0.0 0.0 0.0 0.0
ANTMAIN  28267   0    0  18  158 X SC 0.0  2054  2285 0.00 0.00 0.0 0.0 0.0 0.0
ANTAS000 0.742   0    0   4   68 X    0.0  0.18  0.22 0.00 0.00 0.0 0.0 0.0 0.0
OMVS      2438   0    0 201 2655 X    0.0 772.2 856.4 0.00 0.00 0.0 0.0 0.0 0.0
```

Figure 3-6. ARD Report in z/Architecture

```
                RMF - ARDJ Address Space Resource Data           Line 1 of 14
Command ===>                                                   Scroll ===> HALF

                  CPU= 37/ 35 UIC=2540 PR=   0              System= SYS1 Total

BGBO     DEV    FF PRV LSQA LSQA X CR SRM TCB   CPU   EXCP SWAP LPA CSA NVI V&H
  TIME   CONN  16M  FF  CSF  ESF M    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 3-7. ARDJ Report

Field descriptions

Table 3-2. Fields in the ARD and ARDJ Report

Field Heading	Meaning
hh:mm:ss JOBNAME (ARD report)	The time the report was requested, in the form hh:mm:ss, and the one to eight character jobname associated with the address space.
jobname TIME (ARDJ report)	The name of the job you requested. This column contains the time the report was requested, in the form hh:mm:ss.
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 below 2 gigabytes for the job.
PRV 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.
LSQA ESF	The number of private LSQA pages for the job in expanded storage. This value includes private DREF pages.

Table 3-2. Fields in the ARD and ARDJ Report (continued)

Field Heading	Meaning
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.
CR	An indication whether WLM managed the address space as *storage critical* and/or *CPU critical* during the reporting interval. **S** Storage critical **C** CPU critical **SC** 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, in the form nnn.dd, where nnn represents the number of seconds and dd represents hundredths of a second. 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.
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

```
                    RMF Monitor II - Address Space Options
Command ===>

Change or verify parameters. The input entered on this panel applies to
ARD, ASD, and ASRM. To exit press END.

 Class         ===> T       Specify one of the following workloads:
                            A=All, B=Batch/STC, T=TSO, AS=ASCH, O=OMVS
 Inactive      ===> NO      Specify YES to include inactive address spaces.
```

Figure 3-8. 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.

In compatibility mode, there is an additional field **Domain**, which allows you to specify a particular domain.

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

```
              RMF Monitor II - Address Space Options - Single Job
Command ===>

The jobname entered here applies to ARDJ, ASDJ, and ASRMJ.
To exit press END.

 Jobname  ===> ________    Specify a 1 to 8 character jobname.
```

Figure 3-9. 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

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)
  ```

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

Purpose of the report

The 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.

Using the information given in the report

You can use the 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 address space state data 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.

Contents of the report

In the ASD report of Figure 3-10 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 3-11, 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.

```
                  RMF - ASDJ Address Space State Data            Line 1 of 14
Command ===>                                                   Scroll ===> HALF

                  CPU= 37/ 35 UIC=2540 PR=   0              System= SYS1 Total

08:13:57          S C  R  DP   CS   ESF   CS   TAR  X  PIN   ES   TX  SWAP  WSM
JOBNAME  SRVCLASS P L  LS PR    F        TAR   WSS  M   RT   RT   SC    RV   RV

ELF      ADMIN005 1 NS    88  140    90          0       0    0    0    99
PLOT3602 ADMIN001 1 >> DW FF   27    44          0    ---- ----    1    98
DB2CMSTR ADMIN001 1 NS    5B   91   436          0  X    0    0    0   100
YYMU#421 BATCHLNG 1 IN    3B 1675    22          0       0    0    0    99
DB2CDBM1 ADMIN001 1 NS    5B  912  4462          0  X    0    0    0   100
YYBO#426 BATCHLNG 1 IN    3B 1638    79          0       0    0    0    99
SMTP     ADMIN001 1 NS    5B  202   322          0       0    0    1    98
SNALINK  ADMIN001 1 NS    5B   51    67          0       0    0    0   100
BJKE     TSO      4 IN    20 14.9K 5277          0       0    1    0  -145
BTEU     TSO      1 IN    4B  680     3          0       0    0    0    95
PSF3800  ADMIN001 1 NS    5B  304   291          0       0    0    0     0
```

Figure 3-10. ASD Report

```
                  RMF - ASDJ Address Space State Data            Line 1 of 4
Command ===>                                                   Scroll ===> HALF

                  CPU= 37/ 35 UIC=2540 PR=   0              System= SYS1 Total

SMF               S C  R  DP   CS   ESF   CS   TAR  X  PIN   ES  TX  SWAP   WSM
  TIME   SRVCLASS P L  LS PR    F        TAR   WSS  M   RT   RT  SC    RV    RV

12:43:37 SYSTEM*  1 NS    FF  103   250          0  X ---- ----   0     0
12:43:39 SYSTEM*  1 NS    FF  103   250          0  X  0.0  0.0   0     0
12:43:40 SYSTEM*  1 NS    FF  103   250          0  X  0.0  0.0   0     0
12:43:40 SYSTEM*  1 NS    FF  103   250          0  X ---- ----   0     0
```

Figure 3-11. ASDJ Report

Field descriptions

Table 3-3. 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. **Note:** For systems running in compatibility mode, the fields DMN, PG, and PP are shown instead of SRVCLASS, they display the domain, the performance group, and the performance group period for the specified job.
S P	Service class period.

Table 3-3. Fields in the ASD and ASDJ Report (continued)

Field Heading	Meaning
C L	The field showing the current location of the job. The possible contents and their meanings are: **IN** In storage **OT** Swapped out and ready **LO** Logically swapped out **NS** Non-swappable **WM** Waiting for a resource (e.g. DB2 latch or HSM recall): job is swapped in, is eligible for dispatching, and has accumulated no CPU time after some seconds **WL** Wait queue: long wait as a result of either WAIT TYPE=LONG or of STIMER for more than 0.5 seconds **WT** Wait queue: terminal wait **WO** Wait queue: reasons other than WM, WL, or WT **DL** TSO user delayed by SRM to meet response time objective **PR** Privileged **>>** In the process of being swapped out of storage **<<** In the process of being swapped into storage
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: **TI** Terminal input wait **TO** Terminal output wait **LW** Long wait **XS** Auxiliary storage shortage **RS** Central storage shortage **DW** Detected wait **RQ** Requested swap **NQ** Enqueue exchange **EX** Exchange based on recommendation value **US** Unilateral **TS** Transition Swap **AW** APPC wait **IC** Improve central storage **IP** Improve system paging rate **MR** Make room to swap in an out-too-long user **IW** OMVS input wait **OW** OMVS output wait **SR** In-real swap See Table 5-51 on page 5-121 and Table 5-52 on page 5-121 for detailed descriptions of the swap reasons.
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.
ESF	The number of expanded storage frames occupied by virtual storage pages associated with this address space. If no expanded storage is installed, this field is blank.
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.

Table 3-3. Fields in the ASD and ASDJ Report (continued)

Field Heading	Meaning

PIN
RT
	The page-in rate (PIN). RMF calculates the page-in rate in one of two ways:
	• For cross-memory address spaces, the calculation is:
	`PIN = # Page-ins / Transaction Residency Rime`
	• For all other address spaces, the calculation is:
	`PIN = # Page-ins / Seconds of CPU Time`
	The number of page-ins is the sum of the following:
	• Number of pages brought into central storage one at a time
	• Number of pages brought into central storage in blocks
	• Hiperspace read miss count
	• Number of hiperspace pages brought into central storage.
	• Number of shared storage page-ins
	This field always contains the rate since the last report request. An asterisk (*) in this field indicates that RMF is unable to calculate a value.
	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.

ES RT	The rate of page movement from expanded storage due to demand paging (page faults). This includes both single and blocked pages but does not include hiperspace, VIO, or shared storage pages.
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 *z/OS MVS Initialization and Tuning Guide*.
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 3-12 for a description.

ASRM/ASRMJ - Address Space SRM Data report

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)
  ```

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

Purpose of the report

The 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. Thus the report gives for example information on processor service, storage service, and I/O service.

Using the information given in the report

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.

Contents of the report

In the ASRM report shown in Figure 3-12 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             Line 1 of 84
Command ===>                                                    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 3-12. ASRM Report

```
                   RMF - ASRMJ Address Space SRM Data            Line 1 of 14
Command ===>                                                    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 ------
15:44:29 TSOPRD   1   00:00 00:00:00  27   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 ------
15:44:55 TSOPRD   1   00:00 00:00:00 ---   0      0      0      0      0 ------
15:45:05 TSOPRD   1   00:00 00:00:00 ---   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 ------
15:45:32 TSOPRD   1   00:00 00:00:00 ---   0      0      0      0      0 ------
15:45:33 TSOPRD   1   00:00 00:00:00 ---   0      0      0      0      0 ------
```

Figure 3-13. ASRMJ Report

Field descriptions

Table 3-4. 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.

Table 3-4. Fields in the ASRM and ASRMJ Report (continued)

Field Heading	Meaning	
JOBNAME TIME	The job name for the ASRMJ report. This column contains the time the report was requested, in the form hh:mm:ss.	
SRVCLASS	The service class name. **Note:** For systems running in compatibility mode, the fields DMN, PG, and PP are shown instead of SRVCLASS, they display the domain, the performance group, and the performance group period for the specified job.	
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 3-12 for a description.

CHANNEL - Channel Path Activity report

How to request a CHANNEL 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
  ```

Purpose of the 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.

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.

Using the information given in 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.

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          Line 1 of 129
Command ===>                                                  Scroll ===> HALF

                  CPU= 37/ 35 UIC=2540 PR=   0              System= SYS4 Total

15:15:41  Channel   Utilization(%)   Read(B/s) Write(B/s)    MSG  MSG Send Recv
ID No  G  Type  S   Part  Tot  Bus   Part  Tot  Part  Tot   Rate Size Fail Fail

    8    *FC_SM     21.7 38.9 40.2    10M  24M   11M  18M
   12    *FCV_M     13.6 26.8 33.3     9M  16M    7M  12M
    1  1 *FCV_M     12.4 31.9 43.3     8M  22M    7M  19M
    1    *CNC_M     17.2 34.4
80        FC    Y   22.0 43.9 44.0    11M  24M   11M  20M
81        FC_SM Y   21.7 38.9 40.2    10M  24M   11M  18M
82        FCV   Y   14.7 28.5 31.1     8M  17M    6M  11M
83        FCV_M Y   13.6 26.8 33.3     9M  16M    7M  12M
85        CTC_S Y    0.1  3.9
89        CNC_S      0.7  5.8
8C        CNC   Y    1.7 13.5
8D        CNC_S Y    1.8 13.7
DA        CNC_M Y   17.2 34.4
AE        IQD   Y                              705M   4G   923K  764    0    0
AF        IQD   Y                                1G   1G     1M  838    0    0
B0     1  FCV_M Y   12.4 31.9 43.3     8M  22M    7M  19M
```

Figure 3-14. CHANNEL Report

Field descriptions

Table 3-5. Fields in the CHANNEL Report

Field Heading	Meaning
Channel ID	Hexadecimal number of the 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 * preceeding the channel path type, and the number of channels of the group is displayed in column No.
Channel G	Generation. 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, ...). For example, for a FICON channel, a number 1 indicates that the channel has an auto-negotiated throughput of 1Gbit/sec, or a number 2 indicates a throughput of 2Gbit/sec.
Channel 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, this field is blank. RMF continues to measure channel path activity. Check the operator console for messages.
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.

Table 3-5. Fields in the CHANNEL Report (continued)

Field Heading	Meaning
Note: 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).	

Utilization (%) Part	The channel path utilization percentage for an individual logical partition. RMF uses the values provided by CPMF.
	
	In LPAR mode, the calculation is:
	`Part Channel Path Busy Time`
	`Utilization (%) = ------------------------- * 100`
	`Channel Path Elapsed Time`
	
	In BASIC mode or if EMIF is not installed, blanks are shown.
	
	For OSAEGbE, the value reflects the microprocessor utilization.

Utilization (%) Tot	The channel path utilization percentage for the entire system during an interval.
	
	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:
	`# SRM Observations of`
	`Total Channel Path Busy`
	`Utilization (%) = ---------- * 100`
	`# Samples`
	
	For unshared channels in LPAR mode, the value for total utilization is the same as partition utilization.
	
	For all channels in BASIC mode with CPMF available, the calculation is:
	`Part Channel Path Busy Time`
	`Utilization (%) = ------------------------- * 100`
	`Channel Path Elapsed Time`
	
	For OSAEGbE, the value reflects the microprocessor utilization.

Utilization (%) Bus	Percentage of bus cycles, the bus has been found busy for this channel in relation to the theoretical limit. For OSAEGbE, the value reflects the PCI bus utilization.
Read(B/s)	**Part** Data transfer rates from the control unit to the channel for this partition. **Total** Data transfer rates from the control unit to the channel for the entire system.
Write(B/s)	**Part** Data transfer rates from the channel to the control unit for this partition. **Total** Data transfer rates from the channel to the control unit for the entire system.
Note: The following fields are available only for HiperSockets.	
MSG Rate	Rate of messages sent by this partition.
MSG Size	Average size of messages sent by this partition (in bytes).
Send Fail	Rate of messages (sent by this partition) that failed, but not caused by an unavailable buffer in the receiving partition.

Table 3-5. Fields in the CHANNEL Report (continued)

Field Heading	Meaning
Recv Fail	Rate of messages (received by this partition) that failed due to an unavailable buffer. The value could indicate that more receive buffers are required.

DDMN - Domain Activity report

This report is available only if your system is running in compatibility mode.

How to request a DDMN report

- In ISPF, specify **6** on the Resource Report Selection menu.
- In TSO/E, use **PF5** to select the DDMN report.
- Command interface:

  Display session

  ```
  DDMN
  ```

  Background session

  ```
  DDMN
  ```

Purpose of the report

The DDMN report includes the SRM data in the domain descriptor table at the time you requested the report. If you want to include system response data in the report, use the DISPLAY operator command with the DMN option.

Using the information given in the report

The information in this report allows you to verify the effectiveness of the current domain definitions.

Contents of the report

The heading of the report contains the name of the current installation performance specification (IPS) parmlib member followed by the CPU, I/O, central storage, and SRB service coefficients, and the time when the report was requested.

```
                    RMF - DDMN Domain Activity                Line 1 of 13
Command ===>                                                 Scroll ===> HALF

                  CPU= 37/ 35 UIC=2540 PR=   0             System= SYS1 Total

 IPS = IEAIPSAQ   CPU=10.0 IOC= 5.0 SRB=10.0 MSO= 3.0000              14:00:52
 DMN  MIN MAX IMPLT OMPLT CMPL RUA INC NSW OUTU  TWSR    CIDX  RTOB

   0  999 999   999   999    5   5   0   5    0    4K    0.00   0.0
   1    2  50     2     8    0   0   0   0    0  531     0.00   0.0
   2    1  50     3     7    1   2   1   0    0   16K    1.00   0.0
   3   35  50    35    70    1   1   1   0    0    3K    0.00   0.0
   4   10  50    10    17    1   1   0   0    0    3K    0.00   0.0
   5   10  50    10    20    6   7   6   0    0   23K    1.00   0.0
   6    2  50     3     6    3   1   3   0    0    7K   21.24   0.0
   7   23  50    23    41   14  14   0  14    0    6K    0.00   0.0
   8    2  50     2     5    0   0   0   0    0    0     0.00   0.0
   9    1   2     2     6    2   2   0   2    0  867     1.00   0.0
  10    0   0     0     1    0   0   0   0    0    0     0.00   0.0
  11    0   4     0     2    0   0   0   0    0    0     0.00   0.0
  12    2  20     2     4    0   0   0   0    0  110     0.00   0.0
```

Figure 3-15. DDMN Report

Field descriptions

Table 3-6. Fields in the DDMN Report

Field Heading	Meaning
DMN	The domain number.
MIN	The minimum multiprogramming level specified for the domain in the IPS.
MAX	The maximum multiprogramming level specified for the domain in the IPS.
IMPLT	The swap-in level target the SRM has computed for the domain.
OMPLT	The swap-out level target the SRM has computed for the domain. IMPLT will always be less than or equal to the OMPLT. SRM will attempt to maintain the current MPL within the swap-in and swap-out range.
CMPL	The current actual multiprogramming level for the domain.
RUA	The average number of ready users, rounded off to the nearest integer.
INC	The current number of swappable in-storage users.
NSW	The current number of non-swappable in-storage users.
OUTU	The current number of users on the out-queue.
TWSR	The time-weighted average service rate for the domain.

CIDX	The contention index for the domain.
	• If the actual service rate for the domain is less than either the low average service rate for domain (ASRV) or total service rate for domain (DSRV) value, the contention index is greater than 100.
	• If the actual service rate for the domain is in the specified ASRV or DSRV range, the calculation is:
	`High Value - Actual Service`
	`CIDX = ------------------------- * 100`
	`High Value - Low Value`
	
	For example, if the ASRV was specified as (1000 - 3000) the service range is 2000 (high value - low value). If the domain is currently using 1500 service units, the contention index would be ((3000 - 1500) / 2000) * 100 = 75.
	• If the actual service rate for the domain is greater than either the high ASRV or DSRV value, the contention index is less than 1.

RTOB	The average elapsed time (response time) in tenths of a second for first period TSO transactions, if RTO was specified in the IPS for this IPL; otherwise, the field contains zero.

DEV/DEVV - Device Activity report

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))}
  ```

Purpose of the 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.

Using the information given in the report

The 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.

```
    # Samples Taken * Cycle Time
I = ---------------------------
      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.

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.

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 Name	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 **CHGD** 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 5-59 for an explanation of this condition.

Contents of the report

```
                        RMF - DEV Device Activity               Line 1 of 1513
Command ===>                                                    Scroll ===> CSR

                         CPU=  8/  6 UIC=1190 PR=   0         System= SYSD Total

15:32:09 I=14%  DEV             ACTV RESP IOSQ -DELAY- PEND DISC CONN %D %D
STG GRP  VOLSER NUM  PAV  LCU   RATE TIME TIME CMR DB  TIME TIME TIME UT RV

         SYSLIB D540      008B 14.39  1.0  0.0 0.2 0.3  0.5  0.1  0.4  1  1
         110309 0701      0099 13.45  2.1  0.0 0.0 0.0  0.3  0.1  1.7  2  0
         SYSXCA 732A      0078 2.276  2.1  0.0 0.2 0.0  0.3  0.1  1.8  0  0
         MVSLIB 7700      007D 2.161  1.0  0.0 0.1 0.0  0.2  0.0  0.8  0  0
         SYSUSR D440      0089 1.776  0.7  0.0 0.2 0.0  0.2  0.0  0.5  0  0
SMS      SYSSM2 05C2      004A 1.638  1.7  0.0 0.0 0.0  0.2  0.0  1.5  0  0
         SYSSPL 220A  42  0071 1.476  1.7  0.0 0.0 0.1  0.3  0.1  1.3  0  0
         SYSSP2 220B  44  0071 1.300  1.1  0.0 0.0 0.0  0.3  0.1  0.9  0  0
SMS      SYSSME 067B      004C 1.084  3.9  0.0 0.0 0.0  0.8  0.1  3.1  0  0
         SYSXCP D32A      0087 0.646  8.9  0.0 0.2 1.8  2.1  0.0  6.9  0  0
         SYSDPP D43D      0088 0.576  1.0  0.0 0.2 0.0  0.2  0.8  0.0  0  0
         MVSSCF D700      008A 0.476  1.1  0.0 0.2 0.0  0.2  0.1  0.9  0  0
SMS      SYSSMZ D606      0087 0.430  0.9  0.0 0.2 0.0  0.3  0.0  0.6  0  0
SMS      SYSSM7 05C7      004A 0.338  1.4  0.0 0.0 0.0  0.3  0.0  1.1  0  0
```

Figure 3-16. 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. The above sample report is sorted by ACTV RATE.

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.

```
                        RMF - DEVV Device Activity              Line 1 of 11
Command ===>                                                    Scroll ===> CSR

                         CPU=  4/  4 UIC=2190 PR=   0         System= SYSD Total

         I=28%  DEV             ACTV RESP IOSQ -DELAY- PEND DISC CONN %D %D
  TIME   VOLSER NUM  PAV  LCU   RATE TIME TIME CMR DB  TIME TIME TIME UT RV

15:34:08 SYSLIB D540      008B 9.224  0.9  0.0 0.2 0.2  0.5  0.1  0.4  1  0
15:34:52 SYSLIB D540      008B 8.365  0.9  0.0 0.2 0.2  0.5  0.1  0.4  1  0
15:35:44 SYSLIB D540      008B 10.05  0.8  0.0 0.2 0.2  0.4  0.1  0.4  1  0
15:36:14 SYSLIB D540      008B 9.532  0.8  0.0 0.2 0.2  0.4  0.1  0.4  1  0
15:36:46 SYSLIB D540      008B 9.477  0.8  0.0 0.2 0.1  0.4  0.1  0.4  1  0
15:41:34 SYSLIB D540      008B 8.453  0.8  0.0 0.2 0.1  0.4  0.1  0.4  1  0
15:41:44 SYSLIB D540      008B 8.386  0.8  0.0 0.2 0.1  0.4  0.1  0.4  1  0
15:42:30 SYSLIB D540      008B 8.053  0.8  0.0 0.2 0.1  0.4  0.1  0.4  1  0
15:42:31 SYSLIB D540      008B 8.053  0.8  0.0 0.2 0.1  0.4  0.1  0.4  1  0
15:49:05 SYSLIB D540      008B 1.654  1.0  0.0 0.2 0.0  0.3  0.1  0.7  0  0
15:49:14 SYSLIB D540      008B 1.603  1.0  0.0 0.2 0.0  0.3  0.1  0.7  0  0
```

Figure 3-17. DEVV Report

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 3-7. 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 **CHGD** in this column.

I%	The percentage of the Monitor I interval that has elapsed when RMF generates the Monitor II session report.
	`    # Samples Taken * Cycle Time`
	`I = ----------------------------`
	`      Monitor I Interval Length`
	
	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.
	
	See page 3-26 for additional information about the relationship between a Monitor II report period and a Monitor I interval.

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	The number of parallel access volumes (base and alias) which were available at the end of the reporting interval. If the number has been changed during the reporting interval, it is followed by an '*'.
LCU	The hexadecimal number that identifies 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. 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. There are two reasons that this field is blank: • RMF encountered an error while gathering data, check the operator console for messages. • This is a non-dedicated device in a VM guest system environment.

ACTV RATE	The rate at which start subchannel (SSCH) instructions to the device completed successfully.
	`            # Successful SSCH Instructions`
	`ACTV RATE = ------------------------------`
	`                    Interval`

RESP TIME	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.
	`            Total Device Active Time`
	`ACT TIME  = ------------------------`
	`             Measurement Event Count`
	
	`RESP TIME = ACT TIME + IOSQ TIME`
	
	The active time is the sum of connect, disconnect, and pending time as described later.

Table 3-7. Fields in the DEV and DEVV Report (continued)

Field Heading	Meaning
IOSQ TIME	The average number of milliseconds an I/O request must wait on an IOS queue before a SSCH instruction can be issued. IOSQ Count IOSQ TIME = ------------------ # Monitor I Samples --------------------- Device Activity Rate
DELAY CMR	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. Initial Command Response Time DELAY CMR = ---------------------------- Measurement Event Count
DELAY DB	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. A value is reported every 10 seconds. In a PR/SM environment in LPAR mode, this value is updated every 20 seconds. Device Busy Delay Time DELAY DB = ---------------------- Measurement Event Count If the data is not valid, a dash (-) will be displayed.
PEND TIME	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. Device Pending Time PEND TIME = ---------------------- Measurement Event Count
DISC TIME	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. Device Disconnect Time DISC TIME = ---------------------- Measurement Event Count 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.
CONN TIME	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. Device Connect Time CONN TIME = ---------------------- Measurement Event Count

Table 3-7. Fields in the DEV and DEVV Report (continued)

Field Heading	Meaning

% D UT	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.
	
	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.
	`CONN + DISC RSV`
	`% D UT = (----------- + ---) * 100`
	`INT * PAV N`
	CONN Device connect time
	DISC Device disconnect time
	RSV Number of Monitor I samples when the device was reserved but not involved in an I/O operation
	INT Monitor I interval time (seconds)
	PAV The number of parallel access volumes (or 1 for a non-PAV device)
	N Total number of Monitor I samples
	
	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%.

%D
RV
	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.
	
	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.
	`# Device-reserved Samples`
	`% D RV = ------------------------`
	`# 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 3-18. 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:zzzzzz), 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.

```
            RMF Monitor II - Device Activity Options - Single Device

Command ===>

Specify one of the options below. To exit press END.

 Volume         ===> ______     Specify a volume serial number.

 Device Number  ===> ___        Specify a hexadecimal device number.
```

Figure 3-19. 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,cccccc,... 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-character, hexadecimal device number.

HFS - Hierarchical File System Statistics report

How to request an HFS report

- In ISPF, specify **5** on the I/O Report Selection menu.
- Command interface:

  Display session

  ```
  HFS [hfsname]
  ```

Purpose of the report

The 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.

Contents of the report

```
                    RMF - HFS File System Statistics            Line 1 of 12
Command ===>                                                   Scroll ===> PAGE

                  CPU= 37/ 35 UIC=2540 PR=   0              System= SCLM Delta

File System Name: OMVS.SYS1.ROOT
Mount Date: 03/15/2002   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 3-20. HFS Report

Field descriptions

Table 3-8. 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).

Table 3-8. Fields in the HFS Report (continued)

Field Heading	Meaning
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 **Count** and **Rate** (counts per second).	
Cache	Number of times the first page of a data file was requested and found in virtual storage (cache).
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 **Count** and **Rate** (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 **Count** and **Rate** (counts per second).	
Cache	Number of index page read or write hits.

Table 3-8. Fields in the HFS Report (continued)

Field Heading	Meaning
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 3-21. HFS Report Options Panel

Table 3-9. 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).
Display	**Yes/No** specifies whether the names of the currently active file systems should be listed for selection. **No** is the default. **Note:** 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

How to request an ILOCK report

- In ISPF, specify **9** on the Resource Report Selection menu.
- Command interface:

  Display session

  ```
  ILOCK [ALL]
  ```

Purpose of the 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.

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

```
                  RMF - ILOCK IRLM Long Lock Detection          Line 1 of 15
Command ===>                                                   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 09/05/2002 13:02:10 Deadlock Cycle 00002EC7
--------------------------------------------------------------------------------
TOP      BMP     09C943CFA7800101D700000000000000 DFSSAMB1  00:06:04
BLOCKER  ACO3    ACO3    0000000300000000   0006  IRLMTOPZ
--------------------------------------------------------------------------------
TOP      BMP     09C3614505800101D700000000000000 DFSSAMB1  00:06:09
BLOCKER  ACO1    ACO1    0000000600000000   0006  IRLMTOPA
--------------------------------------------------------------------------------
WAITER   BMP     09C3614505800101D700000000000000 DFSSAMB2
         ACO2    ACO2    0000000800000000   0007  IRLMWTA1  00:05:52  DI21PART
--------------------------------------------------------------------------------
WAITER   BMP     09C943CFA7800101D700000000000000 DFSSAMB7
         ACO2    ACO2    0000000900000000   0008  IRLMWTZ2  00:05:42  DI21PART
--------------------------------------------------------------------------------
```

Figure 3-22. ILOCK Report

Field descriptions

Table 3-10. 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 deadlock cycle number (in hexadecimal format) 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: **BLOCKER** Transaction holds a lock at the time the exit is driven. **TOP BLOCKER** 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. **WAITER** Transaction is waiting for a lock. **BLOCKER/WAITER** The transaction was found as a blocker and waiter. **Note:** 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), BATCH, and CICS.
IMS_ID	Name given to the IMS region at the time it is brought up.
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.

Table 3-10. Fields in the ILOCK Report (continued)

Field Heading	Meaning
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 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

How to request an IOQUEUE report

- In ISPF, specify **2** on the I/O Report Selection menu.
- Command interface:

  Display session

  ```
  IOQUEUE [type]
  ```

  Background session

  ```
  IOQUEUE [(type)]
  ```

Purpose of the 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.

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).

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 LNGTH
- 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 LNGTH
- %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 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.

Note: Starting with z/OS V1R2 RMF, the percentage about all channel paths being busy has been removed due to data inconsistencies for different channel types.

Contents of the report

```
                    RMF - IOQUEUE I/O Queuing Activity            Line 1 of 6
Command ===>                                                   Scroll ===> HALF

                         CPU= 37/ 35 UIC=2540 PR=   0        System= SYS1 Total

16:56:45  I= 78%   DCM Group       Cont  Del Q  AVG  CHPID  %DP  %CU   AVG  AVG
Path DCM CTL Units MN MX DEF LCU   Rate  Lngth  CSS  Taken  Busy Busy  CUB  CMR

11       400A 400B           0031                     2.86   0.0  1.1  0.0  0.0
         400C 400D
2E    Y  400A              0031                     5.88  10.7  0.2  0.0  0.0
20       400B              0031                     2.64  18.2  2.5  0.5  0.2
                   1  2  2  0031                     5.88  10.7  0.2  0.0  0.0
                              0031   1.1   4.78  1.7  11.39  10.2  1.0  0.0  0.0
11       440A              0032                     0.03   0.0 33.3  0.0  0.0
92    Y  440A              0032                     2.59  37.7  4.5  0.6  0.1
20       440B              0032                     0.00   0.0  0.0  0.0  0.0
                   0  1  2  0032                     2.59  37.7  4.5  0.6  0.1
                              0032   2.4   3.61  0.0   2.63  37.3  4.8  0.3  0.1
12       500A              0033                     0.29   5.5  0.0  0.0  0.0
93       500B              0033                     0.22  18.7  0.0  0.0  0.0
                              0033   0.0   0.00  0.0   0.51  11.7  0.0  0.0  0.0
12       700A              0034                     0.68  29.8  0.0  0.0  0.0
93       700B              0034                     0.46  43.7  0.0  0.0  0.0
                              0034   0.1   0.00  0.0   1.14  36.1  0.0  0.0  0.0
11       580A              0035                     0.03   0.0 33.3  0.0  0.0
20       580B              0035                     0.07   0.0 33.3  0.0  0.0
                   0  0  2  0035                     0.00   0.0  0.0  0.0  0.0
                              0035   0.0   0.00  0.0   0.10   0.0 33.3  0.0  0.0
11       5C0A              0036                     0.05   0.0  0.0  0.0  0.0
2E    Y  5C0A              0036                     1.39   2.3  0.0  0.0  0.0
```

Figure 3-23. IOQUEUE Report

Field descriptions

Table 3-11. Fields in the IOQUEUE Report

Field Heading	Meaning
Path	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.
DCM	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.
CTL Units	The hexadecimal identifier of each physical control unit associated with an online channel path in the LCU group.
DCM Group	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. 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.

Table 3-11. Fields in the IOQUEUE Report (continued)

Field Heading	Meaning
LCU	The hexadecimal number that identifies the logical control unit (LCU). 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.

Cont
Rate
	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.
	
	`# Enqueued Requests`
	`Cont Rate = ----------------------`
	`Monitor I Interval Time`

Del
Q Lngth
	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.
	
	`Accumulated Queue Length -`
	`# Enqueued Requests`
	`Del Q Lngth = ----------------------------`
	`# Enqueued Requests`

AVG
CSS
	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.
	
	`Channel Subsystem Time`
	`AVG CSS = ------------------------`
	`# I/O Operations Accepted`

CHPID
Taken
	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.
	
	`# I/O Operations on that Path`
	`CHPID Taken = ----------------------------`
	`Monitor I Interval Time`

% DP
Busy
	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.
	
	`DPB`
	`% DP Busy = --------------- * 100`
	`DPB + CUB + SUC`
	
	DPB Number of deferred I/O requests due to director port busy
	CUB Number of deferred I/O requests due to control unit busy
	SUC Number of successful I/O requests on that path

% CU
Busy
	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.
	
	`CUB`
	`% CU Busy = --------------- * 100`
	`DPB + CUB + SUC`
	
	DPB Number of deferred I/O requests due to director port busy
	CUB Number of deferred I/O requests due to control unit busy
	SUC Number of successful I/O requests on that path

Table 3-11. Fields in the IOQUEUE Report (continued)

Field Heading	Meaning

AVG CUB	The average number of milliseconds of delay that an I/O request encountered for the channel path because the control unit was busy.
	`Control Unit Busy Time`
	`AVG CUB = ----------------------------------`
	`# I/O Operations Accepted on that Path`

AVG CMR	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).
	`Initial Command Response Time`
	`AVG CMR = ----------------------------------`
	`# 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 3-24. 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 hexadecimal and have three characters. 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

How to request an LLI report

- In ISPF, specify **L** on the Primary menu, this leads to the Library List menu. Here, you can select:

  1 LNKLSTxx - Link Library List

  2 LPALSTxx - LPA Library List

  3 IEAAPFxx - Authorized Program Library List

- Command interface:

  Display session

  ```
  LLI      Link Library List
  LLI LPA  LPA Library List
  LLI APF  Authorized Program List
  ```

  In addition, an optional parameter `A` can be specified to create a report with more details. Table 3-12 on page 3-47 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.

Purpose of the 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.

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 3-25. LLI Report - Link Library List

The link library list is determined by Parmlib members LNKLSTxx 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 3-26. 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 3-27. LLI Report - APF Library List

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 3-12. 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	**Y** The link library is APF-authorized. **N** 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: **STATIC** IEALPAxx is used, or STATIC has been selected in PROGxx. **DYNAMIC** DYNAMIC has been selected in PROGxx.
SMS	**Y** The library is SMS-managed. **N** The library is not SMS-managed.
RACF	**Y** The library is defined to RACF. **N** The library is not defined to RACF. **Sec=?** 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.

PGSP - Page/Swap Data Set Activity report

How to request a Page Data Set Activity 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)
  ```

How to request a Swap Data Set Activity report

In z/OS, swap data sets are not supported. Therefore, the option SWAP will create an empty report. You may use this option only if you have systems in your sysplex with releases prior to OS/390 2.10. Please, refer to *OS/390 2.7 RMF Report Analysis* for a description of swap data set activity.

Purpose of the report

The PGSP report provides information on page data set activity. The reporting interval is the period between any two consecutive Monitor II requests.

Using the information given in the report

The PGSP report can help you to determine whether the optimum size has been allocated for each page data set.

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/Swap Data Set Activity            Line 1 of 8
 Command ===>                                                  Scroll ===> HALF

                   CPU= 37/ 35 UIC=2540 PR=   0              System= SYS1 Total

S VOLUME DEV  DEV      %SLOTS   PAGE   I/O REQ AVG PAGES        14:39:33
T SERIAL NUM  TYPE     IN USE TRAN TIME  RATE   PER I/O V  DATA SET NAME

P RES21B 06C5 33903    64.40   0.000     0.000    0.000    SYS1.PLPA
C RES21B 06C5 33903     8.24   0.000     0.000    0.000    SYS1.COMMONP
L PGE120 0120 33903    99.99   0.000     0.500    1.000 N  SYS1.PAGE11
L PGE130 0130 33903    100.0   0.000     0.000    0.000 N  SYS1.PAGE13
L PAG8F3 08F3 33903    14.97   0.000     0.000    0.000 Y  SYS1.PAG8F3
L PAG2B7 02B7 33903     8.23   0.000     1.500    5.000 Y  SYS1.PAG2B7
L PAG714 1714 33903     8.67   0.000     2.000    7.500 Y  SYS1.PAG714
L PAG6C4 26C4 33903    17.75   0.000     2.000   12.750 Y  SYS1.PAG6C4
```

Figure 3-28. PGSP Page Report

Field descriptions

Table 3-13. Fields in the PGSP Report

Field Heading	Meaning	
S T	The one-letter identifier of the type of paging space. The identifiers are: **P** PLPA **C** Common **L** Local	
VOLUME SERIAL	The volume serial number of the volume on which the data set resides.	
DEV NUM	The device number.	
DEV TYPE	The device type.	
% 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. `% SLOTS IN USE = (# Slots in the Data Set - # Available Slots) / # Slots in the Data Set * 100`	
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. `PAGE TRAN TIME = ((USE*INT)/N)/XFER` **USE** Number of samples when the data set is in use **XFER** Total number of pages transferred **N** Total number of samples **INT** Monitor I interval time (seconds)	This field always contains a value that reflects the activity since the last report request.
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.	

Table 3-13. 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. 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 *preceding* 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

How to request an SDS report

- In ISPF, specify **8** on the Resource Report Selection menu.
- Command interface:

  Display session

  ```
  SDS
  ```

Purpose of the 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.

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 04/29/2002 16:49:24, End 05/06/2002 12:28:21, Duration   6d 19h


SMF Buffer Statistics, Start 04/29/2002 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 3-29. SDS Report

Field descriptions

Table 3-14. 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: **ERBDSQRY** RMF Query Available Sysplex SMF Data Service **ERBDSREC** RMF Request Sysplex SMF Record Data Service **ERB2XDGS** RMF Monitor II Sysplex Data Gathering Service **ERB3XDRS** RMF Monitor III Sysplex Data Retrieval Service
Requests Arrived	Number of calls to the data service.
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

How to request an SENQ 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])}
```

Purpose of the 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.

Using the information given in the 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*.

Different types of SENQ reports

You can request the SENQ report as a:

- 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 3-30 on page 3-54 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 3-32 on page 3-55 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 3-31 on page 3-55 shows a sample resource report for a specific system.

Contents of the report

```
                RMF - SENQ System Enqueue Contention              Line 1 of 4
 Command ===>                                                  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 3-30. SENQ Summary Report

```
                RMF - SENQ System Enqueue Contention              Line 1 of 14
 COMMAND ===>                                                  Scroll ===> HALF

                   CPU= 37/ 35 UIC=2540 PR=   0              System= SYS1 Total

11:44:26                     MAJOR NAME
JOBNAME  SYSTEM  ASID  REQ   MINOR NAME

                             ADMPRNTQ
ADMPRINT M303    185   EO    BACKGROUND DRIVER EXECUTING.                (SYSS)
                             ARCGPA
HSM      M303     13   SO    ARCMCDS                                     (SYSS)
                             BLXDASDS
D10MCW1  M303    116   SO    D83INFO.BLGPNLS                             (SYSS)
                             CLR12.UN
D32VJF1  M303     88   SO    IUSER                                       (SYSS)
                             GIMSMP
D31BAR1  M303    217   EO    C87JPLX.SMPCSI.CSI                          (SYSS)
                             SPFDSN
D31BAR1  M303     35   EO    D31BAR1.IAREND.P06.PLS                      (SYSS)
D75DVB1  M303     41   EO    D75DVB1.IATOSDR.W1A.ASM                     (SYSS)
D96JRF1  M303    163   EO    D96JRF1.IEFDB413.PS1.PLS                    (SYSS)
```

Figure 3-31. SENQ GRS Report

```
                RMF - SENQ System Enqueue Contention              Line 1 of 4
 COMMAND ===>                                                  Scroll ===> HALF

                   CPU= 37/ 35 UIC=2540 PR=   0              System= SYS1 Total

11:48:44                    MAJOR NAME
JOBNAME  SYSTEM  ASID REQ   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 3-32. SENQ Detail Report

Field descriptions

Table 3-15. 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 3-15. 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: **E** The request was for exclusive use of the resource **S** The request was for shared use of the resource The second character indicates the status of the request: **O** The requestor owns the resource **W** The requestor is waiting for the resource	
MAJORNAME MINORNAME	The name and scope of the resource. The major name, which corresponds to the *qname* 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 *rname* 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 *name* 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
                               or NO for exclusively owned resources only.
   System ID    ===> ________  Specify the system holding the resources.

 Enqueue report by major-/minorname:

   Majorname    ===> ________
   Minorname    ===> ____________________________
```

Figure 3-33. 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

How to request an SENQR 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) }
```

Purpose of the 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.

Figure 3-34 shows a sample SENQR report requested for all devices.

Using the information given in 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.

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    SO  TSO020 0AF1 CNV SYSIGGV2 CATALOG.VTSO020
BMORRISP  70 AQTS    EO  RES84Z 0D4E  ON SYSVTOC  RES84Z
JES2      31 AQTS    EO  JES2PK 0702 OFF SYSZJES2 JES2PKSYS1.BKUPCKPT
CATALOG   34 AQTS    SO  CAT212 04B1 OFF SYSZVVDS CAT212
```

Figure 3-34. SENQR Report

Field descriptions

Table 3-16. 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.

Table 3-16. Fields in the SENQR Report (continued)

Field Heading	Meaning
ASID	The address space identifier of the job that issued the RESERVE macro instruction for the device identified under DEV.
REQ	The two-character field that describes the request. The first character indicates the type of the request: **E** The request was for exclusive use of the device **S** The request was for shared use of the device The second character indicates the status of the request: **O** The requestor owns the device **W** The requestor is waiting for the device
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	The indicator of the reserve status of the device. **ON** The device is reserved by the processor on which RMF is running. **OFF** The device is being serialized via RESERVE macro instructions, but is currently not reserved. **CNV** The device has been converted to a GRS ENQ. If reserves are issued on systems other than the one on which you request the report, the field is blank.
MAJOR MINOR	The name used to control access to the device by means of RESERVE macro instructions. The major name, which corresponds to the *qname* field in the RESERVE macro instruction, is one to eight characters in length. It is aligned under the MAJOR heading. The minor name, which corresponds to the *rname* 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 *name* 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.

Report options

```
                RMF Monitor II - System Enqueue Reserve Options

Command ===>

Specify a volume serial number or ALLSERV. To exit press END.

 Volume  ===> ALLSERV   Reports reserve requests for a particular volume.
                        Specify ALLSERV for all requests.
```

Figure 3-35. SENQR Report Options Panel

Enter the reserved volume name. For a report of all reserved volumes enter ALLVSER.

SPAG - Paging Activity report

How to request an SPAG 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
  ```

Purpose of the report

The Paging Activity report presents overview information on system paging activity.

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 3-36 shows how you can repeat requests for the report to build a table of system paging activity.

Using the information given in the report

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).

Contents of the report

```
                    RMF - SPAG Paging Activity                 Line 1 of 9
Command ===>                                                  Scroll ===> HALF

                 CPU= 37/ 35 UIC=2540 PR=   0              System= SYS1 Total

        LPA CSA SWP  PGS-SWPD PRIV_IN PRV V&H TAR        HI   ES MIG  ESF  MIG
  TIME   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 3-36. SPAG Report

Field descriptions

Table 3-17. Fields in the SPAG Report

Field Heading	Meaning
TIME	The time the report was requested in the form of hh.mm.ss.
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 on all available frame queues.
HI UIC	The highest unreferenced interval count (UIC).
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.
MIG AGE	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. If there is no unit specified for MIG AGE, the value is in migration hours. Other values are indicated as follows: **M** Migration minutes **S** 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.
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

How to request an SRCS 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
  ```

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 3-37 is an example of repeated requests.

Using the information given in the report

This report detects abnormal situations when they occur so you can request a more detailed report to further analyze the situation.

Contents of the report

```
               RMF - SRCS Central Storage / Processor / SRM          Line 1 of 8
 Command ===>                                                        Scroll ===> HALF

                    CPU= 37/ 35 UIC=2540 PR=   0              System= SYS1 Total

                HI SQA  LPA  LPA CSA  L+C  PRI LSQA LSQA CPU  IN OUT OUT OUT
  TIME   AFC  UIC   F    F   FF   F   FF   FF  CSF  ESF UTL   Q LOG  RQ  WQ

11:16:42 298  255 2.0K 1.4K  61 1.8K 245  728 8.7K 1.9K 105  57 181   1 340
11:16:44 366  255 2.0K 1.4K  61 1.8K 245  788 8.7K 1.9K 102  61 177   1 340
11:16:46 325  255 2.0K 1.4K  61 1.8K 245  760 8.8K 1.9K  98  54 184   1 340
11:16:47 297  255 2.0K 1.4K  61 1.8K 245  734 8.7K 1.9K 104  54 184   1 340
11:16:49 466  255 2.0K 1.4K  61 1.8K 245  754 8.9K 1.9K  98  53 187   2 337
11:16:51 283  255 2.0K 1.4K  61 1.8K 245  789 8.8K 1.9K  97  50 190   1 338
11:16:52 243  255 2.0K 1.4K  61 1.8K 245  764 8.8K 1.9K  99  50 190   1 338
11:16:54 695  255 2.0K 1.4K  61 1.8K 245  771 8.8K 1.9K  87  49 190   2 338
```

Figure 3-37. SRCS Report

Field descriptions

Table 3-18. Fields in the SRCS Report

Field Heading	Meaning	
TIME	The time the report was requested in the form hh.mm.ss.	
AFC	The average length of the available frame queue.	
HI UIC	The highest unreferenced interval count (UIC).	
SQA F	The total number of SQA frames, including frames in central storage and expanded storage.	If invalid data occurs, dashes appear in this field.
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.	
CPU UTL	The average processor utilization percentage for all processors currently online. Depending on Monitor I CPU data gathering, this is either the MVS or the SRM view of the CPU utilization and is identical to the first value in the CPU= field in the header. For details, please refer to Table 3-1 on page 3-4.	
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.	

TRX - Transaction Activity report

This report is only available when your system is running in compatibility mode.

How to request a TRX report

- In ISPF, specify **7** on the Resource Report Selection menu.

  By default, you get all performance groups, other reports can be selected via the Report Options panel.
- In TSO/E, use **PF12** to select the TRX report.
- Command interface:

  Display session

```
TRX {ALLPGN                      }
    {sname,nnnn[,nnnn,nnnn:nnnn]}
```

  Background session

```
TRX {(ALLPGN)                      }
    {(sname,nnnn[,nnnn,nnnn:nnnn])}
```

Purpose of the report

The report is a “snapshot” of transaction activity at the time the report was requested. It shows, by performance group period, the transaction rate (the number of ended transactions per second) and the average elapsed time of each ended transaction.

Using the information given in the report

The TRX report enables you to evaluate the effectiveness of the performance objectives in the IPS and the workload classification in the installation control specification (ICS).

You can compare the transaction rate and the average elapsed time required for each transaction to the objectives for the performance group specified in the IPS.

When the installation control specification associates a performance group with a subsystem, a transaction class, a user identifier, or a transaction name, the report contains this data. Using this data you can relate the contents of the report to the work classification established in the installation control specification.

Special considerations of report output

You can request transaction data for one or more individual performance groups or for all performance groups associated with a subsystem name.

Only performance groups with active work are included in the report; any performance group that had no activity does not appear. Any performance group specified in the report request that is not defined in the IPS or the installation control specification appears in the report marked as *****PGN NOT DEFINED*****.

RMF requires two “snapshots” in order to calculate the transaction rate and the average time for an ended transaction. Thus, TRANS RATE and AVG TRANS

TIME contain dashes the first time the report is requested and the first time a performance group or performance group period appears in a series of report requests.

Contents of the report

The name of the installation control specification is in the format ICS=**name**, where **name** is the eight-character name of the IEAICSxx parmlib member or NONE if no IEAICSxx member is currently in effect.

The name of the IPS is in the format IPS=**name**, where **name** is the eight-character name of the IPS or NONE if the coded defaults for the IPS are currently in effect.

```
                RMF - TRX Transaction Activity                    Line 1 of 5
Command ===>                                                   Scroll ===> HALF

                  CPU= 37/ 35 UIC=2540 PR=   0               System= SYS1 Total

16:56:45  ICS=IEAICS00  IPS=IEAIPST0     ACCT    PERF    TRANS  AVG TRANS TIME
SUBSYS  TRXCLASS   USERID     TRXNAME    INFO   GRP PER   RATE   HHH.MM.SS.TTT

TSO                                        NO C   2   1   1.611  000.00.00.024
TSO                                        NO         2   0.388  000.00.00.095
TSO                                        NO         3   0.055  000.00.00.587
TSO                                        NO         4   0.222  000.00.29.636
JES3    CODEM                              NO C  13   1   4.111  000.00.00.924
```

Figure 3-38. TRX Report

Field descriptions

Table 3-19. Fields in the TRX Report

Field Heading	Meaning
SUBSYS	The name of the subsystem associated with the performance group in the installation control specification. When a subsystem is specified by name in the report request, all performance groups associated with that subsystem are grouped in the report. The field is blank when there is no installation control specification, when it does not define a subsystem name for the performance group, or when the name is not unique (that is, when more than one subsystem name is associated with the performance group).
TRXCLASS	The name of the transaction class defined in the installation control specification for the performance group. If a sub-string was used to define the name, the name reported includes the sub-string position number. The field is blank when there is no installation control specification, when it does not define a transaction class for the performance group, or when the name is not unique (that is, when more than one transaction class name is associated with the performance group).
USERID	The user identifier defined in the installation control specification for the performance group. If a sub-string was used to define the user identifier, the user identifier reported includes the sub-string position number. The field is blank when there is no installation control specification, when it does not define a user identifier for the performance group, or when the user identifier is not unique (that is, when more than one user identifier is associated with the performance group).

Table 3-19. Fields in the TRX Report (continued)

Field Heading	Meaning
TRXNAME	The transaction name defined in the installation control specification for the performance group. If a sub-string was used to define the same, the name reported includes the sub-string position number. The field is blank when there is no installation control specification, when it does not define a transaction name for the performance group, or when the transaction name is not unique (that is, when more than one transaction name is associated with the performance group).
ACCT INFO	The account number specifications. YES indicates that an account number was specified in the installation control specification for the performance group. Blanks indicate that no account number has been assigned.
PERF GRP	The performance group. The number is preceded by a C when the performance group is a control performance group and by an R when the performance group is a report performance group.
PERF PER	The performance group period. When more than one period is active for a performance group, this field, TRANS RATE, and AVG TRANS TIME are repeated for each active performance group period.
TRANS RATE	The transaction rate (number of ended transactions per second) for the performance group period. Because RMF requires two “snapshots” to compute the value for the field, TRANS RATE contains dashes the first time the report is requested or the first time a performance group period appears in a series of reports. TRANS RATE is repeated for each performance group period included for a performance group.
AVG TRANS TIME HHH.MM.SS.TTT	The average elapsed time of all ended transactions for the performance group period, where HHH is the hours, MM is the minutes, SS is the seconds, and TTT is thousandths of a second. Because RMF requires two “snapshots” to compute the value for this field, AVG TRANS TIME contains dashes the first time the report is requested or the first time a performance group appears in a series of reports. AVG TRANS TIME is repeated for each performance group period includes for a performance group.

Report options

```
                RMF Monitor II - Transaction Activity Options

Command ===>

Specify one or more subsystems and/or individual performance groups
or ranges of performance groups. To include all performance groups,
specify ALLPGN. To exit press END.

 Performance groups  ===> ALLPGN_____________________________   Ex: 1,5:10
```

Figure 3-39. TRX Report Options Panel

Type the name of the performance groups you want a report for. If you leave the panel empty, ALLPGN is used.

Chapter 4. Real-time reporting with Monitor I

Monitor I Summary

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
CPU	CPU Activity Report
CRYPTO	Crypto Hardware Activity Report
DEVICE	Device Activity Report
ENQ	Enqueue Activity Report
FCD	FICON Director Activity Report
IOQ	I/O Queuing Activity Report
PAGESP	Page/Swap Data Set Activity Report
PAGING	Paging Activity Report
TRACE	Trace Activity Report
VSTOR	Virtual Storage Activity Report
WKLD	Workload Activity Report

See the following chapter for a description of these reports.

© Copyright IBM Corp. 1990, 2003

Chapter 5. Long-term overview reporting with the Postprocessor

All about Postprocessor reports

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 chapter describes the following 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.
- **Plot report**

  Presents a graphic summary of a number of key system performance indicators.
- **Summary report**

  Presents an overview of system activity.

© Copyright IBM Corp. 1990, 2003

Interval and duration reports

Interval 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:

Single-System report

```
REPORTS(option)
```

or as sysplex reports with the report option:

Sysplex report

```
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 4-1.

Samples of the interval reports printed during a Monitor II session appear in Chapter 3, "Snapshot reporting with Monitor II".

Duration reports

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.

This chapter describes all interval reports in the following order:

Table 5-1. 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(FCD)	FICON director activity	Monitor I	74.7
REPORTS(HFS)	HFS statistics	Monitor III	74.6
REPORTS(HTTP)	HTTP server activity	HTTP server	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
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
REPORTS(WKLD)	Workload activity	Monitor I	72.1
SYSRPTS(WLMGL)	Workload activity	Monitor I	72.3
REPORTS(XCF)	XCF activity	Monitor III	74.2
Note: The ENQ report and the TRACE report are available only as interval reports.			

For each of the reports in the list above, this chapter provides a general introduction, shows a sample report, and describes the meaning of each field in the report.

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 3-1.

Table 5-2. 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

Table 5-2. Monitor II Interval Reports (continued)

Report Option	Report Name	SMF Record
REPORTS(DDMN)	Domain activity	79.10
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
REPORTS(TRX)	Transaction activity	79.8
Note: The DDMN report and the TRX report are available only in compatibility mode.		

Single-system report header

```
                              R E P O R T   T I T L E
                                                                        PAGE nnnn
z/OS V1R4            SYSTEM ID   cccc          DATE mm/dd/yyyy         INTERVAL mm.ss.ttt
                     RPT VERSION V1R2 RMF      TIME hh.mm.ss           CYCLE s.ttt SECONDS
```

Figure 5-1. 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 V1R4	The version of the operating system.
SYSTEM ID cccc	The SMF identifier associated with this system.
RPT VERSION V1R2 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

```
                              R E P O R T   T I T L E
                                                                              PAGE nnnn
z/OS V1R4             SYSPLEX sysplex          DATE mm/dd/yyyy          INTERVAL mmm.ss.ttt
                      RPT VERSION V1R2 RMF     TIME hh.mm.ss            CYCLE ss.ttt SECONDS
```

Figure 5-2. 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 identifier.

Note: The release information will be shown as '*****' if different releases of z/OS or OS/390 are running concurrently in the sysplex.

Duration report header

```
                           R E P O R T   T I T L E
                                                                              PAGE nnnn
z/OS V1R4             SYSTEM ID   cccc         START mm/dd/yyyy-hh.mm.ss        INTERVAL hhh.mm.ss
                      RPT VERSION V1R2 RMF     END   mm/dd/yyyy-hh.mm.ss        CYCLE s.ttt SECONDS
```

Figure 5-3. 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 condition names 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

How to produce 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 the report, specify
`REPORTS(CACHE(options))`

Purpose of the report

The Cache Subsystem Activity report provides cache statistics on a subsystem basis as well as on a detailed device-level basis.

Using the information given in the report

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.

Different report levels

The contents of the report depends on the reporting options:

`REPORTS(CACHE(SUMMARY))` - Cache Summary Reporting

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).

`REPORTS(CACHE(SUBSYS))` - Subsystem-level Reporting

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.

For each 2105 subsystem, a RAID Rank Activity section is part of the report. A RAID rank is a set of physical volumes implementing RAID-5 in the 2105 subsystem. Several logical volumes as well as parallel access volumes are associated with a single RAID rank. Such a subsystem consists of multiple (up to 16) 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 suboptions `SSID/EXSSID` can be used to select or exclude specific control units.

`REPORTS(CACHE(DEVICE))` - Device-level Reporting

This generates, in addition to the report described above, 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.

These two ways can lead to higher I/O rates in the Cache Subsystem Activity report than in the DASD Activity report.

Contents of the report

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 V1R4               SYSTEM ID   SYS1           DATE 06/19/2002           INTERVAL 15.00.055
                                   RPT VERSION V1R2 RMF       TIME 12.35.59

SSID CU-ID  TYPE      CACHE  NVS    I/O   OFF  --CACHE HIT RATE-   --------DASD I/O RATE--------   ASYNC  TOTAL   READ  WRITE      %
                                   RATE  RATE   READ   DFW   CFW   STAGE DFWBP   ICL   BYP OTHER    RATE    H/R    H/R    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 5-4. 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.

```
                                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 V1R4               SYSTEM ID   SYS1           DATE 06/19/2002           INTERVAL 15.00.055
                                   RPT VERSION V1R2 RMF       TIME 12.35.59

*** DEVICE LIST BY DASD I/O RATE ***

VOLUME     DEV  SSID     %    I/O   ---CACHE HIT RATE--   ----------DASD I/O RATE----------   ASYNC   TOTAL    READ   WRITE      %
SERIAL     NUM          I/O   RATE    READ    DFW    CFW   STAGE  DFWBP    ICL    BYP  OTHER    RATE     H/R     H/R     H/R   READ

NP0LBD    6212  0020  54.9   80.8    60.9    0.2    0.0    19.7    0.0    0.0    0.0    0.0     0.1   0.756   0.755   0.975   99.8
NP0LBG    6215  0020  17.6   25.9    11.3    0.0    0.0    14.6    0.0    0.0    0.0    0.0     0.0   0.436   0.436     N/A  100.0
NP1MHI    1E85  3007  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
NP0LB3    6208  0020   8.3   12.2     1.5    5.5    0.0     5.2    0.0    0.0    0.0    0.0     5.0   0.570   0.217   0.999   54.9
DSMP30    1F43  3008  22.8   11.6     4.2    2.7    0.0     4.7    0.0    0.0    0.0    0.0     4.0   0.596   0.474   1.000   76.8
DSMP25    1EC7  3007   6.3   13.0     4.4    4.4    0.0     4.2    0.0    0.0    0.0    0.0     6.6   0.678   0.513   1.000   66.2
 ...

*** DEVICE LIST BY TOTAL I/O RATE ***

VOLUME     DEV  SSID     %    I/O   ---CACHE HIT RATE--   ----------DASD I/O RATE----------    ASYNC   TOTAL    READ   WRITE      %
SERIAL     NUM          I/O   RATE    READ    DFW    CFW   STAGE  DFWBP    ICL    BYP  OTHER     RATE     H/R     H/R     H/R   READ

WSF213    091F  0010  96.5  272.5   272.5    0.0    0.0     0.0    0.0    0.0    0.0    0.0      0.0   1.000   1.000   1.000   100.0
NP1LSU    2F92  3039  99.8  146.1   142.5    0.0    0.0     3.6    0.0    0.0    0.0    0.0      0.0   0.975   0.975     N/A   100.0
IDP507    0149  0041  80.0  140.1     0.7  139.4    0.0     0.0    0.0    0.0    0.0    0.0      6.4   1.000   1.000   1.000     0.5
APPP13    6762  1021  39.2  123.2   123.0    0.0    0.0     0.2    0.0    0.0    0.0    0.0      0.0   0.999   0.999     N/A   100.0
 ...
```

Figure 5-5. Cache Subsystem Summary Report - Top-20 Device Lists

Table 5-3. 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 5-4 on page 5-10 and Table 5-8 on page 5-16.	
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).
% 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.

Cache subsystem satus 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
                                                
```

PP - Cache Report

Table 5-4. 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.
Note: 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 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.	
CDATE	Date on which the cache interval started.
CTIME	Time at which the cache interval started.
CINT	Cache interval time. In interval reports, the format is *mm.ss*, while in duration reports the format is *hh.mm.ss*.
TYPE-MODEL	Hardware type and model.
VOLSER	Volume serial number of the reported DASD device (only for device-level reporting).
NUM	Volume serial number of the reported DASD device (only for device-level reporting).

Table 5-5. Fields in the Cache Subsystem Activity Report - Subsystem Status

Field Heading	Meaning
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.	

Table 5-5. Fields in the Cache Subsystem Activity Report - Subsystem Status (continued)

Field Heading	Meaning
CACHING	Overall caching status of the subsystem. **ACTIVE** Caching is active (online and usable). **DISABLED FOR MAINTENANCE** Cache has been disabled for maintenance. **PENDING ACTIVE STATE** Caching is pending active, that is, cache is in the process of being brought online. **INTERNAL ERROR TERMINATION** An internal error stopped cache (cache is offline) **EXPLICIT HOST TERMINATION** Cache has been deactivated by request from host system or support facility. **DEACTIVATION IN PROCESS** A command requesting deactivation of cache has been received, and destaging from cache to DASD is still in progress. **DEACTIVATION FAILED** A command requesting deactivation of cache has been received, but destaging to DASD has failed.
NON-VOLATILE STORAGE	Overall status of the non-volatile storage (NVS). **ACTIVE** NVS is online and usable. **PENDING DUE TO ERROR** A command requesting deactivation of NVS has been received but transfer from NVS to DASD has failed. **DEACTIVATION IN PROCESS** A command requesting deactivation of NVS has been received, and destaging to DASD is still in progress. **DISABLED FOR MAINTENANCE** NVS has been disabled for maintenance by the support facility. **INTERNAL ERROR TERMINATION** An internal error caused termination of NVS. **EXPLICIT HOST TERMINATION** NVS has been deactivated by request from host system or support facility. **DASD FAST WRITE INHIBITED** DASD FAST WRITE is inhibited because the battery is defective.
CACHE FAST WRITE	Status of the CACHE FAST WRITE (CFW) option. **ACTIVE** CFW is active. **DEACTIVATED** CFW is deactivated. **Note:** CFW does not use NVS.
IML DEVICE AVAILABLE	Status of the IML device. **NO** Device containing a diskette drive for loading the microcode is not available. **YES** Device is operational.

PP - Cache Report

Table 5-6. 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	Total number of cacheable I/O requests to cached devices in the storage subsystem. This value excludes INHIBIT CACHE LOAD and CACHE BYPASS I/O requests.
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	Cache will be managed by *least-recently-used (LRU)* 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 described above.
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).

Table 5-6. Fields in the Cache Subsystem Activity Report - Subsystem Overview (continued)

Field Heading	Meaning
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 *inhibited* 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.	
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. **Note:** 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.

Table 5-6. Fields in the Cache Subsystem Activity Report - Subsystem Overview (continued)

Field Heading	Meaning
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 not found in the cache, and access to DASD was required.

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.

The second section in this part of the report 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 it lists the volumes with the highest activity.

Due to hardware restrictions, RAID rank data are only available, if the device with the lowest unit address in a RAID rank is online. For details, see *IBM TotalStorage ESS DFSMS Software Support Reference (SC26-7440-00)*.

```
                                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 V1R4               SYSTEM ID   SYS1           DATE 06/19/2002            INTERVAL 15.00.055
                                   RPT VERSION V1R2 RMF       TIME 12.35.59

SUBSYSTEM  2105-01     CU-ID  1E91     SSID 3007     CDATE  06/02/2002     CTIME 23.29.20      CINT  15.04
TYPE-MODEL 2105-E20

------------------------------------------------------------------------------------------------------------------------------------------------------
                                                        CACHE SUBSYSTEM DEVICE OVERVIEW
------------------------------------------------------------------------------------------------------------------------------------------------------

VOLUME     DEV  RRID     %    I/O   ---CACHE HIT RATE--   ----------DASD I/O RATE----------    ASYNC   TOTAL    READ   WRITE       %

SERIAL     NUM         I/O   RATE    READ    DFW    CFW   STAGE  DFWBP    ICL    BYP  OTHER     RATE     H/R     H/R     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

...

------------------------------------------------------------------------------------------------------------------------------------------------------
                                                          RAID RANK ACTIVITY
------------------------------------------------------------------------------------------------------------------------------------------------------

ID    RAID     DA   HDD    -------- READ REQ -------     ------- WRITE REQ -------   ---------- HIGHEST UTILIZED VOLUMES ----------
      TYPE                 RATE  AVG MB  MB/S  RTIME     RATE  AVG MB  MB/S  RTIME

*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 5-7. Cache Subsystem Activity Report - Device Overview

Table 5-7. 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	RAID rank 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.

Table 5-7. Fields in the Cache Subsystem Activity Report - Device Overview (continued)

Field Heading	Meaning
DFWBP	Rate of requests that caused DFW BYPASS.
ICL	Rate of Inhibit cache load requests.
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 5-8. 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.
READ and WRITE Requests	
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.
Volume Utilization	
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 5-6.

```
                                    C A C H E   D E V I C E   A C T I V I T Y
                                                                                                                                  PAGE     2


           z/OS V1R4               SYSTEM ID   SYS1           DATE 06/19/2002           INTERVAL 15.00.055
                                   RPT VERSION V1R2 RMF       TIME 12.35.59

SUBSYSTEM  2105-01     CU-ID  1E91     SSID 3007     CDATE  06/02/2002     CTIME 23.29.20      CINT  15.04
TYPE-MODEL 2105-E20
VOLSER     NP1MHI      NUM    1E85     RRID 0700
-------------------------------------------------------------------------------------------------------------------------------------
                                                  CACHE DEVICE STATUS
-------------------------------------------------------------------------------------------------------------------------------------

CACHE STATUS                                    DUPLEX PAIR STATUS

CACHING             - ACTIVE                    DUPLEX PAIR        - NOT ESTABLISHED
DASD FAST WRITE     - ACTIVE                    STATUS             - N/A
PINNED DATA         - NONE                      DUAL COPY VOLUME   - N/A

-------------------------------------------------------------------------------------------------------------------------------------
                                                  CACHE DEVICE ACTIVITY
-------------------------------------------------------------------------------------------------------------------------------------

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
                                                                    DFW BYPASS          0    0.0          ICL             0    0.0
NORMAL        6730     7.4       0     0.0    6899     7.6          CFW BYPASS          0    0.0          BYPASS          0    0.0
SEQUENTIAL     172     0.2       0     0.0    6026     6.7          DFW INHIBIT         0    0.0          TOTAL           0    0.0
CFW DATA         0     0.0       0     0.0                          ASYNC (TRKS)        0    0.0

TOTAL         6902    RATE     7.6

---CKD STATISTICS---     ---RECORD CACHING---

WRITE              0     READ MISSES        0

WRITE HITS         0     WRITE PROM         0
```

Figure 5-8. Cache Subsystem Activity Report - Device-level Reporting

PP - Cache Report

Table 5-9. 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	Caching status of the device. **ACTIVE** Caching is active; requests to the reported device can be processed without DASD access. **DEACTIVATE PENDING** Cache has been deactivated on request from host system or support facility, but transfer of modified data to DASD has failed. **DEACTIVATED** Caching has been deactivated for the reported device.
DASD FAST WRITE	Status of the DASD FAST WRITE option. **ACTIVE** DASD FAST WRITE requests can be processed for this device. **DEACTIVATION PENDING** 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. **DEACTIVATED** DASD FAST WRITE requests are ignored for this device.
PINNED DATA	A device has failed, and data that has not yet been written to DASD is pinned in cache or NVS for later recovery. **NONE** No data is pinned for the reported device. **EXISTS, DFW NOT SUSPENDED** Pinned data exists, but DASD FAST WRITE has not been suspended. **EXISTS, DFW SUSPENDED** Pinned data exists, and DASD FAST WRITE has been suspended.
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.	
DUPLEX PAIR	Overall dual copy status. **NOT ESTABLISHED** The device is configured as a simplex device. **ACTIVE** The device reported is either the primary or the secondary device of a duplex pair. **PENDING** The control unit copies data to establish the duplex pair to which the reported device belongs. **SUSPENDED** An error in one device of the duplex pair caused the duplex pair to be suspended.

Table 5-9. Fields in the Cache Subsystem Activity Report - Cache Device Status (continued)

Field Heading	Meaning
STATUS	The status of this device as part of a duplex pair. **PRIMARY** The device is the primary device of the duplex pair. All channel operations are associated with this device. **SECONDARY** 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. **N/A** The device has not been established as part of a duplex pair.
DUAL COPY VOLUME	Identification of the other device of a duplex pair. **nnnn** Device address of the other device of the duplex pair, if a duplex pair is established. **N/A** The device has not been established as part of a duplex pair.

The report contains less information for devices which had the cache offline at the end of the interval.

```
                                C A C H E   S U B S Y S T E M   A C T I V I T Y
                                                                                                                    PAGE    01

           z/OS V1R4               SYSTEM ID   SYS1           DATE  04/15/2002       INTERVAL  12.00.000
                                   RPT VERSION V1R2 RMF       TIME  07.12.00

SUBSYSTEM  3990-06     CU-ID  02A0     SSID  0016     CDATE  04/15/2002    CTIME  07.12.00     CINT  12.00
TYPE-MODEL 9396-001
VOLSER     RZ3381      NUM    02B6     RRID  N/A
-------------------------------------------------------------------------------------------------------------------------------
                                                  CACHE DEVICE STATUS
-------------------------------------------------------------------------------------------------------------------------------

CACHE STATUS                                    DUPLEX PAIR STATUS

CACHING             - DEACTIVATED               DUPLEX PAIR NOT ESTABLISHED
DASD FAST WRITE     - DEACTIVATED
PINNED DATA         - NONE

-------------------------------------------------------------------------------------------------------------------------------
                                                  CACHE DEVICE ACTIVITY
-------------------------------------------------------------------------------------------------------------------------------

TOTAL I/O     ---     CACHE I/O     ---     CACHE OFFLINE     8772
TOTAL H/R     ---     CACHE H/R     ---
```

Figure 5-9. 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 5-10. Overview Names in the Cache Subsystem Activity Report

Column Heading	Subsystem Report	Device Report
Subsystem Status / Device Status.		
CACHING	CASSC	CADSC
NON-VOLATILE STORAGE	CASSNVS	

Table 5-10. Overview Names in the Cache Subsystem Activity Report (continued)

Column Heading	Subsystem Report	Device Report
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
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

Table 5-10. Overview Names in the Cache Subsystem Activity Report (continued)

Column Heading	Subsystem Report	Device Report
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
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

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)`

Contents of the report

A Coupling Facility Activity report is produced for each coupling facility attached to the sysplex. Figure 5-10 gives an example of the overall structure of the Coupling Facility Activity report. It shows the sequencing of the report sections:

- Usage Summary
- Structure Activity
- Subchannel Activity
- CF to CF Activity

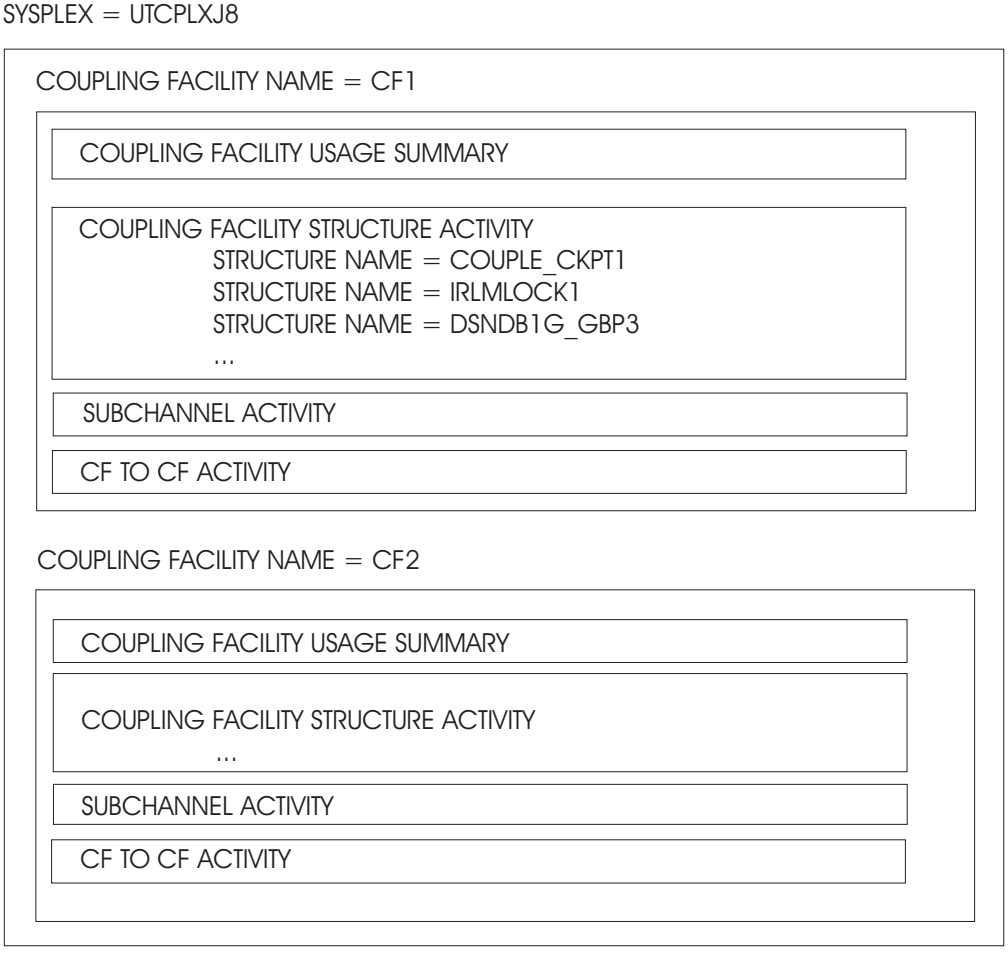

Figure 5-10. Structure of the Coupling Facility Report

Usage Summary section

This section of the Coupling Facility Activity report gives a snapshot of coupling facility storage usage at the end of the reporting interval. It lists all the structures occupying space in the coupling facility, including those with no currently active connections. The structure summary data is grouped in the report by structure type (LIST, LOCK, CACHE and UNKN in that order). Within 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.

```
                                C O U P L I N G   F A C I L I T Y   A C T I V I T Y
                                                                                                                                                                PAGE   1
     z/OS V1R4              SYSPLEX UTCPLXJ8          DATE 04/11/2002          INTERVAL 030.00.000
                            RPT VERSION V1R2 RMF      TIME 12.00.00            CYCLE 01.000 SECONDS

 ------------------------------------------------------------------------------------------------------------------------------------------------------------------
 COUPLING FACILITY NAME = CF1
 TOTAL SAMPLES(AVG) =  1781  (MAX) =  1799  (MIN) =  1671
 ------------------------------------------------------------------------------------------------------------------------------------------------------------------
                                              COUPLING  FACILITY  USAGE  SUMMARY
 ------------------------------------------------------------------------------------------------------------------------------------------------------------------
 STRUCTURE SUMMARY
 ------------------------------------------------------------------------------------------------------------------------------------------------------------------

                                                 % OF                 % OF      AVG    LST/DIR  DATA      LOCK     DIR REC/
       STRUCTURE                      ALLOC      CF            #      ALL       REQ/   ENTRIES  ELEMENTS  ENTRIES  DIR REC
 TYPE  NAME              STATUS CHG   SIZE       STORAGE     REQ      REQ       SEC    TOT/CUR  TOT/CUR   TOT/CUR  XI'S

 LIST  DSNDB1G_SCA       ACTIVE         16M          0.8     3395      0.1      1.89      26K       52K      N/A       N/A
                                                                                            92       158      N/A       N/A
       IXCPLEX_PATH1     ACTIVE         59M          2.9   257609      8.6    143.12      14K       14K      N/A       N/A
                                                                                             1        48      N/A       N/A
...

 LOCK  DSNDB1G_LOCK1     ACTIVE         31M          1.5    64373      2.2     35.76     114K         0     4194K       N/A
                                                                                            41         0     2730       N/A
       IRLMLOCK1         ACTIVE         31M          1.5     1676K    56.0    930.89     114K         0     4194K       N/A
                                                                                             6         0      256       N/A

 CACHE DSNDB1G_GBP0      ACTIVE         21M          1.0     3289      0.1      1.83      22K      4321      N/A         0
                          PRIM                                                              58        35      N/A         0
       DSNDB1G_GBP4      ACTIVE         21M          1.0     3289      0.1      1.83      22K      4321      N/A         0
                          SEC                                                               58        35      N/A         0
...
                                      -------       -----   -------   ------   -------
            STRUCTURE TOTALS           263M         12.9     3077K      100   1709.2

  STORAGE SUMMARY
 ------------------------------------------------------------------------------------------------------------------------------------------------------------------

                                          ALLOC     % OF CF     ------- DUMP SPACE -------
                                          SIZE      STORAGE     % IN USE   MAX % REQUESTED

  TOTAL CF STORAGE USED BY STRUCTURES      263M       12.9
  TOTAL CF DUMP STORAGE                      6M        0.3          0.0              0.0
  TOTAL CF STORAGE AVAILABLE                 2G       86.8
                                         -------
  TOTAL CF STORAGE SIZE                      2G

                                          ALLOC     % ALLOCATED
                                          SIZE

  TOTAL CONTROL STORAGE DEFINED              2G       13.2
  TOTAL DATA STORAGE DEFINED                 0K        0.0

  PROCESSOR SUMMARY
 ------------------------------------------------------------------------------------------------------------------------------------------------------------------

  COUPLING FACILITY         9672        MODEL C01      CFLEVEL 3
  AVERAGE CF UTILIZATION  (% BUSY)           7.8      LOGICAL PROCESSORS:   DEFINED   6    EFFECTIVE   6.0
```

Figure 5-11. 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 5-11. Fields in the Coupling Facility Activity Report - Usage Summary

Field Heading	Meaning
TOTAL SAMPLES (AVG), (MIN), (MAX)	Average (AVG), minimum (MIN) and maximum (MAX) sample count for all systems connected to this Coupling Facility.
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: **ACTIVE** 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. **ACTIVE PRIM** The structure is the rebuilt-old (primary) structure in a duplexing rebuild process. **ACTIVE SEC** The structure is the rebuilt-new (secondary) structure in a duplexing rebuild process. **INACTV** 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. **UNALLOC** 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.
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 STORAGE	The percentage of the total coupling facility storage allocated to this structure.

Table 5-11. 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.
AVG REQ/SEC	The average number of requests per second for this structure.
LST/DIR ENTRIES TOT/CUR	**TOT** The total number of list or directory entries in the structure. **CUR** The number of list or directory entries currently in use. "N/A" in this field and the following indicates that the information is not applicable.
DATA ELEMENTS TOT/CUR	**TOT** The total number of data elements in the structure. **CUR** The number of data elements currently in use.
LOCK ENTRIES TOT/CUR	**TOT** The total number of lock table entries. **CUR** 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 *Usage Summary Report*). A high value is an indicator for a performance problem in this structure.
STORAGE SUMMARY	
TOTAL CF STORAGE USED BY STRUCTURES	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. It should be noted that these totals don't 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.
TOTAL CF DUMP STORAGE	Amount of coupling facility space allocated as dump space.
DUMP SPACE: % IN USE	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.

Table 5-11. Fields in the Coupling Facility Activity Report - Usage Summary (continued)

Field Heading	Meaning
DUMP SPACE: MAX % REQUESTED	The maximum percentage of dump space requested since the coupling facility dump storage was allocated. 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.
TOTAL CF STORAGE AVAILABLE	The amount of coupling facility space that is not allocated to any structure and not allocated as dump space.
TOTAL CF STORAGE SIZE	The total amount of storage in the coupling facility, including both allocated and available space.
TOTAL CONTROL STORAGE DEFINED, TOTAL DATA STORAGE DEFINED, % ALLOC	The amount of coupling facility storage that is allowed to be occupied by control information (CONTROL STORAGE) or data (DATA STORAGE). Each structure plus the dump area is allocated some control and some data storage. 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. If % ALLOC 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: • Changing structure preference lists in the coupling facility policy specification to direct some structures away from this facility. • Adding another coupling facility to the sysplex.
PROCESSOR SUMMARY	
COUPLING FACILITY	Coupling facility processor type.
MODEL	Coupling facility processor model.
CFLEVEL	Coupling facility CF level.
AVG. CF UTILIZATION (% BUSY)	Average value of CPU utilizations within the coupling facility. The utilization of the individual CPs in the coupling facility is recorded in the SMF 74, Subtype 4, Processor Data Section. 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. If the average utilization is high, you can take the following actions: 1. In a PR/SM environment, you can dedicate the CP to the integrated coupling facility or assign additional CPs to the partition. 2. Move structures to a coupling facility with lower utilization. 3. Consider additional or larger coupling facilities.

Table 5-11. Fields in the Coupling Facility Activity Report - Usage Summary (continued)

Field Heading	Meaning
LOGICAL PROCESSORS DEFINED	Number of logical processors defined for the coupling facility.
LOGICAL PROCESSORS EFFECTIVE	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 the time waiting for work. The reported value shows the ratio between the LPAR dispatch time (CFCC execute and wait time) and the RMF interval length. 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 Please, refer to the CPU Activity report in case of an ICMF LPAR.

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.

```
                              C O U P L I N G   F A C I L I T Y   A C T I V I T Y
                                                                                                                                  PAGE   3
    z/OS V1R4              SYSPLEX PLEXPERF          DATE 04/11/2002          INTERVAL 030.00.000
                           RPT VERSION V1R2 RMF      TIME 12.00.00            CYCLE 01.000 SECONDS

-------------------------------------------------------------------------------------------------------------------------------------
 COUPLING FACILITY NAME = CF1
-------------------------------------------------------------------------------------------------------------------------------------
                                        COUPLING  FACILITY  STRUCTURE  ACTIVITY
-------------------------------------------------------------------------------------------------------------------------------------

   STRUCTURE NAME = COUPLE_CKPT1      TYPE = LIST             STATUS = ACTIVE
             # 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   /DEL    STD_DEV   /ALL   CONTENTIONS

   J80        8463    SYNC   2927     5.8   306.6    127.8     NO SCH 1240   22.4  508.4      597.7  113.9   REQ TOTAL     8339
              4.70    ASYNC  5535    11.0  1502.1   1263.2     PR WT     0    0.0    0.0        0.0    0.0   REQ DEFERRED   123
                      CHNGD     1     0.0  INCLUDED IN ASYNC   PR CMP    0    0.0    0.0        0.0    0.0
                                                               DUMP      0    0.0    0.0        0.0    0.0

...

  -----------------------------------------------------------------------------------------------------------------------------------
   TOTAL     50370    SYNC     17K   34.3   342.2    124.5     NO SCH 8797   26.6   1047       1245  278.4   REQ TOTAL       50K
             27.98    ASYNC    33K   65.5  557551G     0.0     PR WT     0    0.0    0.0        0.0    0.0   REQ DEFERRED   497
                      CHNGD    69     0.1                      PR CMP    0    0.0    0.0        0.0    0.0
                                                               DUMP      0    0.0    0.0        0.0    0.0
...
   STRUCTURE NAME = IRLMLOCK1         TYPE = LOCK             STATUS = ACTIVE
             # 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   /DEL    STD_DEV   /ALL   CONTENTIONS

...

   JF0         128K   SYNC    128K    7.7   221.9     37.5     NO SCH    0    0.0    0.0        0.0    0.0   REQ TOTAL      128K
             71.26    ASYNC     0     0.0     0.0      0.0     PR WT     0    0.0    0.0        0.0    0.0   REQ DEFERRED    80
                      CHNGD     0     0.0  INCLUDED IN ASYNC   PR CMP    0    0.0    0.0        0.0    0.0   -CONT           80
                                                               DUMP      0    0.0    0.0        0.0    0.0   -FALSE CONT     25
...

  -----------------------------------------------------------------------------------------------------------------------------------
   TOTAL      1676K   SYNC   1676K    100   198.1     44.7     NO SCH    0    0.0    0.0        0.0    0.0   REQ TOTAL     1674K
             930.9    ASYNC     0     0.0     0.0      0.0     PR WT     0    0.0    0.0        0.0    0.0   REQ DEFERRED   906
                      CHNGD     0     0.0                      PR CMP    0    0.0    0.0        0.0    0.0   -CONT          906
                                                               DUMP      0    0.0    0.0        0.0    0.0   -FALSE CONT    225
...
   STRUCTURE NAME = DSNDB1G_GBP3      TYPE = CACHE            STATUS = ACTIVE PRIMARY
             # 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   /DEL    STD_DEV   /ALL
...

   JA0         1032    SYNC   1025    19.1   448.8     84.3     NO SCH    7    100   2060       1671   2060
               0.57    ASYNC     0     0.0  4508.9   1409.1     PR WT     0    0.0    0.0        0.0    0.0
                       CHNGD     7     0.1  INCLUDED IN ASYNC   PR CMP    0    0.0    0.0        0.0    0.0
                                                               DUMP      0    0.0    0.0        0.0    0.0

...

  -----------------------------------------------------------------------------------------------------------------------------------
   TOTAL       5356    SYNC   5342     100   432.8     89.0     NO SCH   14    100   1510       1374   1510    -- DATA ACCESS ---
               2.98    ASYNC     0     0.0  3685.2   1634.2     PR WT     0    0.0    0.0        0.0    0.0    READS         1331
                       CHNGD    14     0.3                      PR CMP    0    0.0    0.0        0.0    0.0    WRITES        1709
                                                               DUMP      0    0.0    0.0        0.0    0.0    CASTOUTS        44
                                                                                                              XI'S          1296
```

Figure 5-12. Coupling Facility Activity Report - Structure Activity

The following table explains the field headings in the XES Structure Activity section.

Table 5-12. 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 5-11 on page 5-24.
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 AVG/SEC	The sum of synchronous requests and asynchronous requests completed, and requests that changed from SYNC to ASYNC. **TOTAL** Total number of requests **AVG/SEC** 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 three categories.	
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 started synchronously but completed asynchronously to the coupling facility on behalf of connectors to the structure.
# REQ % OF ALL REQ	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, the coupling facility configuration should be analyzed for possible path or coupling facility bottlenecks. (See the Subchannel Activity report.)
DELAYED REQUESTS - These columns list possible contention reasons for requests sent to the coupling facility.	

Table 5-12. 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 duplexed requests, also peer subchannel wait contention (PR WT) and waiting-for-peer-completion contention (PR CMP) is reported. A duplexed request requires two subchannels. A 'PR WT' contention means that for the reported conditions a subchannel for the operation targeted to the peer structure was not available. A 'PR CMP' contention means that one of the two duplexed operations has completed, but the completed subchannel remains unavailable for use until the peer operation completes.
# REQ % of REQ	The total number and the percentage of requests delayed in the interval.
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. **Example**: 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. Note: 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.

Table 5-12. Fields in the Coupling Facility Activity Report - Structure Activity (continued)

Field Heading	Meaning
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.
CASTOUTS	The number of times CASTOUT processing occurs. This is the process of writing changed cache data to permanent storage. 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.
XI'S	The number of times a data item residing in a local buffer pool was marked invalid by the coupling facility. 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.

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.

```
                              C O U P L I N G   F A C I L I T Y   A C T I V I T Y

  z/OS V1R4              SYSPLEX PLEXDT            DATE 09/11/2002          INTERVAL 030.00.000
                         RPT VERSION V1R2 RMF      TIME 13.00.00            CYCLE 01.000 SECONDS
-------------------------------------------------------------------------------------------------------------------------------
 COUPLING FACILITY NAME = CF1
-------------------------------------------------------------------------------------------------------------------------------
                                                  SUBCHANNEL  ACTIVITY
-------------------------------------------------------------------------------------------------------------------------------
          # REQ                              ---------- REQUESTS ----------   ------------------ DELAYED REQUESTS --------------
 SYSTEM   TOTAL    -- CF LINKS --  PTH            #   -SERVICE TIME(MIC)-             #     % OF  ------ AVG TIME(MIC) ------
 NAME     AVG/SEC TYPE  GEN  USE  BUSY           REQ     AVG    STD_DEV              REQ    REQ   /DEL    STD_DEV     /ALL

*W04       30740  ICP     3    3     0   SYNC       957    26.8      95.5   LIST/CACHE    2   0.0  174.0       50.9    348.0
            51.2  SUBCH  42   21         ASYNC    22791   736.7      2811   LOCK          0   0.0    0.0        0.0      0.0
                                          CHANGED      0  INCLUDED IN ASYNC  TOTAL         2   0.0
                                          UNSUCC       0     0.0       0.0
*W05        9793  CFS     2    2   354   SYNC       790    43.3     116.5   LIST/CACHE 1537  19.1   1797       5626     2761K
             9.6  SUBCH   4    4         ASYNC     8100   937.8      3543   LOCK          4   0.5   69.5       15.6    278.0
                                          CHANGED      1  INCLUDED IN ASYNC  TOTAL      1541  17.3
                                          UNSUCC       0     0.0       0.0
```

Figure 5-13. Coupling Facility Activity Report - Subchannel Activity

Table 5-13. Fields in the Coupling Facility Activity Report - Subchannel Activity

Field Heading	Meaning
SYSTEM NAME	The name of the system attached to the coupling facility (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.

Table 5-13. Fields in the Coupling Facility Activity Report - Subchannel Activity (continued)

Field Heading	Meaning
# REQ TOTAL AVG/SEC	**TOTAL** 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. **AVG/SEC** Average number of requests per second for this facility. 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.
CF LINKS	**TYPE** Channel path type. **GEN** Number of subchannels that are defined. **USE** Number of subchannels MVS is currently using for coupling facility requests.
PTH BUSY	Path busy - the number of times a coupling facility request was rejected because all paths to the coupling facility were busy. 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. 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. 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.
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.

Table 5-13. Fields in the Coupling Facility Activity Report - Subchannel Activity (continued)

Field Heading	Meaning
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.
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.

CF to CF Activity section

```
------------------------------------------------------------------------------------------------------------------------
  COUPLING FACILITY NAME = CF1
------------------------------------------------------------------------------------------------------------------------
                                        CF TO CF  ACTIVITY
------------------------------------------------------------------------------------------------------------------------
          # REQ                       ---------- REQUESTS ------------   ------------- DELAYED REQUESTS --------------
PEER      TOTAL  -- CF LINKS --              #   -SERVICE TIME(MIC)-              #    % OF  ---- AVG TIME(MIC) ---
CF       AVG/SEC TYPE       USE             REQ     AVG    STD_DEV              REQ   REQ   /DEL  STD_DEV    /ALL

CF2       86830  CBP          1       SYNC   86830    53.0        33.8   SYNC         19   0.2%   14.5     2.7     0.0
           86.2  ICP          1
```

Figure 5-14. CF to CF Activity

Table 5-14. Fields in the CF to CF Activity Section

Field Heading	Meaning
PEER CF	Name of the remote coupling facility.
# REQ TOTAL	The total number of requests delayed in the interval. This is the sum of the following signals which have been sent from the subject CF to this remote CF: • Number of halt execution signals • Number of ready to complete signals • Number of ready to execute signals • Number of request for suppression signals • Number of request for suppression accepted signals
#REQ AVG/SEC	Average number of requests per second.
CF LINKS TYPE	CF to CF link type.
CF LINKS USE	The number of links used for coupling facility communication.
REQUESTS - The requests are synchronous (SYNC).	
# REQ	The number of requests for this coupling facility.

Table 5-14. Fields in the CF to CF Activity Section (continued)

Field Heading	Meaning
SERVICE TIME - AVG	The average service time in microseconds (based on the difference between the sum of signal delay time and sum of signal service time).
SERVICE TIME - STD_DEV	The standard deviation to 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.

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 5-15. Overview Names in the Coupling Facility Activity Report

Criterion	Overview Name
Average service time of SYNC operations	SYNCST
SYNC operation rate	SYNCRT
Average service time of ASYNC operations	ASYNCST
Ended ASYNC operation rate	ASYNCRT
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
False contention on lock structure	LCKFCONT

CHAN - Channel Path Activity report

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 a report, specify
REPORTS(CHAN)

Purpose of the 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 and the appropriate data cannot be gathered, dashes(---) will be reported instead of data.

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 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.

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.

Using the information given in the report

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.

Contents of the report

```
                                C H A N N E L   P A T H   A C T I V I T Y
                                                                                                                    PAGE    1
         z/OS V1R4               SYSTEM ID   CB88          DATE 07/22/2002           INTERVAL 22.54.336
                                 RPT VERSION V1R2 RMF      TIME 03.07.05             CYCLE 1.000 SECONDS

IODF = 76   CR-DATE: 05/22/2002   CR-TIME: 18.32.08   ACT: ACTIVATE     MODE: LPAR        CPMF: EXTENDED MODE             CSSID: 0
 ------------------------------------------------------------------------------------------------------------------------------------------
                                            OVERVIEW FOR DCM-MANAGED CHANNELS
 ------------------------------------------------------------------------------------------------------------------------------------------

     CHANNEL          UTILIZATION(%)    READ(MB/SEC) WRITE(MB/SEC)
    GROUP  G NO    PART  TOTAL    BUS   PART  TOTAL   PART  TOTAL

    FC_SM  1  8   15.36  55.86   6.00  15.36  60.00  15.36  60.36
    FCV_M    12   30.00  45.00   5.00  45.00  50.00  45.00  50.00
    CNC_M     1   17.23  34.45

 ------------------------------------------------------------------------------------------------------------------------------------------
                                            DETAILS FOR ALL CHANNELS
 ------------------------------------------------------------------------------------------------------------------------------------------

    CHANNEL PATH     UTILIZATION(%)    READ(MB/SEC) WRITE(MB/SEC)     CHANNEL PATH     UTILIZATION(%)    READ(MB/SEC) WRITE(MB/SEC)
  ID TYPE   G SHR  PART  TOTAL    BUS   PART  TOTAL   PART  TOTAL   ID TYPE   G SHR  PART  TOTAL    BUS   PART  TOTAL   PART  TOTAL

  78 CVC_P        OFFLINE                                          80 CTC_S        OFFLINE
  79 CNC_S        OFFLINE                                          81 CNC_S         0.04   0.04
  7A FC     1  Y  20.00  30.00   5.00  20.00  30.00  20.00  50.00  82 FC     1  Y  20.00  30.00   6.00  20.00  30.00  20.00  50.00
  7B FC_SM     Y  15.36  55.86   6.00  15.36  60.00  15.36  60.36  83 FC        Y  15.36  55.66   7.00  15.36  60.00  15.36  60.00
  7C FCV       Y  10.00  30.00   5.00  10.00  50.00  10.00  50.00  84 FCV       Y  10.00  30.00   5.00  10.00  50.00  50.00  10.00
  7D FCV_M     Y  30.00  45.00   5.00  45.00  50.00  45.00  50.00  85 FCV       Y  30.00  45.00   6.00  45.00  50.00  45.00  50.00
  7E CNC_M        17.23  34.45                                      86 CNC_S         0.00   0.00
  7F CNC_S        OFFLINE                                          8C CNC_S         0.00   0.00

  98 CNC_P     Y   0.76   5.91                                      AD CNC_?        OFFLINE
  99 CNC_S     Y   0.75   5.82                                      AE CNC_S        OFFLINE
  9A CNC_?        OFFLINE                                          AF CNC_P        OFFLINE
  9B CNC_?        OFFLINE                                          A4 ISD          OFFLINE
  9C CNC       Y   1.79  13.50                                      A6 CVC_S        OFFLINE
  9D CNC_S     Y   1.87  13.78                                      A7 CVC_P        OFFLINE
  9E CNC_S     Y   1.12   6.79
  9F CNC_P     Y   2.42   3.51


    CHANNEL PATH     WRITE(B/SEC)     MESSAGE RATE     MESSAGE SIZE   SEND FAIL    RECEIVE FAIL
  ID TYPE   G SHR     PART   TOTAL     PART   TOTAL     PART   TOTAL     PART       PART   TOTAL

  AB IQD       Y   712.26M   1.91G  887.16K   2.41M   803.01  792.44     0.13          0       0
```

Figure 5-15. Channel Path Activity Report

Table 5-16. 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.
MODE	The mode of the central processing complex (CPC): **BASIC** The report shows all channels configured in the system. **LPAR** The report shows both the total utilization and the individual partition's utilization of all channels configured to the logical partition.

Table 5-16. Fields in the Channel Path Activity Report (continued)

Field Heading	Meaning
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: **COMPATIBILITY MODE** CPMF is running in compatibility mode. **EXTENDED MODE** CPMF is running in extended mode. **NOT AVAILABLE** CPMF is not available on the system. The indication (CHANGED) will be shown if the CPMF mode has been 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 *z/OS MVS Data Areas, Vol 2 (DCCB-ITZYRETC)*.
CSSID	This field is shown only for z990 processors or later 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 1Gbit/sec, or a number 2 indicates a throughput of 2Gbit/sec. The number of channels of the group is given in column NO.
CHANNEL PATH ID	The hexadecimal number of the 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 2GB/sec.
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.
Note: 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).	

Table 5-16. Fields in the Channel Path Activity Report (continued)

Field Heading	Meaning

UTILIZATION (%) PART	The channel path utilization percentage for an individual logical partition. RMF uses the values provided by CPMF.
	
	In LPAR mode, the calculation is:
	`PART Channel Path Busy Time`
	`UTILIZATION (%) = ------------------------- * 100`
	`Channel Path Elapsed Time`
	
	In BASIC mode or if EMIF is not installed, blanks are shown.
	
	For OSAEGbE, the value reflects the microprocessor utilization.

UTILIZATION (%) TOTAL	The channel path utilization percentage for the entire system during an interval.
	
	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:
	`# SRM Observations of`
	`Total Channel Path Busy`
	`Utilization (%) = -------------------- * 100`
	`# Samples`
	
	For unshared channels in LPAR mode, the value for total utilization is the same as partition utilization.
	
	For all channels in BASIC mode with CPMF available, the calculation is:
	`TOTAL Channel Path Busy Time`
	`UTILIZATION (%) = ------------------------- * 100`
	`Channel Path Elapsed Time`
	
	For OSAEGbE, the value reflects the microprocessor utilization.

UTILIZATION (%) BUS	Percentage of bus cycles, the bus has been found busy for this channel in relation to the theoretical limit. For OSAEGbE, the value reflects the PCI bus utilization.
READ(MB/SEC)	**PART** Data transfer rates from the control unit to the channel for this partition. **TOTAL** Data transfer rates from the control unit to the channel for the entire system.
WRITE(MB/SEC)	**PART** Data transfer rates from the channel to the control unit for this partition. **TOTAL** Data transfer rates from the channel to the control unit for the entire system.
The following fields are for HiperSockets:	
WRITE(B/SEC)	**PART** Data transfer rates from the channel to the control unit for this partition. **TOTAL** Data transfer rates from the channel to the control unit for the entire system. The values are shown in bytes/second.
MESSAGE RATE	**PART** Rate of messages sent by this partition. **TOTAL** Rate of messages sent by the entire system.

Table 5-16. Fields in the Channel Path Activity Report (continued)

Field Heading	Meaning	
MESSAGE SIZE	**PART**	Average size of messages sent by this partition.
	TOTAL	Average size of messages sent by the entire system.
SEND FAIL PART	Rate of messages (sent by this partition) that failed, but not caused by an unavailable buffer in the receiving partition.	
RECEIVE FAIL	**PART**	Rate of messages (received by this partition) that failed due to unavailable buffers. The value could indicate, that more receive buffers are required.
	TOTAL	Rate of messages (received by the entire system) that failed due to unavailable 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 5-17. Overview Names in the Channel Path Activity Report

Column Heading	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:	
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
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

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 a report, specify
`REPORTS(CPU)`

Purpose of the report

The report is divided into the following sections:

CPU Activity
Provides information on the active processors

System Address Space Analysis
Provides overall address space information

The following two section are available only, if the system is running in a Processor Resource/Systems Manager (PR/SM) environment in LPAR mode, they are not available if the system is running in a VM guest environment.

Partition Data Report
Provides data about all configured partitions. This section is described in "Using the information in the Partition Data report" on page 5-46.

LPAR Cluster Report
Provides data about each LPAR cluster. This section is described in "Using the information in the LPAR Cluster report" on page 5-50.

Please, refer to the *z/OS RMF Performance Management Guide* for a more detailed discussion of the Partition Data report and the LPAR Cluster report.

Contents of the report

CPU activity

The first column of the CPU Activity report is headed by `CPU MODEL`. It identifies the machine which RMF is running on by processor family and model number.

For each processor, the report identifies the processor number, serial number, and the following calculations:
- Processor online time percentage of the interval
- Processor LPAR busy time percentage of the interval
- Processor MVS busy time percentage of the interval
- I/O interrupt rate
- Percentage of I/O interrupts handled by the I/O supervisor without re-enabling

If RMF is running as a guest under z/VM, it only reports the processor 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 VM monitor.

Following the data lines for each processor is a TOTAL/AVERAGE line.

Dependent on the mode, a system is running, the CPU utilization shown in RMF reports has a different meaning.

In Basic mode, the CPU busy time is calculated as the difference between the reporting interval time and the MVS WAIT time. This can be considered as the **MVS view** of the CPU time. It is also the physical (real) CPU time because there are only two possible states:
- CPU WAIT state
- CPU BUSY state (which means NON WAIT state)

In LPAR mode, the calculation of the CPU utilization is much more complex due to the different states that are possible:
- 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 partition has dedicated processors

However, performance analysts need both views of CPU utilization. The MVS view of the CPU utilization is a much better and direct indicator to see a CPU bottleneck. The LPAR view is important with respect to capacity aspects.

The **MVS CPU time** for one processor is calculated as:

```
CPU Time = Online Time - Wait Time
```

The **LPAR CPU Time** for one processor is calculated depending on the status of the logical processor:

Dedicated	`CPU Time = Online Time - Wait Time`
Wait Completion = YES	`CPU Time = Dispatch Time - Wait Time`
Wait Completion = NO	`CPU Time = Dispatch Time`

The CPU percentage is calculated as

```
          CPU Time
CPU(%) = ----------- * 100
         Online Time
```

The CPU Activity report presents both views, the LPAR view and the MVS view of the CPU utilization. In case of a Basic mode environment, the LPAR view of the CPU utilization contains dashes.

System address space analysis

The System Address Space Analysis section of the CPU activity report provides overall address space information. The number of SAMPLES taken during the RMF reporting interval is displayed at the top of this section.

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.

Using the information in the CPU report

A high BUSY TIME PERCENTAGE could indicate contention for CPU. To check this, add the IN READY numbers in the columns N+1,...,14+ (where N is the number of CPs in your system). This sum is the percentage of time when at least one task could not be dispatched. If the value is higher than 60%, it implies contention for CPU.

A low BUSY TIME PERCENTAGE can indicate that other bottlenecks in the system are preventing work from being processed.

An OUT READY value of more than 1 could reflect processor storage constraints.

```
                                        C P U   A C T I V I T Y   R E P O R T                    PAGE  1

           z/OS V1R4               SYSTEM ID   SYS1           DATE 03/14/2002         INTERVAL 13.37.614
                                   RPT VERSION V1R2 RMF       TIME 16.30.00           CYCLE 1.000 SECONDS

CPU  2064   MODEL  109

CPU     ONLINE TIME   LPAR BUSY       MVS BUSY      CPU SERIAL  I/O TOTAL          % I/O INTERRUPTS
NUMBER  PERCENTAGE    TIME PERC       TIME PERC     NUMBER      INTERRUPT RATE     HANDLED VIA TPI
3       100.00        64.24           100.0         3A1534      413.0               0.64
4       100.00        64.21           100.0         4A1534      410.9               0.66
5       100.00        64.19           100.0         5A1534      410.6               0.60
6       100.00        64.22           100.0         6A1534      415.6               0.66
TOTAL/AVERAGE         64.22           100.0                      1650               0.64

SYSTEM ADDRESS SPACE ANALYSIS         SAMPLES =    817

           NUMBER OF ASIDS                                DISTRIBUTION OF QUEUE LENGTHS  (%)
TYPE       MIN    MAX     AVG      0      1      2      3      4      5      6     7-8   9-10   11-12  13-14   14+
----     ------ ------ --------   -----  -----  -----  -----  -----  -----  -----  -----  -----  -----  -----  -----
IN
READY         8     23     13.6     0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.1    1.9   20.4   55.5   21.9

                                    0     1-2    3-4    5-6    7-8   9-10   11-15  16-20  21-25  26-30  31-35   35+
                                   -----  -----  -----  -----  -----  -----  -----  -----  -----  -----  -----  -----
IN            61     65     62.0     0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0  100.0
OUT
READY          0      0      0.0   100.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0
OUT
WAIT           0      0      0.0   100.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0
LOGICAL
OUT RDY        0      0      0.0   100.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0
LOGICAL
OUT WAIT      42     45     44.0     0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0  100.0

BATCH         15     15     15.0     0.0    0.0    0.0    0.0    0.0    0.0  100.0    0.0    0.0    0.0    0.0    0.0
STC           89     91     89.0     0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0  100.0
TSO            2      3      2.0     0.0   97.4    2.5    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0
ASCH           0      0      0.0   100.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0
OMVS           0      0      0.0   100.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0    0.0
```

Figure 5-16. CPU Activity Report

Table 5-18. Fields in the CPU Activity Report

Field Heading	Meaning
CPU Activity	
CPU/MODEL	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 version number (in hexadecimal) will be shown.
CPU NUMBER	The internal address of the processor for which values are being reported (logical CPU number).
ONLINE TIME PERCENTAGE	The percentage of time, the processor was online.

Table 5-18. Fields in the CPU Activity Report (continued)

Field Heading	Meaning

LPAR BUSY TIME PERC	• For a dedicated partition:
	`LPAR BUSY TIME % = (Online Time - Wait time) / Online Time * 100`
	• For a non-dedicated partition when Wait Completion is NO:
	`LPAR BUSY TIME % = Partition Dispatch Time / Online Time * 100`
	The partition dispatch time is the elapsed time that PR/SM dispatched this processor during the interval.
	• For a non-dedicated partition when Wait Completion is YES:
	`LPAR BUSY TIME % = (Partition Dispatch Time - Wait Time) / Online Time * 100`

MVS BUSY TIME PERC	The percentage of the online time that the processor was busy. The MVS busy time percentage calculation is: `MVS BUSY TIME % = (Online Time - Wait Time) / Online Time * 100`
CPU SERIAL NUMBER	The contents depends on the mode: • Basic mode The six-digit number that matches the physical serial number stamped on the processor that, along with the model number, uniquely identifies the processor. • LPAR mode For processors up to z900, this field has the format LPnnnn with **L** Logical CPU address **P** Logical partition identifier **nnnn** Digit derived from the serial number For z9xx processors or later, this field has the format PPnnnn, where PP indicates the logical partition identifier. The CPU address is no longer part of the serial number.
I/O TOTAL INTERRUPT RATE	The total rate per second that this processor handled I/O interruptions. The rate reported reflects the processing for the entire interval. This might include periods of time when the SRM enabled or disabled this processor for I/O interruptions. The rate includes interruptions that the second level interrupt handler (SLIH) handled as well as those the Test Pending Interrupt (TPI) instruction handled. `RATE = (SLIH + TPI) / INT` **SLIH** Interruptions that the second level interrupt handler handled **TPI** Interruptions that the Test Pending Interrupt instruction handled **INT** Interval time (seconds)

Table 5-18. Fields in the CPU Activity Report (continued)

Field Heading	Meaning

% I/O INTERRUPTS HANDLED VIA TPI	The percentage of the total interrupts for this processor during the RMF interval that are handled by the I/O supervisor without re-enabling.
	`% TPI INTERRUPTS = ---------- * 100`
	`TPI / (SLIH + TPI)`
	TPI Interruptions that the Test Pending Interrupt instruction handled
	SLIH Interruptions that the second level interrupt handler handled

TOTAL/AVERAGE	The average or total value for all processors.
System Address Space Analysis	
TYPE	The type of measurement data. The following types of address spaces are analyzed: **IN READY** Address spaces that are in central storage and ready to execute. **IN** Address spaces that are in central storage (corresponds to SRM in queue). This count includes the IN READY count. **OUT READY** Address spaces on the SRM out queue that are physically swapped out of central storage and ready to execute. **Note:** 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. **OUT WAIT** Address spaces on the SRM wait queue that are physically swapped out of central storage and not ready to execute. **LOGICAL OUT RDY** Address spaces on the SRM out queue that are physically in central storage but logically swapped out of central storage and ready to execute. **LOGICAL OUT WAIT** 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. **BATCH** Address spaces used for batch jobs. **STC** Address spaces used for started task controls. **TSO** Address spaces used for TSO/E users. **ASCH** APPC/MVS transaction scheduler (ASCH) address spaces. **OMVS** Address spaces for z/OS UNIX System Services.
NUMBER OF ASIDS	The number of address space identifications (ASIDs) of each type detected during the reporting interval, reported as: **MIN** Minimum number of address spaces for each type **MAX** Maximum number of address spaces for each type **AVG** Average number of address spaces for each type
DISTRIBUTION OF QUEUE LENGTHS (%)	The percentage of samples when the number of ASIDs for a type fell within a certain range. The range is indicated in the subheadings.

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 5-19. Overview Names in the CPU Activity Report

Column Heading	Overview Name
ONLINE TIME PERCENTAGE	CONTPER
LPAR BUSY TIME PERC	CPUBSY (LPAR mode only)

Table 5-19. Overview Names in the CPU Activity Report (continued)

Column Heading	Overview Name
MVS BUSY TIME PERC	MVSBSY (LPAR mode) CPUBSY (BASE mode)
TYPE (IN READY)	AVGIARDY
TYPE (other)	MXBATCH, AVGBATCH, MXSTC, AVGSTC, MXTSO, AVGTSO, MXASCH, AVGASCH, MXOMVS, AVGOMVS, AVGOARDY, AVGUIN, AVGUOWT, AVGULRDY, AVGULWT
Number of processors online	NUMPROC
Percentage of the interval during which at least one job could not be dispatched	OCPU1
Percentage of the interval during which at least two jobs could not be dispatched	OCPU2
Percentage of the interval during which at least three jobs could not be dispatched	OCPU3

Using the information in the Partition Data report

The Partition Data Report section of the CPU Activity report provides data about all configured partitions, independent of the type of operating system running in each partition. This section only appears when RMF is running in a Processor Resource/Systems Manager environment in LPAR mode, and is only for partitions active at the end of the duration interval. The section is not available if RMF is running in a VM guest environment.

RMF provides information about the partitions that are sharing the processor. If you want further information about a particular partition, you can run RMF separately on that partition.

The report is divided into four parts:
- Header information
- Partition data (first block)
- Logical partition processor data (second block)
- Average processor utilization percentages (third block)

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 —split up into general purpose processors (CP) and integrated coupling facilities (ICF)
- Wait completion indicator
- Dispatch interval

The partitions are grouped according to the type of their physical processors (CP or ICF), and all percentages which are based on the number of physical processors are calculated using the corresponding number of processors.

The section `PARTITION DATA` provides the following information for each partition:

- Name
- Status
- Weighting share of resources
- Defined and consumed service units
- Capping information

If shared and dedicated processor have been assigned to a partition (exploiting the dynamic ICF expansion function), two lines will be reported for this partition.

The section LOGICAL PARTITION PROCESSOR DATA provides the following information about the partition's processors:
- Number and type of logical 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
  - 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
  - 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

```
                                    P A R T I T I O N   D A T A   R E P O R T
                                                                                                                    PAGE    2
           z/OS V1R4               SYSTEM ID   NP1           DATE 04/15/2002         INTERVAL 14.59.678
                                   RPT VERSION V1R2 RMF      TIME 09.30.00           CYCLE 1.000 SECONDS

MVS PARTITION NAME                    NP1
IMAGE CAPACITY                        100
NUMBER OF CONFIGURED PARTITIONS         9
NUMBER OF PHYSICAL PROCESSORS           9
                 CP                     9
                 ICF                    0
WAIT COMPLETION                        NO
DISPATCH INTERVAL                 DYNAMIC

--------- PARTITION DATA -----------------  -- LOGICAL PARTITION PROCESSOR DATA --   -- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --

                  ----MSU----  -CAPPING--  PROCESSOR-  ----DISPATCH TIME DATA----   LOGICAL PROCESSORS  --- PHYSICAL PROCESSORS ---
NAME       S   WGT  DEF    ACT  DEF   WLM%  NUM   TYPE   EFFECTIVE       TOTAL      EFFECTIVE    TOTAL  LPAR MGMT  EFFECTIVE  TOTAL

NP1        A    20  100     10  NO    62.2  1.2   CP    00.04.29.502  00.04.27.519       25.10    24.92      ****       3.33   3.30
NP2        A     1    0      1  YES    0.0    4   CP    00.00.22.680  00.00.22.083        0.75     0.73      ****       0.28   0.27
NP3        A    10    5      8  NO     3.3  1.0   CP    00.03.37.761  00.03.35.859       24.20    23.99      ****       2.69   2.67
NP4        A   300   95    155  NO     0.0  0.3   CP    01.12.08.405  01.12.06.405       80.18    80.15      ****      53.46  53.43
NP5        A   200   50     52  NO     0.0    4   CP    00.24.13.447  00.24.11.311       40.39    40.33      ****      17.95  17.92
CFC1       A   DED    0     32         0.0    1   CP    00.14.59.611  00.14.59.625       99.99    99.99      0.00      11.11  11.11
CFC2       A   DED    0      0         0.0    1   CP    00.00.00.000  00.00.00.000        0.00     0.00      0.00       0.00   0.00
*PHYSICAL*                                             00.00.03.603                                          0.04               0.04
                                                       ------------  ------------                          ------     ------ ------
TOTAL                                                  01.59.51.408  01.59.46.408                          0.04      88.81  88.75

CB88       D
CB89       D
```

Figure 5-17. CPU Activity Report - Partition Data Report Section

Note: An asterisk (*) next to any value indicates a change to this value during the measurement interval.

PP - CPU Report

Table 5-20. Fields in the Partition Data Report

Field Heading	Meaning
Header Information	
MVS PARTITION NAME	The partition in the MVS system which requested this report.
IMAGE CAPACITY	CPU capacity available to the MVS image measured in MSUs (millions of service units) per hour. The field is calculated as follows: • If the partition has a defined capacity limit, this value is that limit. • If the partition has no defined capacity limit, this value is calculated based on the partition's logical CP configuration and capping state.
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*.
NUMBER OF PHYSICAL PROCESSORS	The number of physical processors in PR/SM system. In addition, the number of general purpose processors (CP) and the number of integrated coupling facilities (ICF) is shown. The value for ICF indicates also processors that are running Linux partitions (IFL). Due to the fact that ICFs and IFLs are a single pool of resources, everything is reported as ICF.
WAIT COMPLETION	The wait completion option of the partition: YES, NO, or MIX **YES** Implies that the processors assigned to each partition will remain dispatched to the partition until the time slice period has ended. **NO** 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. **MIX** Indicates that a mix of YES and NO is used for processors in the partition where RMF is running.
DISPATCH INTERVAL	The time (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 *PR/SM Planning Guide* for more information.
Partition Data	
NAME	The name that identifies a partition. **Note:** The partition identified by the name PHYSICAL is not a configured partition (an '*' before and after the name PHYSICAL accentuates this point). Data reported for *PHYSICAL* is shown only in columns PARTITION TOTAL DISPATCH TIME, AVG PHYS PROC LPAR MGMT and AVG PHYS PROC TOTAL, and it includes all of the uncaptured time which was used by LPAR but could not be attributed to a specific logical partition.
S	The current status of the partition: **A** Activated **D** Deactivated If a partition is deactivated, the rest of the report line is blank.

Table 5-20. Fields in the Partition Data Report (continued)

Field Heading	Meaning
WGT	Either the partition's current weighting of the shared processor resources or one of the following indicators: **DED** Indicates that the partition is dedicated. **DMX** Indicates that a mix of dedicated and non-dedicated processors is used in this partition. **WMX** Indicates that different share values are assigned to logical processors used in this partition.
MSU DEF	Defined capacity limit of the partition in MSUs per hour. 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.
MSU ACT	Actual consumption.
CAPPING DEF	The capping option of the partition: YES or NO This field indicates whether the operator has set 'Initial Capping ON' in the logical partition controls of the Hardware Management Console (HMC) for the partition. The field is only useful to logical partitions with shared processors. Dedicated partitions can be thought of as automatically capped.
CAPPING WLM%	Percentage of time when WLM capped the partition.
Logical Partition Processor Data	
PROCESSOR NUM TYPE	The number and type of logical processors assigned to this partition.
EFFECTIVE DISPATCH TIME	The sum of all logical 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 all of the uncaptured time during the measurement interval which was used by LPAR, but could not be attributed to a specific logical partition.
Average Processor Utilization Percentages 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.	
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 5-20. Fields in the Partition Data Report (continued)

Field Heading	Meaning

PHYSICAL PROCESSORS - LPAR MANAGEMENT	The average LPAR management time on behalf of the partition reported as a percentage of the measurement interval.
	`Total Dispatch Time - Effective Dispatch Time`
	`------------------------------------------ * 100`
	`  # Physical Processors * Interval Time`
	
	If the total dispatch time is smaller than the effective dispatch time (see `TOTAL DISPATCH TIME`), `****` is shown in this column.
	
	The calculation for the *PHYSICAL* partition is:
	`       Total Time *PHYSICAL*`
	`---------------------------------- * 100`
	`# Physical Processors * Interval Time`

PHYSICAL PROCESSORS - EFFECTIVE	The effective utilization of the physical processor resource by the partition.
	`       Effective Dispatch Time`
	`---------------------------------- * 100`
	`# Physical Processors * Interval Time`

PHYSICAL PROCESSORS - TOTAL	The total utilization of the physical processor resource by the partition.
	`        Total Dispatch Time`
	`---------------------------------- * 100`
	`# Physical Processors * Interval Time`

TOTAL	The total amount of time the physical processor resource was assigned to a configured partition and to partition *PHYSICAL*. The sum of the AVERAGE PHYSICAL PROCESSOR UTILIZATION - LPAR MGMT field represents the total utilization of PR/SM physical processor resource by PR/SM. 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. 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*.

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 5-21. Overview Condition Names Applicable to the Partition Data Report

Column Heading	Overview Name
PARTITION DATA - MSU DEF	LDEFMSU
PARTITION DATA - MSU ACT	LACTMSU
PARTITION DATA - CAPPING WLM%	WCAPPER

Using the information in the LPAR Cluster report

The Workload Manager is extended to work with PR/SM on zSeries 900 (z900) servers 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.

```
                                  L P A R  C L U S T E R  R E P O R T
                                                                                                                PAGE    3
          z/OS V1R4               SYSTEM ID   NP1           DATE 04/15/2002         INTERVAL 14.59.678
                                  RPT VERSION V1R2 RMF      TIME 09.30.00           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

SVPLEX1   NP1        NP1         20   20   40   20    100    0.0         4     1.2   24.92   3.30         1024            N/A
          NP2        NP2          1              1                      4       4    0.73   0.27         1024            N/A
          NP3        NP3         10   10   50   10    100    0.0         4     1.0   23.99   2.67         1024            N/A
-------------------------   ------------------------------------   ------------------------------   ------------------------------
                      TOTAL     31                                      12           49.64   6.24         3072            N/A

SVPLEX2   NP4        NP4        300  100  500  300    0.0    0.0         6     0.3   80.15  53.43         1024            N/A
          NP5        NP5        200            200                      4       4   40.33  17.92         1024            N/A
-------------------------   ------------------------------------   ------------------------------   ------------------------------
                      TOTAL     500                                      10           120.5  71.36         2048            N/A
```

Figure 5-18. CPU Activity Report - LPAR Cluster Report Section

Table 5-22. 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 name are grouped together.
PARTITION	Name of the logical partition.
SYSTEM	z/OS system name.
Weighting Statistics 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 minimum weighting, or 10% below the maximum weighting.
Processor Statistics	

Table 5-22. Fields in the LPAR Cluster Report (continued)

Field Heading	Meaning
NUMBER DEFINED / ACTUAL	Defined and average actual number of logical 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:
	`Partition Total Dispatch Time`
	`---------- * 100`
	`∑ Logical Processor Online Times`
	
	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:
	`Partition Total Dispatch Time`
	`---------- * 100`
	`# Physical Processors * Interval Time`
	
	The contents of this field is equal to the `PHYSICAL PROCESSOR UTILIZATION - TOTAL` column in the Partition Data report.

Storage Statistics	
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 5-23. Overview Condition Names Applicable to the LPAR Cluster Report

Column Heading	Overview Name
WEIGHTING - DEFINED INIT (Cluster)	WDEFC
WEIGHTING - DEFINED INIT (Partition)	WDEFL
WEIGHTING - DEFINED MIN	WMINL
WEIGHTING - DEFINED MAX	WMAXL
WEIGHTING - ACTUAL AVG	WACTL
WEIGHTING - ACTUAL MIN%	WMIPL
WEIGHTING - ACTUAL MAX%	WMAPL
PROCESSOR - NUMBER DEFINDED (Cluster)	NLDEFC
PROCESSOR - NUMBER DEFINDED (Partition)	NLDEFL
PROCESSOR - NUMBER ACTUAL	NLACTL
PROCESSOR - TOTAL% LBUSY (Cluster)	LBUSYC
PROCESSOR - TOTAL% LBUSY (Partition)	LBUSYL
PROCESSOR - TOTAL% PBUSY (Cluster)	PBUSYC
PROCESSOR - TOTAL% PBUSY (Partition)	PBUSYL

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

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 print this report, specify
`REPORTS(CRYPTO)`

Purpose of the report

The Crypto Hardware Activity report provides information about the activities in the various cryptographic hardware features.

The Internet is rapidly becoming the basis for electronic commerce. More businesses are automating their data processing operations. Online databases are becoming increasingly large and complex. Many businesses transmit sensitive data on open communication networks and store confidential data offline. Every day the potential for unauthorized persons to access sensitive data increases.

Only cryptographic services can provide the data confidentiality and the identity authentication that is required to protect business commerce on the Internet.

The Integrated Cryptographic Service Facility (ICSF) is a software element of z/OS that works with the cryptographic hardware and the z/OS Security Server to provide secure, high-speed cryptographic services in the z/OS environment. ICSF provides the application programming interfaces by which applications request the cryptographic services. The cryptographic feature is secure, high-speed hardware that performs the actual cryptographic functions. The cryptographic feature available to your applications depends on the server or processor hardware. The Crypto Hardware Activity report supports the following hardware:

- **PCI Cryptographic Coprocessor (PCICC)**

  The PCICC is a data encryption coprocessor card designed to provide data security functions which is attached via an internal PCI (Peripheral Component Interconnect) bus to the system. The card operation is controlled by an on-board 486-type processor. It has special hardware for various cryptographic operations (Rivest-Shamir-Adleman (RSA) algorithm, DES, or Secure Hash Algorithm 1 (SHA-1)). Other cryptographic algorithms are carried out by using the on-board processor. The card works in parallel and asynchronously to the zSeries processors.

  The PCICC card is aimed at a high security environment where the keys are encrypted when outside the protected card. It is mainly used for RSA public key algorithm operations including RSA key generation, special PIN operations, and special user cryptographic functions (using the User Defined Extension (UDX) capability of the card).

- **PCI Cryptographic Accelerator (PCICA)**

  The PCICA is a cryptographic hardware accelerator card which is attached via the PCI bus to the system. This card provides a competitive option to customers that do not need the high security environment of the PCICC but need high cryptographic performance for RSA public key algorithms which may be required in Web server applications. The PCICA provides clear key RSA operations with a modulo length of 1024 bit or 2048 bits. Two algorithms are

available, modular exponentiation (ME) and Chinese Remainder Theorem (CRT) for both key lengths. The card has five hardware engines that can work in parallel and asynchronously to the zSeries processors. There are always two cards packaged in one adapter that plugs into the system board. This allows for a scalable and cost effective system performance in Web server applications.

- **PCIX Cryptographic Coprocessor (PCIXCC)**

  The PCIXCC is a cryptographic hardware adapter card which is attached to the system via an internal PCIX (Peripheral Component Interconnect Extended) bus. The PCIXCC adapter is a breakthrough concept which supports high security-demanding applications and also serves as an execution environment for customer written programs (user defined extensions: UDX).

  The PCIXCC is programmable, so you can extend the capability of the CCA application program that performs within the PCIXCC to support new applications.

  This hardware can replace the PCI cryptographic coprocessor (PCICC) card and functions previosly executed on the CMOS Cryptographic Coprocessor feature (CCF).

- **ICSF Services**

  The Crypto Hardware Activity report provides performance measurements on selected ICSF activities on either the CCF hardware or on the PCIXCC hardware (for z990 processors):

  – Using the Data Encryption Standard (DES), which is probably the best known and scrutinized encryption algorithm, 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 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.

Contents of the report

For cryptographic 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. The data for cryptographic accelerators is showing details for the two available algorithms, providing by this some information how the usage of these algorithms affects the utilization of the accelerator.

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.

PP - CRYPTO Report

```
                                   C R Y P T O   H A R D W A R E   A C T I V I T Y
                                                                                                                PAGE    6
            z/OS V1R4               SYSTEM ID   SYS1           DATE 03/14/2002          INTERVAL 14.59.946
                                    RPT VERSION V1R2 RMF       TIME 16.30.00            CYCLE 1.000 SECONDS

-------- CRYPTOGRAPHIC COPROCESSOR ------
          ------ TOTAL --------  KEY-GEN
TYPE   ID   RATE  EXEC TIME UTIL%    RATE
PCIXCC  0   0.00        0.0   0.0    0.00
        1   0.01       3205  32.1    0.01
        6  83.04        1.1   8.8       0
        7   0.00        0.0   0.0    0.00

------------------------------------------------------ CRYPTOGRAPHIC ACCELERATOR ------------------------------------------------------
          -------- TOTAL --------   ------ ME(1024) -------   ------ ME(2048) ------   ------ CRT(1024) ------  ------ CRT(2048) ------
TYPE   ID   RATE  EXEC TIME  UTIL%    RATE  EXEC TIME  UTIL%   RATE  EXEC TIME  UTIL%   RATE  EXEC TIME  UTIL%    RATE  EXEC TIME  UTIL%
PCICA   2  165.2        1.3   21.5   107.1        1.1   11.8      0          0      0   58.1        1.7    9.7       0          0      0


----------------------------- ICSF SERVICES EXECUTED ON PCIXCC --------------------------------
        DES ENCRYPTION    DES DECRYPTION    ----- MAC ------    - HASH -    ------ PIN -------
        SINGLE  TRIPLE    SINGLE  TRIPLE    GENERATE  VERIFY               TRANSLATE  VERIFIY
RATE      4975K  497.5     12438    1244K      12438    4975K     497.5          1244K    1244
SIZE      0.75     100K    10.00    0.01       10.00    0.01      10000
```

Figure 5-19. Crypto Hardware Activity Report

Table 5-24. Fields in the CRYPTO Hardware Activity Report

Field Heading	Meaning
Cryptographic Coprocessor	
TYPE	Type that defines the cryptographic coprocessor.
ID	Index that specifies the cryptographic coprocessor.
TOTAL	**RATE** Rate of all operations on this cryptographic coprocessor. **EXEC TIME** Average execution time (milliseconds) of all operations on this cryptographic coprocessor. **UTIL%** Total utilization percentage of this coprocessor.
KEY-GEN RATE	Rate for RSA-key-generation operations.
Cryptographic Accelerator	
TYPE	Type that defines the cryptographic accelerator.
ID	Index that specifies the cryptographic accelerator.
TOTAL	Rate, execution time and utilization for all operations on this cryptographic accelerator.
ME(1024)	Rate, execution time and utilization for all operations in 1024-bit-ME format.
ME(2048)	Rate, execution time and utilization for all operations in 2048-bit-ME format.
CRT(1024)	Rate, execution time and utilization for all operations in 1024-bit-CRT format.
CRT(2048)	Rate, execution time and utilization for all operations in 2048-bit-CRT format.
ICSF services executed on PCIXCC or **ICSF services executed on CCF**	

Table 5-24. Fields in the CRYPTO Hardware Activity Report (continued)

Field Heading	Meaning
DES ENCRYPTION	**SINGLE RATE** Rate of requests (calls) to encipher for single-DES. **SINGLE SIZE** Average number of bytes per request that have been enciphered with single-DES. **TRIPLE RATE** Rate of requests (calls) to encipher for triple-DES. **TRIPLE SIZE** Average number of bytes ´per request that have been enciphered with triple-DES.
DES DECRYPTION	**SINGLE RATE** Rate of requests to decipher for single-DES. **SINGLE SIZE** Average number of bytes per request that have been deciphered with single-DES. **TRIPLE RATE** Rate of requests to decipher for triple-DES. **TRIPLE SIZE** Average number of bytes per request that have been deciphered with triple-DES.
MAC	**GENERATE RATE** Rate of requests to generate MACs. **GENERATE SIZE** Average number of bytes per request for which MAC has been generated. **VERIFY RATE** Rate of requests to verify MACs. **VERIFY SIZE** Average number of bytes per request for which MAC has been verified..
HASH	**RATE** Rate of requests to hash. **SIZE** Average number of bytes to be hashed per request.
PIN	**TRANSLATE RATE** Rate of requests to translate PIN. **VERIFY RATE** Rate of requests to verify PIN.

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 5-25. Overview Names in the CRYPTO Hardware Activity Report

Column Heading	Overview Name
Cryptographic coprocessor TOTAL RATE	CRYCTR
Cryptographic coprocessor TOTAL EXEC TIME	CRYCTE

Table 5-25. Overview Names in the CRYPTO Hardware Activity Report (continued)

Column Heading	Overview Name
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 CRT1024) 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
DES ENCRYPTION SINGLE RATE	CRYISDER
DES ENCRYPTION SINGLE SIZE	CRYISDES
DES ENCRYPTION TRIPLE RATE	CRYITDER
DES ENCRYPTION TRIPLE SIZE	CRYITDES
DES DECRYPTION SINGLE RATE	CRYISDDR
DES DECRYPTION SINGLE SIZE	CRYISDDS
DES DECRYPTION TRIPLE RATE	CRYITDDR
DES DECRYPTION TRIPLE SIZE	CRYITDDS
MAC GENERATE RATE	CRYIMGR
MAC GENERATE SIZE	CRYIMGS
MAC VERIFY RATE	CRYIMVR
MAC VERIFY SIZE	CRYIMVS
HASH RATE	CRYIHAR
HASH SIZE	CRYIHAS
PIN TRANSLATE RATE	CRYIPTR
PIN VERIFY RATE	CRYIPVR

DEVICE - Device Activity report

How to request this report

Monitor I gathers data for this report automatically with the default option `DEVICE(DASD)`. See the *z/OS RMF User's Guide* for details. If you want to suppress gathering, you need to specify `NODEVICE`.

To print this report, specify
`REPORTS(DEVICE(type))`

Purpose of the 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.

Using the information given in the report

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.

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 sub-options 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 5-20 on page 5-63 shows a DASD Activity report.

Device data incomplete or missing

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 CUB DELAY
AVG SPB 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, plot reports and summary reports, overflows are neither detected nor identified. Because Monitor II device activity reports, as well as some exception reports, plot 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 MVS, 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/Swap data set 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 4 seconds will ensure that no pending time is lost. However, some performance penalty does occur because of the 4-second interval.

```
                              D I R E C T   A C C E S S   D E V I C E   A C T I V I T Y

                                                                                                                                                PAGE    6
           z/OS V1R4               SYSTEM ID   SYS1           DATE 04/25/2002           INTERVAL 14.59.946
                                   RPT VERSION V1R2 RMF       TIME 16.30.00             CYCLE 1.000 SECONDS

 TOTAL SAMPLES =    900   IODF = A3   CR-DATE: 04/05/2002   CR-TIME: 07.42.20         ACT: POR

                                          DEVICE   AVG  AVG   AVG  AVG   AVG  AVG  AVG    %      %     %     AVG    %      %
 STORAGE  DEV  DEVICE   VOLUME PAV  LCU  ACTIVITY RESP IOSQ  CMR  DB    PEND DISC CONN   DEV    DEV   DEV   NUMBER ANY    MT
  GROUP   NUM  TYPE     SERIAL          RATE     TIME TIME  DLY  DLY   TIME TIME TIME   CONN   UTIL  RESV  ALLOC  ALLOC  PEND
          0401 3380K    SYSLIB      0032   1.246   4.6  0.0   0.0  2.2   2.5  0.1  2.0   0.25   0.26   0.0   89.6   100.0   0.0
          0402 3380K    SYSUSR      0032   0.250   1.4  0.0   0.0  0.1   0.3  0.0  1.0   0.03   0.03   0.1   14.0   100.0   0.0
          040F 3380K    SCL338      0032   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0    0.0   100.0   0.0
                         LCU        0032   1.496   4.1  0.0   0.0  1.8   2.1  0.1  1.8   0.07   0.07   0.0    104   100.0   0.0
          044F 3380K    MVSPG1      0033   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0    0.0   100.0   0.0
          0460 3380K    RMFLIB      0033   0.036   6.1  0.0   0.0  0.3   1.5  0.0  4.6   0.02   0.02   0.0    0.0   100.0   0.0
          047F 3380K    MVSPLX      0033   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0    0.0   100.0   0.0
                         LCU        0033   0.036   6.1  0.0   0.0  0.3   1.5  0.0  4.6   0.00   0.00   0.0    4.0   100.0   0.0
          0500 33903    MVSLIB      0034   0.082  22.6  0.0   0.0 13.0  16.2  0.3  6.1   0.05   0.05   0.0   20.4   100.0   0.0
          0501 33903    MVSSCF      0034   0.012   1.6  0.0   0.0  0.0   0.5  0.0  1.0   0.00   0.00   0.0    4.6   100.0   0.0
          0502 33903    MVSCI2      0034   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0    0.0   100.0   0.0
 OMVSSYS  0503 33903    MVSOP2      0034   0.008  19.2  0.0   0.0  0.0   0.8  0.0 18.4   0.01   0.01   0.0    0.0   100.0   0.0
          0705 33909    15CY09      0035   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0    0.0   100.0   0.0
          0707 33909    16RJ02      0035   0.036  55.2  0.0   0.0 36.7  44.6  0.0 10.6   0.04   0.04   0.0    0.0   100.0   0.0
          0708 33909    15CYX9      0035   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0    0.0   100.0   0.0
```

Figure 5-20. Direct Access Device Activity Report

The reports for communication equipment, character reader devices, graphic devices, and unit record devices have the same format. Therefore, the Communication Equipment Activity report is shown as example.

```
                            C O M M U N I C A T I O N   E Q U I P M E N T   A C T I V I T Y
                                                                                                                                    PAGE  1
           z/OS V1R4               SYSTEM ID   SYS1           DATE 04/25/2002           INTERVAL 14.59.946
                                   RPT VERSION V1R2 RMF       TIME 16.30.00             CYCLE 1.000 SECONDS

 TOTAL SAMPLES =  1,800    IODF = 12  CR-DATE: 04/05/2002   CR-TIME: 10.23.17     ACT: ACTIVATE

                                  DEVICE   AVG  AVG   AVG  AVG   AVG  AVG  AVG    %      %     %     AVG    %      %    %
         DEV  DEVICE   VOLUME  LCU  ACTIVITY RESP IOSQ  CMR  DB    PEND DISC CONN   DEV    DEV   DEV   NUMBER ANY    MT   NOT
         NUM  TYPE     SERIAL       RATE     TIME TIME  DLY  DLY   TIME TIME TIME   CONN   UTIL  RESV  ALLOC  ALLOC  PEND RDY

         1106                 004B   0.058   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0          100.0         0
         1107                 004B   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0          100.0         0
         111F                 004B   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0            0.0         0
                       LCU    004B   0.058   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0           66.7         0
         2100                 005B   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0            0.0         0
         2101                 005B   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0            0.0         0
         211F                 005B   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0            0.0         0
                       LCU    005B   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0            0.0         0
         8DC1                 0067   1.925   0.5  0.0   0.0  0.0   0.1  0.3  0.2   0.03   0.08   0.0          100.0         0
         8EC1                 0067   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0            0.0         0
                       LCU    0067   1.925   0.5  0.0   0.0  0.0   0.1  0.3  0.2   0.02   0.04   0.0           50.0         0
         9DC1                 006A   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0            0.0         0
         9EC1                 006A   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0            0.0         0
                       LCU    006A   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00   0.0            0.0         0
```

Figure 5-21. Communication Equipment Activity Report

The following figure shows the Magnetic Tape Device Activity report.

PP - DEVICE Report

```
                          M A G N E T I C   T A P E   D E V I C E   A C T I V I T Y
                                                                                                                                PAGE  1
          z/OS V1R4               SYSTEM ID J67D            DATE 03/07/2002           INTERVAL 15.00.029
                                  RPT VERSION V1R2 RMF      TIME 16.30.00           CYCLE 1.000 SECONDS

TOTAL SAMPLES =  1,800   IODF = 02   CR-DATE: 03/03/2002   CR-TIME: 20.49.20     ACT: POR

                                DEVICE   AVG  AVG   AVG  AVG   AVG  AVG  AVG    %      %     %     NUMBER   AVG      TIME
        DEV  DEVICE   VOLUME  LCU  ACTIVITY RESP IOSQ  CMR  DB    PEND DISC CONN   DEV    DEV   DEV   OF       MOUNT    DEVICE
        NUM  TYPE     SERIAL       RATE     TIME TIME  DLY  DLY   TIME TIME TIME   CONN   UTIL  RESV  MOUNTS   TIME     ALLOC


        8000 3590L            0017   0.004   1.6  0.0   0.0  0.0   0.3  1.0  0.3   0.00   0.00  0.0       0        0        0
        8001 3590L            0017   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0       0        0        0
        8002 3590L            0017   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0       0        0        0
        8003 3590L            0017   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0       0        0        0
                      LCU     0017   0.004   1.6  0.0   0.0  0.0   0.3  1.0  0.3   0.00   0.00  0.0       0        0        0
        8100 3590L            0018   0.004   1.5  0.0   0.0  0.0   0.3  1.0  0.3   0.00   0.00  0.0       0        0        0
        8101 3590L            0018   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0       0        0        0
        8102 3590L            0018   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0       0        0        0
        8103 3590L            0018   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0       0        0        0
                      LCU     0018   0.004   1.5  0.0   0.0  0.0   0.3  1.0  0.3   0.00   0.00  0.0       0        0        0
        8200 3590L            0019   0.004   1.5  0.0   0.0  0.0   0.3  1.0  0.3   0.00   0.00  0.0       0        0        0
        8201 3590L            0019   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0       0        0        0
        8202 3590L            0019   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0       0        0        0
        8203 3590L            0019   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0       0        0        0
        8204 3590L            0019   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0       0        0        0
        8205 3590L            0019   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0       0        0        0
                      LCU     0019   0.004   1.5  0.0   0.0  0.0   0.3  1.0  0.3   0.00   0.00  0.0       0        0        0
        8300 3590L            001A   0.004   1.9  0.0   0.0  0.0   0.3  1.4  0.3   0.00   0.00  0.0       0        0        0
        8301 3590L            001A   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0       0        0        0
        8302 3590L            001A   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0       0        0        0
        8303 3590L            001A   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0       0        0        0
                      LCU     001A   0.004   1.9  0.0   0.0  0.0   0.3  1.4  0.3   0.00   0.00  0.0       0        0        0
        8400 3590L            001B   0.004   2.0  0.0   0.0  0.0   0.3  1.5  0.2   0.00   0.00  0.0       0        0        0
        8401 3590L            001B   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0       0        0        0
        8402 3590L            001B   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0       0        0        0
        8403 3590L            001B   0.000   0.0  0.0   0.0  0.0   0.0  0.0  0.0   0.00   0.00  0.0       0        0        0
                      LCU     001B   0.004   2.0  0.0   0.0  0.0   0.3  1.5  0.2   0.00   0.00  0.0       0        0        0
```

Figure 5-22. Magnetic Tape Device Activity Report

Table 5-26. 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.
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	The number of parallel access volumes (base and alias) which were available at the end of the reporting interval. If the number has been changed during the reporting interval, it is followed by an '*'.

Table 5-26. Fields in the Device Activity Reports (continued)

Field Heading	Meaning
LCU	The number of the logical control unit (LCU) to which the device belongs. 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. There are two reasons that this field is blank: • RMF encountered an error while gathering data, check the operator console for messages. • This is a non-dedicated device in a VM guest system environment.

DEVICE ACTIVITY RATE	The rate at which start subchannel (SSCH) instructions to the device completed successfully.
	
	`# Successful SSCH Instructions`
	`ACTV RATE = ------------------------------`
	`Interval Time`
	
	For devices using suspended channel programs, resume I/O requests are included in the SSCH counts.
	
	In the LCU summary line, this field contains the sum of the rates for each individual device.
	
	If the device has been deleted during the last interval, DEVICE DYNAMICALLY DELETED appears in the field instead of the measurement data.
	
	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.

AVG RESP TIME	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.
	
	`Device Active Time`
	`AVG ACT TIME = ------------------------`
	`Measurement Event Count`
	
	`AVG RESP TIME = AVG ACT TIME + AVG IOSQ TIME`
	
	The active time is the sum of connect, disconnect, and pending time as described later.
	
	In the LCU summary line, this field contains the weighted average of the individual average response times for each device.

AVG IOSQ TIME	The average number of milliseconds an I/O request must wait on an IOS queue before a SSCH instruction can be issued.
	
	`IOSQ Count / # Samples`
	`AVG IOSQ TIME = ----------------------`
	`Device Activity Rate`

Table 5-26. Fields in the Device Activity Reports (continued)

Field Heading	Meaning

AVG CMR DLY	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).
	`Initial Command Response Time`
	`AVG CMR DLY = ----------------------------------`
	`# I/O Operations Accepted on that Path`

AVG DB DLY	The average number of milliseconds of delay that I/O requests to this device encountered because the device was busy. Device busy might mean:
	• Another system is using the volume
	• Another system reserved the device
	• Head of string busy conditions caused contention
	• Some combination of these three conditions has occurred.
	`Device Busy Delay Time`
	`AVG DB DLY = ----------------------`
	`Measurement Event Count`

AVG PEND TIME	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.
	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.
	`Device Pending Time`
	`PEND TIME = ------------------------`
	`Measurement Event Count`

AVG DISC TIME	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.
	`Device Disconnect Time`
	`AVG DISC TIME = ----------------------`
	`Measurement Event Count`
	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.

AVG CONN TIME	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.
	`Device Connect Time`
	`AVG CONN TIME = -------------------------`
	`Measurement Event Count`

% DEV CONN	The percentage of time during the interval when the device was connected to a channel path.
	`Device Connect Time`
	`% DEV CONN = ------------------ * 100`
	`Interval Time`

Table 5-26. Fields in the Device Activity Reports (continued)

Field Heading	Meaning

% DEV UTIL	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.
	
	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.
	`CON + DISC RSV`
	`% DEV UTIL = (---------- + ---) * 100`
	`INT N`
	CON Device connect time
	DISC Device disconnect time
	RSV Number of samples when the device was reserved but not involved in an I/O operation
	INT Interval time (seconds)
	N Total number of samples

% DEV RESV	The percentage of time during the interval when a shared device was reserved by the processor on which RMF was started.
	
	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.
	`# Device-reserved Samples`
	`% DEV RESV = ------------------------ * 100`
	`# Samples`

AVG NUMBER ALLOC	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. 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.

% ANY ALLOC	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.
	
	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.
	`# Samples when the Device was Allocated`
	`% ANY ALLOC = ------------------------------------ * 100`
	`# Samples`

% MT PEND	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.
	
	At each cycle, RMF updates a counter when it detects a mount pending condition. At the end of the interval, the percentage is computed.
	`Counter for Mount-Pending Condition`
	`% MT PEND = ---------------------------------- * 100`
	`# Samples`

Table 5-26. Fields in the Device Activity Reports (continued)

Field Heading	Meaning

%NOT RDY	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.
	
	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.
	`# Samples when the Device was not Ready`
	`%NOT RDY = ------------------------------------ * 100`
	`# Samples`

NUMBER OF MOUNTS	The number of tape mounts, shown as an integer value, detected by RMF. 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. If a mount-pending condition is detected at the first cycle of the interval, the mount count for the interval will be incremented by one. In the LCU summary line, this field contains the sum of all mount counts. This field is reported only for magnetic tape devices. **Note:** 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.

AVG MOUNT TIME	The average mount time pending for every device, expressed in the form of HH:MM:SS.
	`# Samples`
	`Tape Mount was Pending`
	`---------- * Interval`
	`# Samples`
	`AVG MOUNT TIME = ----------`
	`# Mounts`
	
	If the mount count or the sample count is zero, the result is zero.
	
	This field is reported only for magnetic tape devices.

TIME DEVICE ALLOC	The total time the device was allocated during the interval, expressed in the form of HH:MM:SS.
	`# Samples Tape Device`
	`was Allocated`
	`TIME DEVICE ALLOC = ---------- * Interval`
	`# Samples`
	
	If the sample count is zero, the result is zero.
	
	This field is reported only for magnetic tape devices.

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 5-27. Overview Names in the DASD Activity Report

Column Heading	Overview Name
DEVICE ACTIVITY RATE	DART
AVG RESP TIME	DRTAVG
AVG IOSQ TIME	DQTAVG
AVG CMR DLY	CMRDL
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

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 print this report, specify

```
REPORTS(DOMINO)
```

Purpose of the report

The Domino Server family is an integrated messaging and Web application software platform. The Domino Server enables Web clients to communicate with Notes 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.

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)

```
                                L O T U S   D O M I N O   S E R V E R   S U M M A R Y
                                                                                                                PAGE  1
           z/OS V1R4              SYSTEM ID   LN21           DATE 03/14/2002          INTERVAL 05.00.000
                                  RPT VERSION V1R2 RMF       TIME 18.40.00

SERVER NAME                    AVAILABLE  ---- USERS -----  TASKS  TRANSACTION  ASYNC  I/O RATE      MAIL RATE       SMTP RATE
                               HHH.MM.SS  CONNECTED ACTIVE             RATE     READS    WRITES  DELIVERED  SENT  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 5-23. Lotus Domino Server Report - Summary

PP - DOMINO Report

```
                                L O T U S   D O M I N O   S E R V E R   D E T A I L S
                                                                                                                                PAGE  2
           z/OS V1R4                SYSTEM ID   LN21           DATE 03/14/2002          INTERVAL 05.00.000
                                    RPT VERSION V1R2 RMF       TIME 18.40.00

NAME: SUT1/COCPOK

--- USER ACTIVITY ---   ------ TASKS ------   ------------- MESSAGES -------------   --- ACCESS RATES ---   --- DATABASE CACHE ---
MAX                0   MAX            2148   MAILBOXES                         1   AS I/O READ    139.3   STATUS              OK
CONNECTED       2034   CURRENT        2148                  COUNT  RATE AVG SIZE   AS I/O WRITE   99.41   MAX ENTRIES        516
ACTIVE            18   MAX UPDATES       0   MAIL DELIVERED   1274  4.25        3   POP3 READ       0.00   CURRENT ENTRIES    216
 WITHIN  1 MIN   361   MAX REPLICS       0   MAIL SENT         211  0.70        1   IMAP READ       0.00   HIGH WATER MARK    774
 WITHIN  3 MIN  1098   COUNT REPLICS     0   SMTP RECEIVED       0  0.00        0   HTTP READ       0.00   INITIAL DB OPENS   408
 WITHIN  5 MIN  1287                         SMTP SENT           0  0.00        0   HTTP WRITE      0.00   REJECTIONS           0
 WITHIN 15 MIN  1814                                                                                       HITS               379
 WITHIN 30 MIN  2001
                      - VIRTUAL THREADS -   - PHYSICAL THREADS -                   --- AVAILABILITY ---   -- NSF BUFFER POOL ---
                      MAX            2034   MAX              76                   THRESHOLD          0   MAX              44032
                      CURRENT        2034   CURRENT          27                   INDEX             98   CURRENT              0
                                            TOTAL           100

------------------------------------------------------------------------------------------------------------------------------------
                                                TRANSACTION ACTIVITY
------------------------------------------------------------------------------------------------------------------------------------

MAXIMUM NUMBER OF CONCURRENT TRANSACTIONS:  NO LIMIT

----------------------- TOP-10 BY COUNT ------------------------    ----------------------- TOP-10 BY R/T TOTAL -------------------
TYPE NAME                    COUNT %TOTAL  RATE    R/T      R/T    TYPE NAME                    COUNT %TOTAL  RATE    R/T      R/T
                                                    AVG    TOTAL                                                          AVG    TOTAL

     TOTAL                   19560  100.0 65.20  22193  434090K         TOTAL                   19560  100.0 65.20  22193  434090K

   6 OPEN_NOTE_RQST           3670  18.76 12.23  23617 86674928      59 NIFOPENNOTE_RQST         2175  11.12  7.25  79485  172880K
  59 NIFOPENNOTE_RQST         2175  11.12  7.25  79485  172880K      8 UPDATE_NOTE_RQST_ALT     1266   6.47  4.22  68560 86796448
   8 UPDATE_NOTE_RQST_ALT     1266   6.47  4.22  68560 86796448      6 OPEN_NOTE_RQST           3670  18.76 12.23  23617 86674928
  55 READ_ENTRIES_RQST        1225   6.26  4.08  275.2   337168    142 START_SERVER_RQST         580   2.97  1.93   149K 86470752
  49 CLOSE_COLLECTION_RQST    1185   6.06  3.95  11.75    13926     51 OPEN_COLLECTION_RQST     1151   5.88  3.84  413.5   475961
  48 CLOSE_DB_RQST_ALT        1183   6.05  3.94  18.44    21815     55 READ_ENTRIES_RQST        1225   6.26  4.08  275.2   337168
   1 OPEN_DB_RQST             1163   5.95  3.88  26.06    30306     15 SEARCHSTART_RQST          135   0.69  0.45   2174   293431
  51 OPEN_COLLECTION_RQST     1151   5.88  3.84  413.5   475961    128 UPDATE_FOLDER_RQST        623   3.19  2.08  76.58    47707
  34 READ_OBJECT_RQST         1030   5.27  3.43   3.53     3635      1 OPEN_DB_RQST             1163   5.95  3.88  26.06    30306
 111 SERVER_TIME_LITE_RQST     860   4.40  2.87   1.20     1030     48 CLOSE_DB_RQST_ALT        1183   6.05  3.94  18.44    21815

------------------------------------------------------------------------------------------------------------------------------------
                                                PORT ACTIVITY
------------------------------------------------------------------------------------------------------------------------------------
MAXIMUM NUMBER OF CONCURRENT SESSIONS:   65535
SESSION TIMEOUT:                         63744

NAME     -- SESSIONS IN --  -- SESSIONS OUT --   SENT  RECEIVED
            COUNT  RATE        COUNT  RATE       (KB)      (KB)

TCPIP         580  1.93          266  0.89     117815      2620
```

Figure 5-24. Lotus Domino Server Report - Details

Table 5-28. 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.

Table 5-28. Fields in the Domino Server Summary Report (continued)

Field Heading	Meaning
SMTP RATE - READS	Rate of SMTP messages received from other servers.
SMTP RATE - WRITES	Rate of SMTP messages sent to other servers.

Table 5-29. Fields in the Domino Server Details Report

Field Heading	Meaning
NAME	Server name
User Activity	
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 **1**, **3**, **5**, **15**, and **30** minutes.
Tasks	
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.
Messages	
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.
Access Rates	
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.
Database Cache	
STATUS	Status of the database cache: either **OK** 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.
Virtual Threads	

Table 5-29. Fields in the Domino Server Details Report (continued)

Field Heading	Meaning
MAX	Maximum number of virtual thread pool threads.
CURRENT	Number of virtual thread pool threads currently in use.
Physical Threads	
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.
Availability	
THRESHOLD	Server availability threshold.
INDEX	Server availability index.
NSF Buffer Pool	
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.
Transaction Activity	
MAXIMAL CONCURRENT	Limit for number of concurrent transactions on a server.
Top-10 List of Transaction Types — 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.
Port Activity	
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

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 a report, specify

```
REPORTS(ENQ)
```

Note: The ENQ report is only available as an interval report, not as a duration report.

Purpose of the 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.

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.

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.

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.

```
                                   E N Q U E U E   A C T I V I T Y                                                PAGE 1
           z/OS V1R4               SYSTEM ID   SYS1           DATE 03/14/2002          INTERVAL 14.59.946
                                   RPT VERSION V1R2 RMF       TIME 16.30.00            CYCLE 1.000 SECONDS

ENQUEUE DETAIL ACTIVITY            GRS MODE: RING
-NAME-  ----- CONTENTION TIME -----  -- JOBS AT MAXIMUM CONTENTION--  -%QLEN DISTRIBUTION- AVG Q -REQUEST TYPE - TOTAL
MAJOR    MIN    MAX    TOT    AVG    ----- OWN ----- ----- WAIT ----    1    2    3    4+  LNGTH -EXCL-  -SHARE- EVENT
 MINOR                              TOT     NAME    TOT     NAME                                  MIN MAX MIN MAX
                                             SYSNAME         SYSNAME
CLRSHARE
 D15.UNIPMOD2                                    (SYSTEMS)
        0.008  0.494  1.579  0.225    1 RINGEL  (S)   1 D75ARP1#(E)   100  0.0  0.0  0.0   1.00   0   1   0   1     7
                                          SYS1            SYS1
 CLRVSAM.PSC1BA.UNILIB                           (SYSTEMS)
        0.000  0.494 19.741  0.111    1 D75ARP1#(E)   1 SHIPPEEP(S)   100  0.0  0.0  0.0   1.00   0   1   0   1   177
                                          AQFT            SYS1
 D15.UNIPLAN2                                    (SYSTEMS)
        0.014  0.724  1.716  0.143    1 HOBBIT  (S)   1 VIRKO   (E)   100  0.0  0.0  0.0   1.00   0   1   0   1    12
                                          SYS1            SYS1
 POKLCG.SHADOW.BANK                              (SYSTEMS)
        0.008  0.445 12.712  0.189    1 SHIPPEEP(E)   1 D75ARP1#(E)   100  0.0  0.0  0.0   1.00   1   1   0   0    67
                                          SYS1            SYS1
 CLRVSAM.SDUMP.UNILIB                            (SYSTEMS)
        0.000  3.332 29.582  0.580    1 ##732517(E)   3 ##732526(E)  25.5 17.6 56.9  0.0   2.31   0   3   0   2    51
                                          AQFT            SYS1
                                                         ##732521(S)
                                                          SYS1
 CLRVSAM.SDUMP.BULK0001                          (SYSTEMS)
        0.022  0.737  1.575  0.393    1 ##732567(E)   1 ##732578(S)   100  0.0  0.0  0.0   1.00   0   1   0   1     4
                                          AQFT            SYS1
 CLRVSAM.JOBLOG.BULK0001                         (SYSTEMS)
        0.236 43.436 84.955 10.619    1 ##PDSRL2(E)   1 ##PDSRL1(E)  87.5 12.5  0.0  0.0   1.13   1   2   0   0     8
                                          SYS1            SYS1
 CLRVSAM.JOBLOG.UNIBCKUP                         (SYSTEMS)
        0.012  0.337  1.853  0.185    1 ##PDSRL1(E)   1 ##PDSRL3(E)   100  0.0  0.0  0.0   1.00   1   1   0   0    10
                                          SYS1            SYS1
 CLRVSAM.DAE.UNILIB                              (SYSTEMS)
        0.008  1.161  3.348  0.418    1 ##732686(E)   1 EMILLER (S)   100  0.0  0.0  0.0   1.00   0   1   0   1     8
                                          AQFT            AQFT
```

Figure 5-25. Enqueue Detail Activity Report

Table 5-30. 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.
CONTENTION TIME (MIN MAX TOT AVG)	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. 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. The calculation used to determine the average contention time is: `AVG CONT TIME = Contention Time for the Resource / # Contention Events`

Table 5-30. Fields in the Enqueue Activity Report (continued)

Field Heading	Meaning
JOBS AT MAXIMUM CONTENTION	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. 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. The field is reported only when the enqueue activity detail report is requested.

% QLEN DISTRIBUTION (1 2 3 4+)	The percentage of contention events during the interval when the number of requestors queued to the resource was one, two, three, or four or more. The samples are taken during the period when contention was occurring. 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.
	
	The calculation used for each queue length is:
	`Accumulator for that Queue Length`
	`% QLEN = ---------------------------------- * 100`
	`# Enqueue Contention Events`

AVG Q LNGTH	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.
	`# Requestors Waiting`
	`AVG Q LNGTH = ----------------------------`
	`# Enqueue Contention Events`

-REQUEST TYPE- -EXCL -- SHARE - (MIN MAX MIN MAX)	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.
TOTAL EVENT	The total number of resource contention events that occur during the measurement 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 5-31. Overview Names in the Enqueue Activity Report

Field Heading	Overview Name
CONTENTION TIME - TOT	ENQT
CONTENTION TIME - AVG	ENQAVG
CONTENTION TIME - MAX	ENQMAX
TOTAL EVENT	ENQNE

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.

FCD - FICON Director Activity report

How to request this report

The default option for Monitor I data gathering is NOFCD. See the *z/OS RMF User's Guide* for details. Therefore, it is required that you specify FCD if you want to get this report.

To print this report, specify

```
REPORTS(FCD(option))
```

Purpose of the 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

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
                                                                                                PAGE    1
          z/OS V1R4               SYSTEM ID   SYS1           DATE  06/03/2002          INTERVAL 15.00.000
                                  RPT VERSION V1R2 RMF       TIME  07.15.00            CYCLE 1.000 SECONDS

IODF = 99   CR-DATE: 06/01/2002   CR-TIME: 08.45.00     ACT: POR

SWITCH DEVICE: 0414   SWITCH ID: 01   TYPE: 005000   MODEL: 001   MAN: MCD   PLANT: 01   SERIAL: 00000MK00109

PORT     -CONNECTION-    AVG FRAME     AVG FRAME SIZE      PORT BANDWIDTH (MB/SEC)     ERROR
ADDR     UNIT      ID       PACING       READ   WRITE      -- READ --  -- WRITE --     COUNT

 05      CHP       FA            0        808     285           50.04        10.50         0
 07      CHP       4A            0        149     964           20.55         5.01         0
 08      CHP-H     F4            0        568  965.5T           70.56         4.02         1
 09      CHP       FC            0        558    1424           50.07        10.53         0
 0B      CHP       F4            0        872     896           50.00        10.56         0
 12      CHP       D5            0         73     574           20.51         5.07         0
 13      CHP       C8            0        868    1134           70.52         2.08         1
 14      SWITCH  ----            0        962     287           50.03        10.59         0
 15      CU      C800            0       1188     731           20.54         5.00         0
         CU      CA00                                           70.55         3.01         1
 16      CHP       CB            0        740    1185           70.50         2.06         1
```

Figure 5-26. FICON Director Activity Report

Table 5-32. 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.

Table 5-32. Fields in the FICON Director Activity Report. (continued)

Field Heading	Meaning
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 ID of the FICON director which is associated with this switch device.
TYPE MODEL MAN PLANT SERIAL	The hardware description which characterizes the switch device.
PORT ADDR	The hexadecimal address of the port.
CONNECTION	Provides information about the connected unit. **UNIT** • CHP: denotes a channel path • CHP-H: denotes a channel path of the system which requested this report • CU: denotes a control unit • SWITCH: denotes a switch 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. **ID** The hexadecimal identifier of the connector. For connection unit SWITCH, dashes are provided. 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.
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 5-33. Overview Names in the FICON Director Activity Report

Field Heading	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

How to produce 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 the report, specify

`REPORTS(HFS)`

Purpose of the 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.

Using the information given in 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.

Contents of the report

HFS Global Statistics

```
                                    H F S   G L O B A L   S T A T I S T I C S
                                                                                                  PAGE    1
          z/OS V1R4              SYSTEM ID   SYS1           DATE  04/03/2002          INTERVAL 15.00.000
                                 RPT VERSION V1R2 RMF       TIME  12.00.00            CYCLE 1.000 SECONDS

       --- STORAGE LIMITS (MB) --                     ----- FILE I/O ----  --- METADATA I/O --
                                                          COUNT     RATE       COUNT     RATE
       VIRTUAL      MAX      1000
                    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    BUFFER   -------- POOL SIZE --------      DATA      ---------- I/O ACTIVITY ---------
NUMBER      BUFFERS      SIZE     PAGES     BYTES    %FIXED    SPACES        TOTAL        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 5-27. HFS Global Statistics Report

Table 5-34. Fields in HFS Global Statistics Report

Field Heading	Meaning
Storage Limits - 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.
File I/O - 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.
Metadata I/O - 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.
Buffer Pool Statistics	
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.

Table 5-34. Fields in HFS Global Statistics Report (continued)

Field Heading	Meaning
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

```
                                H F S   F I L E   S Y S T E M   S T A T I S T I C S
                                                                                                                PAGE    2
           z/OS V1R4               SYSTEM ID   SYS1           DATE  04/03/2002           INTERVAL 15.00.000
                                   RPT VERSION V1R2 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: 04/03/2002   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: 04/03/2002   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 5-28. HFS File System Statistics Report

Table 5-35. 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.
Allocation - 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.
File I/O - The fields are given as COUNT and RATE (count per second).	

Table 5-35. Fields in the HFS File System Statistics Report (continued)

Field Heading	Meaning
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.
Metadata I/O - 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.
HIT RATIO	Percentage of cache-found requests based on total number of requests.
Index I/O - 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.
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.

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

How to request this report

The Postprocessor is using SMF records type 103 as input for the HTTP Server report. These records are not gathered by an RMF monitor but are written by HTTP servers. See the *z/OS RMF User's Guide* for details.

To print this report, specify

```
REPORTS(HTTP)
```

Purpose of the 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.

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                                        PAGE   1
           z/OS V1R4               SYSTEM ID   SYS1           DATE 03/14/2002          INTERVAL 30.00.000
                                   RPT VERSION V1R2 RMF       TIME 08.06.37
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                         000.24.35      0.03     0.03      5.46   30.15      39  0.00   5120  0.90      0  0.00      1
    4F37-CA73-0005-678E
mvs047.tcp.raleigh.ibm.com     000.01.10      2.14     0.07      8.67   41.37      23  0.00   5120  1.06      0  0.00      0
    3D29-CA40-0006-997C
s390server17.wsclab.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 5-29. HTTP Server Report - Summary

PP - HTTP Report

```
                                    H T T P   S E R V E R   D E T A I L S
                                                                                                                    PAGE   2
          z/OS V1R4              SYSTEM ID   SYS1           DATE 03/14/2002          INTERVAL 30.00.000
                                 RPT VERSION V1R2 RMF       TIME 08.06.37

------------------------------------------------------------------------------------------------------------------------------------
                                            SERVER CHARACTERISTICS
------------------------------------------------------------------------------------------------------------------------------------
SERVER NAME:        s390server17.wsclab.washington.ibm.com
SERVER ROOT IN HFS: /webtime/development/pr
IP-ADDRESS:         9.67.116.9                  STARTUP:            02/11/2002-08.03.01
PORT:               8181                        SECURITY TYPE:      1
TYPE:               PROXY                       SSL-PORT:           443
APPL-LVL:           V5R2M0                      SERVER TOKEN:       3D29-CA4C-000B-7E86

--------- FLAGS ---------     ------- RESOURCES ------                      ------- CACHE -------     -- TIMEOUT THRESHOLDS --
DNS LOOKUP            YES     MAX BUFFER          102K                     CACHE             YES     INPUT                330
ACL SETTINGS          YES     MAX THREADS          39                      MAX SIZE         5120     OUTPUT              3600
META FILE              NO                                                  MAX FILES          NO     SCRIPT               600
DIRECTORY ACCESS       NO     -- GARBAGE COLLECTION --                    LIMIT 1           100     IDLE THREADS           0
SERVER IMBEDS HTML     NO     ENABLED               NO                    LIMIT 2          4000     CACHE LOCK           120
NORMAL MODE           YES     INTERVAL           10800                    TIME MARGIN       120
GMT                    NO     MEMORY USE           500                    KEEP EXPIRED       NO
PROXY                 YES                                                 CONNECT            NO

------------------------------------------------------------------------------------------------------------------------------------
                                                SERVER ACTIVITY
------------------------------------------------------------------------------------------------------------------------------------

------------ ACTIVITY ------------     - REQUEST TYPES -     ------ THREADS ------     -- CACHE USAGE --  ---- MISCELLANEOUS ---
                      COUNT    RATE     GET             2     MAX                39     KBYTES READ     0  TIMEOUTS             1
REQUESTS                  1    0.01     POST            0     USED                0     HITS            0  CONNECTIONS          2
REQUESTS DISCARDED        0    0.00     CGI             0     NON-SSL WAIT       17     IN USE      58930
REQUESTS IN ERROR         1    0.01     GWAPI           1     SSL WAIT           16     FILES           0
RESPONSES                 3    0.02                           ASYNC I/O WAIT      0
RESPONSES DISCARDED       0    0.00                           MSG QUEUE WAIT      0

----------- THROUGHPUT ------------     ------------- RESPONSE TIMES --------------     ------------- ERROR STATISTICS ----------
                     BYTES    RATE                          MIN      MAX      AVG     200 (OK)                                1
IN                     893    4.85     DNS LOOKUP         0.000    0.000    0.000     302 (MOVED TEMPORARILY)                 0
OUT                    730    3.97     SERVICE PLUGIN     0.000    0.170    0.060     401 (UNAUTHORIZED)                      0
UNKNOWN                  0    0.00     CGI                1.580    1.580    1.580     403 (FORBIDDEN)                         0
                                       SSL HANDSHAKE      0.000    0.000    0.000     404 (NOT FOUND)                         2
                                       PROXY              9.580    9.580    9.580     407 (PROXY UNAUTHORIZED)                0
                                                                                     500 (INTERNAL SERVER ERROR)             0
```

Figure 5-30. HTTP Server Report - Details

Table 5-36. 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 5-29 on page 5-87.
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	**MAX** Maximum number of threads the server can have in the thread pool (or NO if no limit has been specified). **USED** Number of currently active threads of the server.

Table 5-36. Fields in the HTTP Server Summary Report (continued)

Field Heading	Meaning
CACHE SIZE	**MAX** Maximum cache size (KB) for this server. **USED** Used cache size of this server.
CACHE FILES	**MAX** Maximum number of files to be in the cache of this server. **USED** Number of files in the cache of this server.
TIMEOUTS	Number of timeouts on the server.

Table 5-37. Fields in the HTTP Server Details Report

Field Heading	Meaning
Server Characteristics - Configuration Data	
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. **HTTP** Simple or normal HTTP server **PROXY** Proxy server **CACHING** Caching server **CACHING PROXY** Caching proxy **UNKNOWN** 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.
Server Characteristics - Flags	
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.
Server Characteristics - Resources	
MAX BUFFER	Maximum size of content buffer.
MAX THREADS	Maximum number of threads the server can have in the thread pool.
Server Characteristics - Garbage Collection	
ENABLED	Indication whether garbage collection is enabled.
INTERVAL	Garbage collection interval in seconds.
MEMORY USE	Garbage collection memory usage.
Server Characteristics - Cache	
CACHE	Cache flag.

Table 5-37. Fields in the HTTP Server Details Report (continued)

Field Heading	Meaning
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.
Server Characteristics - Timeout Thresholds (in seconds)	
INPUT	Input timeout.
OUTPUT	Output timeout.
SCRIPT	Script timeout.
IDLE THREADS	Timeout for idle threads.
CACHE LOCK	Cache lock timeout.
Server Activity - Requests - 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.
Server Activity - Request Types	
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.
Server Activity - Threads	
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.
Server Activity - Cache Usage	
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.

Table 5-37. Fields in the HTTP Server Details Report (continued)

Field Heading	Meaning
Server Activity - Miscellaneous	
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.
Server Activity - Throughput - 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.
Server Activity - Response Times - The values are given as minimum, maximum and average response time (in seconds). **Note:** 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.
Server Activity - Error Statistics - The number of responses with a specific error code.	
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

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)
```

Purpose of the 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.

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.

Note: Starting with z/OS V1R2 RMF, the percentage about all channel paths being busy has been removed due to data inconsistencies for different channel types.

Using the information given in the report

The information in the I/O Queuing Activity report, used after the channel path activity and I/O Device Activity reports have shown that a problem exists, can help you 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 LNGTH, 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  Q U E U I N G  A C T I V I T Y
                                                                                                  PAGE    1
          z/OS V1R4               SYSTEM ID   SYS1          DATE 03/14/2002          INTERVAL 15.00.036
                                  RPT VERSION V1R2 RMF      TIME 16.30.00            CYCLE 1.000 SECONDS

TOTAL SAMPLES =    900   IODF = 01   CR-DATE: 05/10/2000   CR-TIME: 21.00.01     ACT: POR


      - INITIATIVE QUEUE -    ------- IOP UTILIZATION -------    -- % I/O REQUESTS RETRIED --    -------- RETRIES / SSCH ---------
 IOP      ACTIVITY    AVG Q     % IOP   I/O START    INTERRUPT          CP    DP    CU    DV             CP      DP      CU      DV
            RATE      LNGTH      BUSY      RATE         RATE       ALL  BUSY  BUSY  BUSY  BUSY     ALL   BUSY    BUSY    BUSY    BUSY

  00       1282.40     0.15     54.58      853.56       984.56    67.7  34.0  15.1  11.3   7.3     2.10   1.05    0.47    0.35    0.23
  02       1188.30     0.11     51.25      797.23       856.23    63.4  27.5   9.2  18.3   8.4     1.74   0.75    0.25    0.50    0.23

 SYS       2470.70     0.13     52.83     1650.79      1840.79    65.8  31.1  12.4  14.5   7.8     1.92   0.91    0.36    0.42    0.23

                                                                                 AVG   AVG                    DELAY   AVG
 LCU    CONTROL UNITS     DCM GROUP    CHAN     CHPID     % DP    % CU    CUB   CMR      CONTENTION      Q     CSS
                        MIN MAX DEF   PATHS    TAKEN     BUSY    BUSY    DLY   DLY          RATE      LNGTH   DLY

 0031   400A                          11      0.573     0.00    0.00    0.0   1.2
        400B                          20      0.583    28.16    0.41    1.0   0.1
                         2   2   2             2.447    10.90    0.24    0.0   0.0
                                    *         3.604    12.79    0.24    0.0   0.2         0.000     0.00    0.0

 0032   440A                          11      0.013     0.00    0.00    0.0   0.1
        440B                          20      0.027    33.33    0.00    0.0   1.2
                         0   0   2             0.000     0.00    0.00    1.0   0.1
                                    *         0.040    25.00    0.00    0.0   0.0         0.000     0.00    0.0

 0033   500A                          12      0.070    34.38    0.00    0.0   0.2
        500B                          93      0.027    76.47    0.00    0.0   0.1
                                    *         0.097    56.06    0.00    0.0   1.2         0.010     0.00    0.0

 0034   700A                          12      0.023    80.56    0.00    1.0   0.1
        700B                          93      0.027     0.00    0.00    0.0   0.0
                                    *         0.050    65.91    0.00    0.0   0.2         0.000     0.00    0.0

 0035   580A                          11      0.033     0.00    0.00    0.0   0.1
        580B                          20      0.017    25.00   12.50    0.0   1.2
                         0   0   2             0.000     0.00    0.00    0.0   1.2
                                    *         0.050    11.11    5.56    0.0   1.2         0.000     0.00    0.0

 0036   5C0A                          11      0.390     0.00    0.85    1.0   0.1
        5C0B                          20      0.497    27.44    3.26    0.0   0.0
                         2   2   2             2.240     9.13    0.67    0.0   0.2
                                    *         3.127    11.78    1.21    0.0   0.1         0.000     0.00    0.0
```

Figure 5-31. I/O Queuing Activity Report

In a PR/SM environment, the top section of the report measures PR/SM system activity. The rest of the report applies to I/O activity for the MVS being measured by RMF.

In a 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.

```
                                    I/O   Q U E U I N G   A C T I V I T Y
                                                                                                                PAGE  1
          z/OS V1R4               SYSTEM ID   MVSI           DATE 06/15/2002          INTERVAL 02.00.003
                                  RPT VERSION V1R2 RMF       TIME 14.16.00            CYCLE 1.000 SECONDS

TOTAL SAMPLES =    300   IODF = 01   CR-DATE: 04/14/2001   CR-TIME: 13.01.15     ACT: POR
       - INITIATIVE QUEUE -
 IOP      ACTIVITY    AVG Q
            RATE      LNGTH
 NO H/W DATA

                                                                                                    DELAY
 LCU    CONTROL UNITS    DCM GROUP   CHAN     CHPID    % DP   % CU  CONTENTION     Q
                       MIN MAX DEF  PATHS    TAKEN    BUSY   BUSY     RATE      LNGTH

 0008   3040                          40      NO H/W DATA
                                      80      NO H/W DATA

 0009   1780                          17      NO H/W DATA
                                      83      NO H/W DATA

 000A   10C0                          10      NO H/W DATA
                                      82      NO H/W DATA

 000B   10D0                          1D      NO H/W DATA
                                      9D      NO H/W DATA

 000F   2680                          36      NO H/W DATA
                                      B6      NO H/W DATA
```

Figure 5-32. IOQ Report - z/OS Running as VM Guest System

Table 5-38. Fields in the I/O Queuing 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.
IOP	The two digit hexadecimal identifier of the I/O processor (IOP). The IOP data sections are sorted according ascending IOP numbers. A summary line which summarizes the the measurement data of the individual IOPs is added following the last IOP data line. This summary line starts with the character string *SYS*, indicating that it contains system wide information.
Initiative Queue	
ACTIVITY RATE	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. $$\text{ACTIVITY RATE} = \frac{\text{\# I/O Requests on the IOP Queue}}{\text{Interval}}$$
AVG Q LNGTH	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. $$\text{AVG Q LNGTH} = \frac{\text{Accumulated Queue Length}}{\text{\# I/O Requests on the IOP Queue}} - 1$$
IOP Utilization	

Table 5-38. Fields in the I/O Queuing Activity Report (continued)

Field Heading	Meaning

% IOP BUSY	The ratio of the number of times the IOP was found busy and the total number of I/O processor samples.
	`# Busy samples`
	`% IOP BUSY = ---------------------------------- * 100`
	`# Busy samples + # Idle samples`

I/O START RATE	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.
	`# I/O functions started`
	`I/O START RATE = -------------------------`
	`Interval`

INTERRUPT RATE	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.
	`# Processed I/O interrupts`
	`INTERRUPT RATE = ----------------------------`
	`Interval`

% I/O Requests Retried	

ALL	The ratio of the number of retries and the number of I/O functions initially started plus the total number of retries.
	`# Retries`
	`%ALL = ------------------------------------ * 100`
	`# I/O functions started + # Retries`

CP BUSY	The ratio of the number of times an I/O operation was retried on the I/O processor because the selected channel path was busy and the number of I/O functions initially started plus the total number of retries.
	`# Retries due to channel path busy`
	`%CP BUSY = ------------------------------------ * 100`
	`# I/O functions started + # Retries`

DP BUSY	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 and the number of I/O functions initially started plus the total number of retries.
	`# Retries due to director port busy`
	`%DP BUSY = ------------------------------------ * 100`
	`# I/O functions started + # Retries`

CU BUSY	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 and the number of I/O functions initially started plus the total number of retries.
	`# Retries due to control unit busy`
	`%CU BUSY = ------------------------------------ * 100`
	`# I/O functions started + # Retries`

DV BUSY	The ratio of the number of times an I/O operation was retried on the I/O processor because the targeted device was busy and the number of I/O functions initially started plus the total number of retries.
	`# Retries due to device busy`
	`%DV BUSY = ------------------------------------ * 100`
	`# I/O functions started + # Retries`

Retries / SSCH	

Table 5-38. Fields in the I/O Queuing Activity Report (continued)

Field Heading	Meaning

ALL	The ratio of the number of retries on the I/O processor and the number of I/O functions initially started.
	`# Retries`
	`ALL = ------------------------`
	`# I/O functions started`

CP BUSY	The ratio of the number of retries on the I/O processor because the selected channel path was busy and the number of I/O functions initially started.
	`# Retries due to channel path busy`
	`CP BUSY = -----------------------------------`
	`# I/O functions started`

DP BUSY	The ratio of the number of retries on the I/O processor because a director port on the path to that device was busy and the number of I/O functions initially started.
	`# Retries due to director port busy`
	`DP BUSY = -----------------------------------`
	`# I/O functions started`

CU BUSY	The ratio of the number of retries on the I/O processor because the control unit of the targeted device was busy and the number of I/O functions initially started.
	`# Retries due to control unit busy`
	`CU BUSY = -----------------------------------`
	`# I/O functions started`

DV BUSY	The ratio of the number of retries on the I/O processor because the targeted device was busy and the number of I/O functions initially started.
	`# Retries due to device busy`
	`DV BUSY = ----------------------------`
	`# I/O functions started`

LCU	The four-digit hexadecimal identifier of the logical control unit (LCU). 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. 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.
CONTROL UNITS	The four-digit hexadecimal identifier of each physical control unit contained in the logical control unit.
DCM GROUP MIN - MAX - DEF	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. 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.

Table 5-38. Fields in the I/O Queuing Activity Report (continued)

Field Heading	Meaning

CHAN PATHS	The two-digit hexadecimal channel path identifier (CHPID) of the channel paths that are attached to the physical control units contained in the logical control unit. There can be up to four 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:
	`OFFLINE`
	`NOW OFFLINE`
	`NOW ONLINE`
	
	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:
	`PATH OFFLINE`
	`PATH NOW OFFLINE`
	`PATH NOW ONLINE`
	
	An '*' in this column indicates a summary line for all channel paths connected to the same LCU.

CHPID TAKEN	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.
	`# I/O Operations Accepted on that Path`
	`CHPID TAKEN = ----------------------------------`
	`Interval`
	
	Note: If vary activity has occurred during the interval, this field is blank.

% DP BUSY	The ratio of the number of times an I/O request was deferred because the director port was busy and the number of attempts to service I/O requests during the measurement interval. This field indicates director port contention.
	`DPB`
	`% DP BUSY = --------------- * 100`
	`DPB + CUB + SUC`
	DPB Number of deferred I/O requests due to director port busy
	CUB Number of deferred I/O requests due to control unit busy
	SUC Number of successful I/O requests on that path

% CU BUSY	The ratio of the number of requests deferred due to control unit busy and 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 logical control unit.
	`CUB`
	`% CU BUSY = --------------- * 100`
	`DPB + CUB + SUC`
	DPB Number of deferred I/O requests due to director port busy
	CUB Number of deferred I/O requests due to control unit busy
	SUC Number of successful I/O requests on that path
	
	RMF reports a value even if the channel path changes status during the interval.

AVG CUB DLY	The average number of milliseconds of delay that an I/O request encountered for the channel path because the control unit was busy.
	`Control Unit Busy Time`
	`AVG CUB DLY = ----------------------------------`
	`# I/O Operations Accepted on that Path`

Table 5-38. Fields in the I/O Queuing Activity Report (continued)

Field Heading	Meaning

AVG CMR DLY	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).
	`Initial Command Response Time`
	`AVG CMR DLY = ------------------------------------`
	`# I/O Operations Accepted on that Path`

CONTENTION RATE	The rate at which the I/O processor places delayed I/O requests on the CU-HDR for this LCU. 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 via an alternate path, the IOP does not count the condition.
	`# Enqueued Requests`
	`CONTENTION RATE = -------------------`
	`Interval`

DELAY Q LNGTH	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.
	`Accumulated Queue Length`
	`DELAY Q LNGTH = ------------------------ - 1`
	`# Enqueued Requests`

AVG CSS DLY	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 first attempts to initiate the operation.
	`Channel Subsystem Time`
	`AVG CSS DLY = -------------------------`
	`# I/O Operations Accepted`

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 5-39. Overview Names in the I/O Queuing Activity Report

Column Heading	Overview Name
ACTIVITY RATE	IOPAC
AVG Q LNGTH	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

Table 5-39. Overview Names in the I/O Queuing Activity Report (continued)

Column Heading	Overview Name
% 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 LNGTH	IODLQ
AVG CUB DLY	IOCBT
AVG CMR DLY	IOCMR
AVG CSS DLY	IOCSS
Percentage of all channel path busy	IOCHPB — This condition is not available anymore.

OMVS - OMVS Kernel Activity report

How to produce 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 the report, specify
`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

```
                                    O M V S   K E R N E L   A C T I V I T Y
                                                                                                                                  PAGE   1

           z/OS V1R4               SYSTEM ID   AQTS           DATE 03/08/2002      INTERVAL 30.00.000
                                   RPT VERSION V1R2 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
-------------------------------------------------------------------------------------------------------------------------------------
                          PROCESSES                       USERS                  PROCESSES PER USER
 MAXIMUM  (TOT)             1200                           50                          12
-------------------------------------------------------------------------------------------------------------------------------------
                 MINIMUM  AVERAGE  MAXIMUM      MINIMUM  AVERAGE  MAXIMUM      MINIMUM  AVERAGE  MAXIMUM
-------------------------------------------------------------------------------------------------------------------------------------
 CURRENT  (TOT)       99      854     1200           12       13       23
 OVERRUNS (N/S)        0     5.5*      333            0     1.5*     4.8            0     1.0*     5.3

                                              OMVS INTER-PROCESS COMMUNICATION
-------------------------------------------------------------------------------------------------------------------------------------
                  MESSAGE QUEUE IDS              SEMAPHORE IDS            SHARED MEMORY IDS          SHARED MEMORY PAGES
 MAXIMUM  (TOT)             500                        500                        500                       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                    SHARED LIB REGION           QUEUED SIGNALS
--------------   ---------------------------------------------------   ------------------------   ----------------------
                 MEMORY MAP STORAGE PAGES       SHARED STORAGE PAGES     MAX SHARED LIBRARY REGION   MAXIMUM QUEUED SIGNALS
 MAXIMUM  (TOT)             4096                       131K                       16M                        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
 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 5-33. 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 5-40. 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 during this interval.
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	

Table 5-40. Fields in the OMVS Kernel Activity Report (continued)

Heading	Meaning
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.
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/Swap Data Set Activity report

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 a report, specify

`REPORTS(PAGESP)`

Purpose of the report

Note: In z/OS, swap data sets are not supported. Therefore, information about swap data sets is available only for SMF data of systems prior to OS/390 2.10. Please, refer to *OS/390 2.7 RMF Report Analysis* for a description of swap data set activity.

The Page/Swap Data Set Activity report provides information about page data set use for each individual data set. The information about the number of slots used is sampled and reported as minimum, maximum, and average values for the interval. Other information in the report gives the time 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 service time or 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

Using the information given in the report

The information in the page data set section of the report can be used, 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.
- Use swap only in special cases to help important non-swappable work at the expense of swappable work.
- 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

```
                                P A G E / S W A P  D A T A S E T  A C T I V I T Y
                                                                                PAGE  1
          z/OS V1R4              SYSTEM ID   SYS1          DATE 03/14/2002         INTERVAL 14.59.946
                                 RPT VERSION V1R2 RMF      TIME 16.30.00           CYCLE 1.000 SECONDS

NUMBER OF SAMPLES =     900                    PAGE DATA SET USAGE
                                              -------------------------

PAGE                                                                  %    PAGE                 V
SPACE    VOLUME  DEV  DEVICE    SLOTS  ---- SLOTS USED ---  BAD       IN   TRANS  NUMBER  PAGES  I
TYPE     SERIAL  NUM  TYPE      ALLOC    MIN    MAX   AVG   SLOTS     USE  TIME   IO REQ  XFER'D 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   1,933  12,191 Y  SYS1.PGT1B5.LOCAL
LOCAL    PGT80A  080A 33903    540000  42338  43947  43012      0    4.67  0.003   2,069  13,129 Y  SYS1.PGT80A.LOCAL
LOCAL    PGT80C  080C 33903    540000  41388  43197  42128      0    6.22  0.004   2,106  13,299 Y  SYS1.PGT80C.LOCAL
LOCAL    PGT857  0857 33903    540000  39245  40736  39858      0    5.33  0.004   1,987  12,575 Y  SYS1.PGT857.LOCAL
LOCAL    PGT859  0859 33903    540000  39809  41273  40438      0    5.89  0.004   1,969  12,546 Y  SYS1.PGT859.LOCAL
```

Figure 5-34. 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, SLOTS ALLOC, and SWAP SETS 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 5-41. Fields in the Page/Swap Data Set Activity Report

Field Heading	Meaning
PAGE SPACE TYPE	Page space type, which can be PLPA, COMMON, or LOCAL.
VOLUME SERIAL	Volume serial number of the volume on which the data set resides.
DEV NUM	Number of the device on which the data set resides.
DEVICE TYPE	Device type on which the data set resides.
SLOTS ALLOC	The total number of slots each page data set contains.
----SLOTS USED ---	The number of slots that were being used for pages. **Note:** The actual minimum and maximum values might occur at a time when RMF is not sampling.
BAD SLOTS	Number of slots that encountered permanent I/O errors.

Table 5-41. Fields in the Page/Swap Data Set Activity Report (continued)

Field Heading	Meaning

% IN USE	Percentage of time during the reporting interval when the data set was considered busy by the Auxiliary Storage Manager (ASM).
	
	At each cycle, RMF tests each data set, and at the end of the interval, the percentage is calculated.
	`# Busy Samples`
	`% IN USE = -------------- * 100`
	`# Samples`

PAGE TRANS TIME	Average number of seconds required to complete a page transfer.
	`(USE*INT)/N`
	`PAGE TRANS TIME = -----------`
	`XFER`
	USE Number of samples when the data set was in use
	XFER Total number of pages transferred
	N Number of samples
	INT Interval time (seconds)

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.
DATA SET NAME	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. **Notes:** • If a data set was dynamically introduced during the interval, its data set name is preceded by an asterisk (*). • 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.
VIO	Indication of whether the local paging data set accepts VIO pages. The symbols are: **Y** VIO pages are accepted **N** VIO pages are not accepted

Overview reference

Table 5-42. Overview Names in the Page/Swap Data Set Activity Report

Field Heading	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

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 a report, specify
REPORTS(PAGING)

Purpose of the report

The Paging Activity report provides information about the demands made on the system paging facilities and the use of central storage, expanded storage, and external page storage during the interval. It also provides swapping statistics for the interval.

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. This might occur because of a specification in the IPS of minimum multiprogramming levels that are too high for the system configuration and current workload mix.

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.

In the Expanded Storage Movement Rates section of the report, a low **AVG HIGH UIC** value might indicate that more central storage is needed.

A low **AVG MIGR AGE** together with high migration rates can suggest expanded storage contention.

Contents of the report

The header of each report section contains the information about the mode the system is running in:

- MODE = ESA - The system is running in 31-bit addressing mode.
- MODE = ESAME - The system is running in 64-bit addressing mode.

The Paging Activity report is formatted into the following four sections, each with a separate subheading:

- CENTRAL STORAGE PAGING RATES
- Depending on the mode the system is running, the second section is different:
  MODE=ESA - EXPANDED STORAGE MOVEMENT RATES
  MODE=ESAME - CENTRAL STORAGE MOVEMENT RATES
- FRAME AND SLOT COUNTS
- SWAP PLACEMENT ACTIVITY

Central Storage Paging Rates section

```
                                        P A G I N G   A C T I V I T Y
                                                                                                                    PAGE    1
          z/OS V1R4               SYSTEM ID   SYS1          DATE 03/16/2002          INTERVAL 60.00.378
                                  RPT VERSION V1R2 RMF      TIME 09.00.00            CYCLE 0.333 SECONDS

OPT = IEAOPT00   MODE = ESA                  CENTRAL STORAGE PAGING RATES - IN PAGES PER SECOND
-------------------------------------------------------------------------------------------------------------------------------------------

                        PAGE IN                                          PAGE OUT
                  ------------------------------------        ----------------------------
                                        NON           %                                 %
                           NON        SWAP,          OF                                OF
                          SWAP,  NON    NON    TOTAL  TOTL              NON    TOTAL   TOTL
CATEGORY          SWAP    BLOCK  BLOCK  RATE   RATE   SUM       SWAP   SWAP   RATE    SUM

PAGEABLE SYSTEM
AREAS (NON VIO)
     LPA                          0.07   0.13    1

     CSA                          0.20   0.28    3                0.00   0.00    0
                                 ------ ------ ----               ------ ------  ----
     SUM                          0.27   0.42    4                0.00   0.00    0

ADDRESS SPACES
     HIPERSPACE            0.00          0.00    0                0.00   0.00    0

     VIO                   0.01          0.01    0                0.00   0.00    0

     NON VIO       1.91    4.51   3.50   9.76   96         0.37   0.00   0.37   100
                  ------  ------ ------ ------ ----       ------ ------ ------  ----
     SUM           1.91    4.52   3.50   9.77   96         0.37   0.00   0.37   100

TOTAL SYSTEM
     HIPERSPACE            0.00          0.00    0                0.00   0.00    0

     VIO                   0.01          0.01    0                0.00   0.00    0

     NON VIO       1.91    4.51   3.76   10.18  100         0.37   0.00   0.37   100
                  ------  ------ ------ ------ ----       ------ ------ ------  ----
     SUM           1.91    4.52   3.76   10.19  100         0.37   0.00   0.37   100

     SHARED                       0.00   0.00                     0.00   0.00


PAGE MOVEMENT WITHIN CENTRAL STORAGE      41.30          PAGE MOVEMENT TIME %     0.0
AVERAGE NUMBER OF PAGES PER BLOCK          5.2
BLOCKS PER SECOND                          0.87
PAGE-IN EVENTS (PAGE FAULT RATE)           4.64
```

Figure 5-35. PAGING Report - Central Storage Paging Rates Section

This section of the Paging Activity report monitors paging rates in central storage. 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
- Total rate
- % of total sum

The page-in non-swap group is further divided into:

- Blocked
- Non-blocked

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
- Address space pages used, broken down into hiperspace data, VIO data and non-VIO data

The rate of page movement within central storage is shown in the bottom left corner of the page.

Table 5-43. Fields in the Paging Activity Report - Central Storage Paging Rates Section

Field Heading	Meaning
CATEGORY	The component parts of paging rates identifying these basic components: • Pageable system area, non-VIO data • Address space data • Total system data.
PAGEABLE SYSTEM AREAS (NON VIO)	The areas of central storage that are not associated with a single address space. This section consists of: **LPA** All values are reported except for swaps **CSA** All values are reported except for swaps **SUM** Sum of LPA and CSA.
ADDRESS SPACES	The areas of central storage that are associated with individual address spaces. This section consists of: **HIPERSPACE** All values are reported except for swaps **VIO** All values are reported except for swaps **NON VIO** All values are reported **SUM** Sum of address space hiperspace, VIO and non-VIO.
TOTAL SYSTEM	The sum of system pageable areas and address space values and the following: **HIPERSPACE** Consists of address space hiperspace values **VIO** Consists only of address space VIO values **NON VIO** Sum of system pageable areas non-VIO and memory non-VIO values **SUM** Sum of system pageable areas sum and address space sum. (The computer system total for paging rates.) **SHARED** Number of shared page group page-ins and page-outs in central storage. The page-in/out rate is included in the SUM values.

Table 5-43. Fields in the Paging Activity Report - Central Storage Paging Rates Section (continued)

Field Heading	Meaning
PAGE IN	The rate of pages read into central storage. **SWAP** 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. **NON SWAP,BLOCK** 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. **NON SWAP,NON BLOCK** 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. **TOTAL RATE** 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. **% OF TOTL SUM** The percentage of the total page-in rate for each part of the total. The total here is reflected in SUM under TOTAL SYSTEM.

Table 5-43. Fields in the Paging Activity Report - Central Storage Paging Rates Section (continued)

Field Heading	Meaning
PAGE OUT	The rate of pages written to auxiliary storage. **SWAP** The rate of pages written to auxiliary storage as a result of address space swap outs. There is no PAGE OUT for shared storage due to SWAP. **NON SWAP** 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). 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. A hiperspace page out occurs when a standard hiperspace page is no longer needed in central storage and is written to auxiliary storage. VIO pages transferred as a result of a swap-out are not included. **TOTAL RATE** 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. **% OF TOTL SUM** The percentage of the total page-out rate for each part of the total. The total here is reflected in the subcategory SUM under TOTAL SYSTEM.
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 processor time expended on page movement to obtain or free a particular type of frame for a page to be fixed (i.e. 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 5-44. Overview Names in the Paging Activity Report - Central Storage Paging Rates Section

Column Heading	Overview Name
PAGE IN - TOTAL RATE	TPAGRT, PAGERT, DPAGRT
PAGE OUT - TOTAL RATE	TPAGRT, PAGERT, DPAGRT
PAGE MOVEMENT WITHIN CENTRAL STORAGE	PGMVRT

Note: The exceptions PAGERT, TPAGRT and DPAGRT apply to the sum of the values in the fields PAGE IN - TOTAL RATE and PAGE OUT - TOTAL RATE.

```
                                        P A G I N G   A C T I V I T Y
                                                                                                                    PAGE    2
          z/OS V1R4              SYSTEM ID   SYS1          DATE 03/16/2002         INTERVAL 60.00.378
                                 RPT VERSION V1R2 RMF      TIME 09.00.00           CYCLE 0.333 SECONDS

OPT = IEAOPT00   MODE = ESA              EXPANDED STORAGE MOVEMENT RATES - IN PAGES PER SECOND
-------------------------------------------------------------------------------------------------------------------------------

    ESF CONFIGURATION                                                                                          HIGH UIC  MIGR AGE
  INSTALLED      ONLINE                                                                                     MIN    12     2020
  ---------      ------                                                                                     MAX   254     5742
    65536        65536                                                                                      AVG    60.6   3209.5

                                          MIGRATED      FREED
                     WRITTEN TO     READ FROM      FROM     WITHOUT      *----- EXPANDED STORAGE FRAME COUNTS -----*
                       EXP STOR      EXP STOR    EXP STOR  MIGRATION          MIN          MAX          AVG

  TOTAL        RT        175.22        150.99       17.84       2.12       59,302       65,529       64,084
   PAGES        %        100.0%        100.0%      100.0%

  HIPERSPACE   RT          0.02          0.00        0.00                      67          140          112
   PAGES        %          0.0%          0.0%        0.0%

  VIO          RT          2.93          1.96        0.03                      18        1,832          169
   PAGES        %          1.7%          1.3%        0.2%

  SHARED       RT          0.00          0.00
   PAGES        %          0.0%          0.0%

                                              FRAME AND SLOT COUNTS
-------------------------------------------------------------------------------------------------------------------------------

                    CENTRAL STORAGE              EXPANDED STORAGE              LOCAL PAGE DATA SET SLOT COUNTS
                  ---------------------------  ---------------------------  ---------------------------------------------
  (1082 SAMPLES)      MIN       MAX       AVG      MIN       MAX       AVG                        MIN         MAX         AVG

  AVAILABLE            42     5,745       841        7     6,234     1,452  AVAILABLE SLOTS     340,407     385,117     358,313

  SQA               2,476     2,553     2,523      194       267       214  VIO SLOTS                 0          90          16
  LPA               1,182     2,011     1,641      256     1,136       640
  CSA               2,584     3,795     3,316      877     2,654     1,449  NON-VIO SLOTS        72,083     116,722      98,870
  LSQA              6,239     8,213     6,876    2,685     4,042     3,759
  REGIONS+SWA      40,425    45,830    44,486   57,390    63,627    61,756  BAD SLOTS                 0           0           0

  TOTAL FRAMES     61,440    61,440    61,440   65,536    65,536    65,536  TOTAL SLOTS         457,200     457,200     457,200

                      FIXED FRAMES                                          SHARED FRAMES AND SLOTS
                  ---------------------------                           ---------------------------------------------

  NUCLEUS           1,754     1,754     1,754                             CENTRAL STORAGE           0           0           0
  SQA               2,389     2,422     2,407                             EXPANDED STORAGE          0           0           0
  LPA                 250       251       250                             FIXED TOTAL               0           0           0
  CSA                 452       495       453                             FIXED BELOW 16 M          0           0           0
  LSQA              5,714     6,668     6,220                             AUXILIARY SLOTS           0           0           0
  REGIONS+SWA       3,529     5,642     3,589
  BELOW 16 M          436     1,202       533                             TOTAL                     0           0           0
  BETWEEN 16M-2G      N/A       N/A       N/A

  TOTAL FRAMES     14,141    16,747    14,674
```

Figure 5-36. PAGING Report - Expanded Storage Movement Rates/Frame and Slot Counts Sections

Expanded Storage Movement Rates section

This section is available only for systems running in 31-bit mode, it includes the expanded storage frame (ESF) configuration, for installed and online storage frames, the system high unreferenced interval count (HIGH UIC), and the migration age (MIGR AGE). It also includes the following for both VIO and HIPERSPACE:

- Pages written to expanded storage
- Pages read from expanded storage
- Migration from expanded storage
- Pages freed without migration
- Expanded storage frame counts

Table 5-45. Fields in the Paging Activity Report - Expanded Storage Movement Rates Section

Field Heading	Meaning
ESF CONFIGURATION	The expanded storage frame configuration consists of two counts: **INSTALLED** Total amount of installed expanded storage frames **ONLINE** Amount of installed expanded storage frames that are actually online to the system
HIGH UIC	The system high unreferenced interval count (HIGH UIC). The HIGH UIC can indicate contention for central storage frames. When the HIGH 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. Because this value varies over the measurement interval, RMF provides minimum, maximum and average HIGH UIC values.
MIGR AGE	The length of time a page remains unreferenced in expanded storage before it migrates to auxiliary storage. A low migration age can indicate an over-use of expanded storage. On the other hand, a high migration age could mean under-use of expanded storage. By comparing the migration age to the criteria age of each category on the criteria table, you can understand which types of pages are being sent to expanded storage and which are not. Because this value varies over the measurement interval, RMF provides minimum, maximum and average migration age values.
WRITTEN TO EXP STOR	**TOTAL PAGES** Total rate (pages per second) sent to expanded storage, and the total percent (100%), including VIO and hiperspace pages. This field indicates how much your system uses expanded storage. **HIPERSPACE PAGES** Rate of hiperspace pages written to expanded storage and percent of total pages that were hiperspace pages sent to expanded storage. **VIO PAGES** Rate and percent of VIO pages sent to expanded storage. **SHARED PAGES** Rate and percent of total pages that were shared pages sent to expanded storage.
READ FROM EXP STOR	**TOTAL PAGES** Total rate and total percent (100%) of pages read in from expanded storage. This field indicates how much your system uses expanded storage. **HIPERSPACE PAGES** Rate and percent of hiperspace pages read from expanded storage. **VIO PAGES** Rate and percent of VIO pages read from expanded storage. **SHARED PAGES** Rate and percent of total pages that were shared storage pages read from expanded storage

Table 5-45. Fields in the Paging Activity Report - Expanded Storage Movement Rates Section (continued)

Field Heading	Meaning
MIGRATED FROM EXP STOR	The rate of pages migrated from expanded storage to auxiliary storage and the hiperspace pages migrated from expanded storage. The migration rate identifies how effectively you are using expanded storage. A low migration rate means that pages that are paged out to expanded storage are referenced quickly enough to avoid migration to an external I/O paging device. Low migration rates mean that the system is avoiding a high paging overhead. A high migration rate, in contrast, might mean contention for limited expanded storage or that the expanded storage criteria table does not match your workload. **TOTAL PAGES** Total rate and the total percent (100%) of pages migrated from expanded storage to auxiliary storage. **HIPERSPACE PAGES** Rate of hiperspace pages migrated from expanded storage to auxiliary storage, and percent of total pages migrated to auxiliary storage that were hiperspace pages. **VIO PAGES** Rate of VIO pages migrated from expanded storage to auxiliary storage, and percent of total pages migrated to auxiliary storage that were VIO pages. When RMF is running on a processor without expanded storage, this field contains zeroes.
FREED WITHOUT MIGRATION	The rate of expanded storage pages freed without being migrated. These are pages freed because the user was finished with the storage.
EXPANDED STORAGE FRAME COUNTS	**TOTAL PAGES** MIN, MAX, and AVG of VIO, HIPERSPACE, and TOTAL allocated expanded storage frame counts. **HIPERSPACE PAGES** Storage frame counts allocated to hiperspace. HIPERSPACE PAGES is a subset of the TOTALS PAGES. **VIO PAGES** Storage frame counts allocated to VIO address space in expanded storage. VIO PAGES is a subset of the TOTALS PAGES.

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 5-46. Spreadsheet Range Names and Overview Names in the Paging Activity Report - Expanded Storage Movement Rates Section

Column Heading	Overview Name
HIGH UIC - MAX	MXHUIC
HIGH UIC - AVG	AVGHUIC

Central Storage Movement Rates section

This section is available for systems running in 64-bit mode.

```
                                        P A G I N G   A C T I V I T Y
                                                                                                                    PAGE    2
          z/OS V1R4               SYSTEM ID   ZB13          DATE 12/16/2002          INTERVAL 30.00.378
                                  RPT VERSION V1R2 RMF       TIME 08.00.00            CYCLE 1.000 SECONDS

OPT = IEAOPT00   MODE = ESAME            CENTRAL STORAGE MOVEMENT RATES - IN PAGES PER SECOND
-------------------------------------------------------------------------------------------------------------------------------

   HIGH UIC (AVG) = 2498.3   (MAX) = 2540   (MIN) = 2133


                     WRITTEN TO       READ FROM        *--- CENTRAL STORAGE FRAME COUNTS ----*
                   CENTRAL STOR    CENTRAL STOR             MIN           MAX           AVG

   HIPERSPACE   RT          0.02            0.00             67           140           112
   PAGES

   VIO          RT          2.93            1.96             18         1,832           169
   PAGES



                                                FRAME AND SLOT COUNTS
-------------------------------------------------------------------------------------------------------------------------------

                      CENTRAL STORAGE                                                LOCAL PAGE DATA SET SLOT COUNTS
                   ---------------------------                                    ---------------------------------------------
   (1801 SAMPLES)      MIN       MAX       AVG                                                      MIN         MAX         AVG

   AVAILABLE       712,579   716,956   715,263                    AVAILABLE SLOTS     610,950     610,950     610,950

   SQA               8,697     8,715     8,709                    VIO SLOTS                 0           0           0
   LPA               4,140     4,141     4,141
   CSA               7,094     7,118     7,110                    NON-VIO SLOTS         1,050       1,050       1,050
   LSQA             27,216    27,234    27,225
   REGIONS+SWA     274,340   278,716   276,033                    BAD SLOTS                 0           0           0

   TOTAL FRAMES    1040362   1040362   1040362                    TOTAL SLOTS         612,000     612,000     612,000

                       FIXED FRAMES                                          SHARED FRAMES AND SLOTS
                   ---------------------------                    ---------------------------------------------
   NUCLEUS           1,880     1,880     1,880                    CENTRAL STORAGE           0           0           0
   SQA               7,485     7,503     7,496
   LPA                  65        65        65                    FIXED TOTAL               0           0           0
   CSA               1,374     1,374     1,374                    FIXED BELOW 16 M          0           0           0
   LSQA             14,337    14,350    14,340                    AUXILIARY SLOTS           0           0           0
   REGIONS+SWA      72,114    72,815    72,298
   BELOW 16 MEG         53        61        53                    TOTAL                     0           0           0
   BETWEEN 16M-2G   21,408    22,154    21,607

   TOTAL FRAMES     97,266    97,973    97,454
```

Figure 5-37. PAGING Report - Central Storage Movement Rates/Frame and Slot Counts Sections

Table 5-47. Fields in the Paging Activity Report - Central Storage Movement Rates Section

Field Heading	Meaning
HIGH UIC	The system high unreferenced interval count (HIGH UIC). The HIGH UIC can indicate contention for central storage frames. When the HIGH 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. In contrary to systems in 31-bit mode with a maximal value of 254, the UIC can have a maximal value of 2540.
WRITTEN TO REAL STOR	**HIPERSPACE PAGES** Rate of hiperspace pages written to central storage. **VIO PAGES** Rate of VIO pages written to central storage.

Table 5-47. Fields in the Paging Activity Report - Central Storage Movement Rates Section (continued)

Field Heading	Meaning
READ FROM REAL STOR	**HIPERSPACE PAGES** Rate of hiperspace pages read from to central storage. **VIO PAGES** Rate of VIO pages read from to central storage.
REAL STORAGE FRAME COUNTS	MIN, MAX, and AVG of allocated frame counts. **HIPERSPACE PAGES** Storage frame counts allocated to hiperspace. **VIO PAGES** Storage frame counts allocated to VIO address space.

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 5-48. Overview Names in the Paging Activity Report - Central Storage Movement Rates Section

Column Heading	Overview Name
HIGH UIC - MAX	MXHUIC
HIGH UIC - AVG	AVGHUIC
HIPERSPACE WRITTEN TO REAL STOR	RSHSPW
HIPERSPACE READ FROM REAL STOR	RSHSPR
VIO WRITTEN TO REAL STOR	RSVIOW
VIO READ FROM REAL STOR	RSVIOR
HIPERSPACE REAL STORAGE FRAME COUNTS - MIN	RSHSPM
HIPERSPACE REAL STORAGE FRAME COUNTS - MAX	RSHSPX
HIPERSPACE REAL STORAGE FRAME COUNTS - AVG	RSHSPA
VIO REAL STORAGE FRAME COUNTS - MIN	RSVIOM
VIO REAL STORAGE FRAME COUNTS - MAX	RSVIOX
VIO REAL STORAGE FRAME COUNTS - AVG	RSVIOA

Frame and Slot Counts section

This section of the Paging Activity report (see Figure 5-36 on page 5-113 and Figure 5-37 on page 5-116) shows the total frames and fixed frames for central storage and expanded storage. The number of samples for this interval is located in the top left corner of this section.

The total frames category is broken down into SQA, LPA, CSA, LSQA, REGIONS+SWA, and AVAILABLE. The fixed frames category is broken down into Nucleus, SQA, LPA, CSA, LSQA, REGIONS+SWA, BELOW 16 M and BETWEEN 16M-2G. The row BETWEEN 16M-2G shows values only if the system is running in ESAME mode (see Figure 5-37 on page 5-116). If it is running in ESA mode, it shows 'N/A' (see Figure 5-36 on page 5-113).

The shared storage category is broken down into central storage, expanded storage, fixed total, fixed below 16M, and auxiliary slots. All of these values are presented as MIN (minimum), MAX (maximum), and AVG (average).

This section also shows counts for available slots, VIO slots, non-VIO slots, bad slots, and total slots. These values are presented as MIN (minimum), MAX (maximum), and AVG (average).

For systems running in 64-bit mode, all fields related to expanded storage will not be shown.

Table 5-49. Fields in the Paging Activity Report - Frame and Slots Counts Section

Field Heading	Meaning
SAMPLES	The number of valid samples taken in this interval is shown in the upper left corner of this report section.
TOTAL FRAMES	This particular category of central and expanded storage is broken down into SQA, LPA, CSA, LSQA, REGIONS + SWA and AVAILABLE. The value for TOTAL FRAMES is not a summation of the previous six categories at the end of the interval, but is derived by adding the six categories during each valid sample and later reporting the MIN sum, MAX sum, and AVG sum for the set of samples. The TOTAL FRAMES count for central storage includes frames both above and below 16 megabytes and nucleus frames. When data is not available to RMF in a particular category, the following message will appear to the right of the category name: `*** NO COUNTS AVAILABLE ***` **Notes:** 1. The actual maximum or minimum value of a field might occur at a time when RMF is not sampling. 2. In rare cases the `AVAILABLE` field can contain a negative number. The negative number indicates a storage shortage situation, that is, the number of available frames is almost exhausted and ″reserved″ frames are already being used.
FIXED FRAMES	This category is broken down into NUCLEUS, SQA, LPA, CSA, LSQA, REGIONS + SWA, BELOW 16 M, and BETWEEN 16M-2G. The row BETWEEN 16M-2G shows values only if the system is running in ESAME mode (see Figure 5-37 on page 5-116). If it is running in ESA mode, it shows 'N/A' (see Figure 5-36 on page 5-113). The value for TOTAL FRAMES is not a summation of the previous seven categories at the end of the interval but is derived by adding the categories during each valid sample and reporting the MIN sum, MAX sum, and AVG sum for the set of samples. Nucleus frames are included in the FIXED FRAMES count. When data is not available to RMF in a particular category, the following message will appear to the right of the category name: `*** NO COUNTS AVAILABLE ***`
AVAILABLE SLOTS	The number of page slots that do not contain any data pages and that are available for use.
VIO SLOTS	The number of local page data set slots that contain pages for VIO data sets.
NON-VIO SLOTS	The number of local page data set slots that contain pages belonging to address-space virtual storage.

Table 5-49. Fields in the Paging Activity Report - Frame and Slots Counts Section (continued)

Field Heading	Meaning
BAD SLOTS	The 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.
TOTAL SLOTS	This value is the MIN, MAX, and AVG sum of the four categories above for the set of samples, not for the interval.
SHARED FRAMES AND SLOTS	This category is broken down into CENTRAL STORAGE, EXPANDED STORAGE, FIXED TOTAL, FIXED BELOW 16M, and AUXILIARY slots.

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 5-50. Overview Names in the Paging Activity Report - Frame and Slot Counts Section

Column Heading	Overview Name
CENTRAL-STORAGE - MAX	STORAGE, MXSQA, MXCSAT, MAXVIOF
CENTRAL-STORAGE - AVG	AVGSQA, AVGCSAT, AVGVIOF
FIXED FRAMES - CSA - AVG	AVGCSAF
FIXED FRAMES - CSA - MAX	MXCSAF
FIXED FRAMES - BETWEEN 16M-2G - MIN	FXBETWM
FIXED FRAMES - BETWEEN 16M-2G - AVG	FXBETWA
FIXED FRAMES - BETWEEN 16M-2G - MAX	FXBETWX

Swap Placement Activity section

```
                                        P A G I N G   A C T I V I T Y
                                                                                                        PAGE    3
          z/OS V1R4                SYSTEM ID   SYS1           DATE 03/16/2002          INTERVAL 60.00.378
                                   RPT VERSION V1R2 RMF       TIME 09.00.00            CYCLE 0.333 SECONDS

OPT = IEAOPT00   MODE = ESA                S W A P   P L A C E M E N T   A C T I V I T Y
                         *------ AUX STORAGE -------* *---LOGICAL SWAP--*  *---------------- EXPANDED STORAGE ----------------*

                                  AUX              AUX STOR                                                MIGRATED              LOG SWAP
                                  STOR     AUX STOR    VIA                LOG SWAP   EXP STOR   EXP STOR    FROM     EXP STOR   /EXP STOR
                        TOTAL     TOTAL    DIRECT   TRANSITION LOG SWAP  EFFECTIVE  DIRECT      TOTAL    EXP STOR  EFFECTIVE  EFFECTIVE

TERMINAL      CT     16,652        33         0         33     16,652     16,426         0        220          27        193     16,619
INPUT/OUTPUT  RT       4.63      0.01      0.00       0.01       4.63       4.56      0.00       0.06        0.01       0.05       4.62
WAIT           %      75.4%      0.2%      0.0%     100.0%     100.0%      98.6%      0.0%       1.3%       12.3%      87.7%      99.8%

LONG          CT      3,832         7         0          7      3,832      3,816         0         13           4          9      3,825
WAIT          RT       1.06      0.00      0.00       0.00       1.06       1.06      0.00       0.00        0.00       0.00       1.06
               %      17.4%      0.2%      0.0%     100.0%     100.0%      99.6%      0.0%       0.3%       30.8%      69.2%      99.8%

DETECTED      CT      1,539         1         0          1      1,539      1,507         0         31           0         31      1,538
WAIT          RT       0.43      0.00      0.00       0.00       0.43       0.42      0.00       0.01        0.00       0.01       0.43
               %       7.0%      0.1%      0.0%     100.0%     100.0%      97.9%      0.0%       2.0%        0.0%     100.0%      99.9%

UNILATERAL    CT         38         0         0          0         38         28         0         10           0         10         38
              RT       0.01      0.00      0.00       0.00       0.01       0.01      0.00       0.00        0.00       0.00       0.01
               %       0.2%      0.0%      0.0%       0.0%     100.0%      73.7%      0.0%      26.3%        0.0%     100.0%     100.0%

TRANSITION    CT         11         1         0          1          1          0        10         10           0         10         10
TO NON-       RT       0.00      0.00      0.00       0.00       0.00       0.00      0.00       0.00        0.00       0.00       0.00
SWAPPABLE      %       0.0%      9.1%      0.0%     100.0%       9.1%       0.0%     90.9%      90.9%        0.0%     100.0%      90.9%

TOTAL         CT     22,072        42         0         42     22,062     21,777        10        284          31        253     22,030
              RT       6.13      0.01      0.00       0.01       6.13       6.05      0.00       0.08        0.01       0.07       6.12
               %     100.0%      0.2%      0.0%     100.0%     100.0%      98.7%      0.0%       1.3%       10.9%      89.1%      99.8%


AUXILIARY STORAGE - AVERAGE PAGES PER SWAP OUT -        32       AVERAGE PAGES PER SWAP IN -        164
OCCURRENCES OF TERMINAL OUTPUT WAIT -     1,406
```

Figure 5-38. PAGING Report - Swap Placement Section

This section of the Paging Activity report (page 3 of the sample report) provides a full analysis of swap placement. This section reports data for physical swapping to auxiliary storage, logical swapping and, physical swapping to expanded storage. The top of the page provides the current OPT member, which contains the current expanded storage criteria table used to tune the use of expanded storage. A line at the bottom of the report indicates the average number of pages per swap-out and swap-in, a second line indicates the number of terminal output wait occurrences.

The section identifies, by reason, the total swap occurrences (logical swaps and physical swaps) and gives an analysis of the resulting swap transitions. The initial swap decisions that involve expanded storage are:

- **Logical swapping in central storage** - LOGICAL SWAP

  Of the total number of swap candidates, some may be logically swapped in central storage. At some later time, some of these may be physically swapped to expanded storage or physically swapped to auxiliary storage. Of those physically swapped to expanded storage, some may be migrated to auxiliary storage at a later time.

- **Physical swapping to expanded storage** - EXPANDED STORAGE

  Of the total number of logical swap candidates, some may be physically swapped directly to expanded storage. Of these, some may be migrated to auxiliary storage at a later time.

  For systems running in 64-bit mode, this block of data will be shown as N/A.
- **Physical swapping to auxiliary storage** - AUX STORAGE

  Of the total number of logical swap candidates, some may be placed directly on auxiliary storage.

The entries in the left column of the Swap Placement Activity section identify the swap reason. For each swap reason, the report provides the count (CT), rate (RT), and percentage (%).

Table 5-51 lists the **common swap reasons** that may occur in the Swap Placement Activity section of the Paging Activity report.

Table 5-51. Common Swap Reasons in the Paging Activity Report

Swap Reason	Definition
TERMINAL INPUT/OUTPUT WAIT	The number of swap-outs that occurred because a terminal was waiting for input or was waiting for output buffers. The address space was not a candidate for logical swapping. Although TERMINAL INPUT/OUTPUT WAIT is identified as a single reason, SRM distinguishes INPUT/OUTPUT as two separate swap reasons, each having its unique set of counters. To derive the fields for TERMINAL INPUT/OUTPUT WAIT, RMF always uses the sum of both counters in its calculation.
LONG WAIT	The number of swap-outs that occurred because of user-requested long waits. Swap-outs of address spaces that were swapped in, executed, and entered a wait state before SRM received control are also counted.
DETECTED WAIT	The number of swap-outs that occurred because of user address spaces that have not issued a WAIT,LONG=YES macro, but they have gone without executing for more than the SRM-specified wait threshold or 2 seconds, whichever is greater. The SRM determines the wait threshold by dividing 8 seconds by the CPU-related SRM second/real time seconds.
UNILATERAL	The number of swap-outs that occurred because the multiprogramming level for a domain became higher than the target multiprogramming level specified for that domain.
EXCHANGE ON RECOMMENDATION VALUE	The number of swap-outs that occurred because one user in a domain must be swapped out to allow another user (with a higher workload level) in the same domain to be swapped in.

Table 5-52 lists the **uncommon swap reasons** that may occur in the Swap Placement Activity section of the Paging Activity report. These swaps rarely occur, but it is important to be aware of them.

Table 5-52. Uncommon Swap Reasons in the Paging Activity Report

Swap Reason	Definition
ENQUEUE EXCHANGE	The number of swap-outs that occurred to make room for users enqueued upon a system resource that is required by other users.

Table 5-52. Uncommon Swap Reasons in the Paging Activity Report (continued)

Swap Reason	Definition
REQUESTED	The number of swap-outs that occurred because the program properties table (PPT) indicated that a job or step is non-swappable. The address space is swapped out, and its LSQA is swapped into non-reconfigurable storage. This field also includes the number of swap outs caused by a CONFIG STOR,OFFLINE command.
AUXILIARY STORAGE SHORTAGE	The number of swap-outs that occurred because 70% of all local page slots are allocated. The swappable user address spaces that acquired the greatest rate of auxiliary storage page slots are swapped out.
CENTRAL PAGEABLE STORAGE SHORTAGE	The number of swap-outs that occurred because of a shortage of central pageable frames. The swappable users that have acquired the most fixed pages are swapped out.
TRANSITION TO NONSWAPPABLE	The number of swap-outs that occurred because an address space is made non-swappable. When the address space is made non-swappable, bits are set to ensure that it will use the preferred area and not the reconfigurable area, after it is made non-swappable.
APPC WAIT	The number of swap-outs that occurred because a job requesting work by an APPC/MVS service must be swapped out.
IMPROVE CENTRAL STORAGE	The number of swap-outs that occurred when SRM determined that a change in address space use would improve system paging and swap costs. For SRM to make this decision, the system must be running within thresholds specified to control the multi-programming level.
IMPROVE SYSTEM PAGING	The number of swap-outs that occurred when the system page fault rate exceeds the threshold.
MAKE ROOM FOR OUT-TOO-LONG USERS	The number of swap-outs that occurred to prepare for the swap-in of an address space that has been kept out longer than its recommendation value would dictate.
OMVS INPUT	The number of swap-outs that occurred due to an OMVS input wait.
OMVS OUTPUT	The number of swap-outs that occurred due to an OMVS output wait.
IN-REAL SWAP	The number of swap-outs that occurred due to recovery of real storage frames from a swappable address space.
Note: If no swaps occurred during the measurement interval for a particular swap reason, the report omits that swap reason and its associated counts, rates, and percentages. For example, few, if any, swaps usually occur for the following swap reasons: • Requested swap • Auxiliary storage shortage • Central pageable storage shortage • Transition to non-swappable state	

Table 5-53 on page 5-123 gives detailed descriptions of the fields in the Swap Placement Activity section of the Paging Activity report.

Column headings (from left to right in the report) describe how SRM handles swap requests and how effective SRM is satisfying these requests. If expanded storage is not available or not used, the rightmost five column headings still appear, but data fields for count, rate, and percentage are blank.

The “TOTAL” row on the left side of the report presents the totals for each field in the report. Unless otherwise noted, percent data is always a percentage of the corresponding TOTAL.

Table 5-53. Fields in the Paging Activity Report - Swap Placement Activity Section

Field Heading	Meaning
TOTAL	The swap requests for this reason in relation to the total number of swap requests. For each reason, it describes the total number of swap requests, the rate per second, and the percentage of all swaps that occurred for that reason.
AUX STOR TOTAL	The total number of swap requests directly or indirectly sent to auxiliary storage.
AUX STOR DIRECT	The swap requests for this reason in relation to the total number of swap requests that were sent to auxiliary storage. Percent data is a percentage of the AUX STOR TOTAL.
AUX STOR VIA TRANSITION	The swap requests for this reason in relation to the total number of swap requests eventually sent to auxiliary storage. These were either logically swapped from central storage to auxiliary storage or swapped from central storage to expanded storage. Percent data is a percentage of the AUX STOR TOTAL.
LOG SWAP	The swap requests that were logically swapped.
LOG SWAP EFFECTIVE	The swap requests that were logically swapped and never physically swapped. Percent data is a percentage of LOG SWAP.
EXP STOR DIRECT	The logical swap requests sent to expanded storage.
EXP STOR TOTAL	The total of all swap requests sent to expanded storage. These include those swaps sent to expanded storage and those logically swapped that were physically swapped and sent to expanded storage.
MIGRATED FROM EXP STOR	Out of the total number of swap requests sent to expanded storage, those eventually moved to auxiliary storage. The percent field is a percentage of EXP STOR TOTAL.
EXP STOR EFFECTIVE	The swap requests sent to expanded storage and not moved to auxiliary storage. The percent field is a percentage of EXP STOR TOTAL.
LOG SWAP/EXP STOR EFFECTIVE	The swap requests that either remained logically swapped in central storage or, after being sent to expanded storage, remained there and were not moved. This value indicates that I/O to auxiliary storage was avoided when the swap request was satisfied. Logically, this value is the sum of LOG SWAP EFFECTIVE and EXP STOR EFFECTIVE.
AUXILIARY STORAGE - AVERAGE PAGES PER SWAP OUT	The average number of pages swapped out for each storage swap-out. The range of values is 0 to 4,096 pages or 0 to 16,777,216 bytes. This value is presented at the bottom of the report.
AUXILIARY STORAGE - AVERAGE PAGES PER SWAP IN	The average number of pages swapped in for each storage swap-in. This value is presented at the bottom of the report.
OCCURRENCES OF TERMINAL OUTPUT WAIT	The number of swap outs while in terminal output wait. There are cases where an address space is not physically swapped, but is being moved on and off the logical swap queue while in terminal wait. The internal SRM logic updates the swapout counts on each successful move. Since the job is still in a terminal output wait its swap reason remains a terminal output wait.

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 5-54. Overview Names in the Paging Activity Report - Swap Placement Activity Section

Column Heading	Overview Name	
TOTAL	**SWART**	Total swaps
	Swaps due to:	
	PSOCLW	Long wait
	PSOCDW	Detected wait
	PSOCU	Unilateral
	PSOCEORV	Recommendation
	PSOCENQE	Enqueue exchange
	PSOCREQ	Requested
	PSOCAXSS	Auxiliary storage shortage
	PSOCRPSS	Central storage shortage
	PSOCTONS	Transition to non-swap
	PSOCTW	Terminal wait

SDEVICE - Shared Device Activity report

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))`

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.

Purpose of the 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 MVS 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.

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 (SYSF and SYSO). Only three devices are shown.

The third device does not have the same device number on all two systems.

```
                S H A R E D   D I R E C T   A C C E S S   D E V I C E   A C T I V I T Y

      z/OS V1R4              SYSPLEX RMFLP4X           DATE 04/25/2002         INTERVAL 15.56.594
                             RPT VERSION V1R2 RMF      TIME 16.15.14           CYCLE 01.000  SECONDS


   TOTAL SAMPLES(AVG) =  900.0  (MAX) =  900.0  (MIN) =  900.0


                         SMF           DEVICE    AVG  AVG   AVG  AVG  AVG  AVG  AVG    %      %     %     AVG
DEV  DEVICE   VOLUME PAV SYS  IODF LCU ACTIVITY  RESP IOSQ  CMR  DB   PEND DISC CONN   DEV    DEV   DEV   NUMBER
NUM  TYPE     SERIAL     ID   SUFF     RATE      TIME TIME  DLY  DLY  TIME TIME TIME   CONN   UTIL  RESV  ALLOC

0401 3380K    SYSLIB     *ALL           10.869   3.0  0.1 10.0  0.7   0.9  0.2  1.8   1.97   2.16   0.2   125
                         SYSF  01  0032    1.098   6.2  0.0  0.0  3.1   3.4  0.4  2.4   0.26   0.31   0.0  36.0
                         SYS1  01  0032    9.771   2.6  0.1  0.0  0.4   0.6  0.1  1.8   1.71   1.85   0.2  88.5

0402 3380K    SYSUSR     *ALL             0.607   1.9  0.0  0.0  0.2   0.8  0.0  1.1   0.07   0.07   0.0  15.0
                         SYSF  01  0032    0.201   2.8  0.0  0.0  0.5   1.8  0.0  1.0   0.02   0.02   0.0   1.0
                         SYS1  01  0032    0.406   1.5  0.0  0.0  0.0   0.3  0.0  1.1   0.05   0.05   0.0  14.0

0460 3380K    RMFLIB     *ALL             0.071   7.6  0.0  0.0  0.8   3.0  0.0  4.6   0.03   0.03   0.0   0.0
0460                     SYSF  01  0033    0.036   7.0  0.0  0.0  0.8   2.4  0.0  4.6   0.02   0.02   0.0   0.0
0461                     SYS1  01  0033    0.036   8.3  0.0  0.0  0.8   3.7  0.0  4.6   0.02   0.02   0.0   0.0
```

Figure 5-39. Shared DASD Activity Report

For the field descriptions, see Table 5-55 on page 5-127.

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.

```
                S H A R E D   M A G N E T I C   T A P E   D E V I C E   A C T I V I T Y
                                                                                          PAGE   1
    z/OS V1R4              SYSPLEX PLEXPERF          DATE 04/25/2002          INTERVAL 30.00.001
                           RPT VERSION V1R2 RMF      TIME 13.39.00            CYCLE 01.000 SECONDS


   TOTAL SAMPLES(AVG) = 1800.0  (MAX) = 1800.0  (MIN) = 1800.0


                     SMF           DEVICE   AVG  AVG   AVG   AVG  AVG  AVG  AVG    %      %     %     NUMBER    AVG    TIME
DEV  DEVICE  VOLUME SYS  IODF LCU ACTIVITY RESP IOSQ  CMR   DB   PEND DISC CONN   DEV    DEV   DEV   OF        MOUNT  DEVICE
NUM  TYPE    SERIAL ID   SUFF     RATE     TIME TIME  DLY   DLY  TIME TIME TIME   CONN   UTIL  RESV  MOUNTS    TIME   ALLOC


05D0 3490    TAP508 *ALL           9.511  43.1  1.0 11.7   2.0  27.5  7.3  7.0   6.64  13.60   0.0      4       19   57:37
             TAP508 P40   A1  0006  0.012   452  227 39.0   131 217.7  0.1  6.6   0.01   0.01   0.0      0        0   13:37
                    P50   A1  0006  9.474  41.2  0.1 11.5   1.0  26.3  7.3  7.0   6.61  13.58   0.0      4       19   12:01
             TAP508 P60   A1  0006  0.012   153 45.2 33.0  29.9 101.5  0.3  5.8   0.01   0.01   0.0      0        0   15:39
             TAP508 P70   A1  0006  0.012  1130  409 78.8   560 714.7  0.1  6.0   0.01   0.01   0.0      0        0   16:18

05D5 3490    TAP606 *ALL           9.829  54.4  3.1 14.1   4.0  34.5  9.9  6.5   6.39  16.17   0.0      8       28 1:04:04
                    P40   A1  0006  0.012  1347  636 31.8   640 702.6  0.1  8.0   0.01   0.01   0.0      0        0   16:27
                    P50   A1  0006  1.109  37.2  3.2  8.1   7.9  25.9  1.9  5.8   0.65   0.86   0.0      1     1:05   17:17
                    P60   A1  0006  8.696  51.1  0.0 14.7   1.2  33.0 11.0  6.6   5.73  15.29   0.0  *   7       22   16:11
             TAP606 P70   A1  0006  0.012  2144  955 60.8  1061  1182  0.1  7.6   0.01   0.01   0.0      0        0   14:07
```

Figure 5-40. Shared Magnetic Tape Device Activity Report

Field descriptions

The table Table 5-55 contains all report fields and the description how the values shown are calculated.

Table 5-55. Fields in the Shared Device Activity Reports

Field Heading	Meaning
SYSPLEX	Sysplex name
DATE	This is the earliest date found in all records used to process this report
INTERVAL	This is the longest interval which can be built by using integer multiple interval lengths of all SMF record interval lengths. **Note:** This value is referenced below as the "Common Interval Length" (abbreviated as **CIL**). The interval length of each system is referenced as the "System Interval Length" (abbreviated as **SIL**). 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.
CYCLE	This is the cycle value found in the first SMF record used to process this report. **Note:** If the CYCLE values of all SMF records used to build the report do not match, no report is generated.
TOTAL SAMPLES	This field shows the total number of samples used to build this report. **Note:** 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.
DEV NUM	The four-digit hexadecimal device number of a physical I/O device. **Note:** 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.
DEVICE TYPE	The device type of the volume.
VOLUME SERIAL	The volume serial number of the volume mounted on the device at the end of the reporting interval. **Note:** In the Shared Direct Access Device report, this field is printed only on the summary line. The system line field is blank. 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.
PAV	The number of parallel access volumes (base and alias) which were active at the end of the reporting interval. If the number has been changed during the reporting interval, it is followed by an '*'.
SMF SYS ID	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. **Note:** Data is also considered to be partial, if the interval length of that system differs from the interval shown in the report header.

Table 5-55. Fields in the Shared Device Activity Reports (continued)

Field Heading	Meaning
IODF SUFF	The IODF suffix in effect for this system **Note:** In the single-system Device Activity report, the header contains the complete IODF name and suffix. Here, only the suffix is presented.
LCU	The number of the logical control unit to which the device belongs. On the system summary line, this field is blank. **Note:** The LCU number for the same device can differ between the different systems even if the device is being shared.

DEVICE ACTIVITY RATE	The rate at which start subchannel (SSCH) instructions to the device completed successfully.
	
	The calculation for the single system line is:
	`SSCH Count`
	`DEVICE ACTIVITY RATE = -----------`
	`CIL`
	
	In the summary line, this field contains the sum of the rates for each single system.
	Note: For multi-exposure devices the field reflects the value of the entire device. This is true also in all following fields.

AVG RESP TIME	The average number of milliseconds the device required to complete an I/O request.
	
	The average response time consist of two parts, the average service time and the average IOS queue time.
	
	For the single system line this is:
	`AVG RESP TIME = Avg IOSQT`
	`+ AVG PEND TIME`
	`+ AVG DISC TIME`
	`+ AVG CONN TIME`
	
	In the summary line, the same formula is used using the corresponding summary line fields.

Table 5-55. Fields in the Shared Device Activity Reports (continued)

Field Heading	Meaning

AVG IOSQ TIME	The average number of milliseconds an I/O request must wait on an IOS queue before a SSCH instruction can be issued.
	
	Using the abbreviation for "Device Activity Rate (**DAR**)", the calculation for the single system line is:
	`IOS Queue Count`
	`AVG IOSQ TIME= ---------------`
	`STS * DAR`
	
	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 (**MEC**). The weighting factor for System SYSi is therefore, assuming we have n systems:
	`MEC(SYSi)`
	`Wi = ----------------------------------`
	`MEC(SYS1) + MEC(SYS2) + ... MEC(SYSn)`
	
	The weighted AVG IOS queue time is therefore calculated as:
	
	`WEIGHTED AVG IOS QT = (AVG IOS QT(SYS1)) * W1 +`
	`(AVG IOS QT(SYS2)) * W2 +`
	`...`
	`(AVG IOS QT(SYSn)) * Wn`
	
	Note: This weighting algorithm is different to the LCU summary line algorithm used in the single-system Device Activity report.
	
	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.

AVG CMR DLY	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.
	`Initial Command Response Time`
	`AVG CMR DLY = ----------------------------`
	`MEC`

AVG DB DELAY	The average number of milliseconds of delay that I/O requests to this device encountered because the device was busy.
	
	The calculation for the single system line is:
	`Total DB DLY Time`
	`AVG DB DLY = -----------------`
	`MEC`
	
	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.

Table 5-55. Fields in the Shared Device Activity Reports (continued)

Field Heading	Meaning

AVG PEND TIME	The average number of milliseconds an I/O request must wait in the hardware.
	
	The calculation for the single system line is:
	`Total Pending Time`
	`AVG PEND TIME= ------------------`
	`MEC`
	
	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.

AVG DISC TIME	The average number of milliseconds the device was disconnected while processing an SSCH instruction.
	
	The calculation for the single system line is:
	`Total Disconnect Time`
	`AVG DISC TIME = ----------------------`
	`MEC`
	
	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.

AVG CONN TIME	The average number of milliseconds the device was connected to a channel path and actually transferring data between the device and central storage.
	
	The calculation for the single system line is:
	`Total Connect Time`
	`AVG CONN TIME = ------------------`
	`MEC`
	
	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.

% DEV CONN	The percentage of time during the interval when the device was connected to a channel path.
	
	The calculation for the single system line is:
	`Device Connect Time`
	`% DEV CONN = 100 * -------------------`
	`CIL`
	
	In the summary line, this field contains the sum of the single system % DEV CONN values.

Table 5-55. Fields in the Shared Device Activity Reports (continued)

Field Heading	Meaning

% DEV UTIL	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.
	
	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:
	
	`(Device Connect Time +`
	`% DEV  Disconnect Time)`
	`UTIL = 100 * ----------------------  +`
	`CIL`
	
	`UTL * SIL`
	`100 * -----------`
	`STS* CIL`
	
	In the summary line, this field contains the sum of the single system % DEV UTIL values.

% DEV RESV	The percentage of time during the interval when a shared device was reserved by the system on which RMF was started.
	
	Using the abbreviations for "number reserved samples (**DRP**)", and "system total samples (**STS**)", the calculation used for the single system line is:
	
	`DRP * SIL`
	`% DEV RESV = 100 * -----------`
	`STS * CIL`
	
	This is the percent of the common interval length time which this system holds an reserve on this device.
	
	In the summary line, this field contains the sum of the single system % DEV RESV values.

AVG NUMBER ALLOC	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 the shared direct access storage devices.
	
	The calculation used for the single system line is:
	
	`AVG  Num Alloc * System Interval Length`
	`NUMBER = ----------------------------------`
	`ALLOC  STS * Common Interval Length`
	
	In the summary line, this field contains the sum of the single system AVG NUMBER ALLOC values.

Table 5-55. Fields in the Shared Device Activity Reports (continued)

Field Heading	Meaning
NUMBER OF MOUNTS	The number of tape mounts, shown as an integer value, detected by RMF. This field is reported only for magnetic tape devices. 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. If a mount-pending condition is detected at the first cycle of the interval, the mount count for the interval will be incremented by one. At the single system line the value is displayed as a integer value allowing a range between 0 and 9999. 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.

AVG MOUNT TIME	The average mount time pending for every device, expressed in the form of HH:MM:SS.
	
	Using the abbreviations "Mount Pending Samples (**MTP**)", and "Mount Total Count (**MTC**)", the calculation used for the single system line is:
	`AVG MOUNT TIME = MTP * SIL / STS * MTC`
	
	Note: The mount time value has units of full seconds, while the interval length usually is shown in units of milliseconds.
	
	If the mount count or the sample count is zero, the result is zero.
	
	In the summary line, the average mount time is calculated as:
	`AVG MOUNT TIME = (AVG Mount Time(SYS1)* MTC(SYS1) + ...) / (∑ All System MTC Counts)`
	
	This field is reported only for magnetic tape devices.

TIME DEVICE ALLOC	The total time the device was allocated during the interval, expressed in the form of HH:MM:SS. This field is reported only for magnetic tape devices. The calculation is: `TIME DEVICE ALLOC = ALC * SIL / STS` If the system total sample count is zero, the result is zero. In the summary line, the field contains the sum of the system line values. **Note:** The mount time value has units of full seconds, while the interval length in the report header is shown in units of milliseconds.

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 5-56. Overview Names in the Shared DASD Activity Report

Criterion	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

TRACE - Trace Activity report

How to request this report

To gather data for this report, specify as a Monitor I gatherer option

```
TRACE(variable [,options list])
```

To produce the report, specify

```
REPORTS(TRACE)
```

Purpose of the 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 5-58 on page 5-136.

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.

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, RCVPTR, DMDTRUA, and DMDTCMPL.

In addition, you can monitor the use of expanded storage by tracing the variables RCEESINU, RCEESWRT,RCENWSF and RCEWSDNE, MCVMGAGE and MCVSTCRI.

By tracing the domain variables, for a system running in compatibility mode, you can determine whether the domain setting in the IPS is appropriate. You can trace the multiprogramming target level over the interval and compare it with the current multiprogramming level that the system is aiming at for each domain.

Note: Domain tracing ends on a SET IPS operator command and will not restart unless the session options are modified to re-specify domain tracing, or the Monitor I session is stopped and restarted. When this occurs, **ENDED** appears on the report for the domain.

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.

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 X 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.

Domain table tracing in compatibility mode

Tracing the domain variables causes the domain numbers to be appended to an abbreviation of the variable on the report:

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

Each variable specified is traced and reported for all domains that were active when RMF was initialized. For instance, if domains 0, 1, 46, and 245 are active when RMF starts, and DMDTINCU is requested, the following names are presented for the options requested:

DMINU000
DMINU001
DMINU046
DMINU245

```
                                        T R A C E  A C T I V I T Y
                                                                                                    PAGE 1
          z/OS V1R4               SYSTEM ID   SYS1           DATE 03/14/2002          INTERVAL 14.59.946
                                  RPT VERSION V1R2 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 * RCVPAGRT * RCVPTR   * 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 5-41. TRACE Report

Table 5-57. 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 5-58. 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 CPUs

Table 5-58. Variables in the Trace Activity Report (continued)

Variable	Value
CCVENQCT	Number of users non-swappable for enqueue reasons
CCVRBSTD	Recent base time of day
CCVRBSWT	Recent base system wait time
CCVUTILP	System CPU utilization
CCVWTDB	Number of wait-time dispatches when the LPAR is the guest
DMDTCIDX	Domain contention index * 100
DMDTCMPL	Current multiprogramming level
DMDTINCU	Users per domain in storage
DMDTOUTU	Users per domain out of storage
DMDTRUA	Average number of ready users * 16
DMDTTWET	Average elapsed time in 1024 microsecond units for first period TSO transaction (does not include delay imposed by SRM to meet response time objective)
DMDTTWSR	Time weighted service rate for a domain
Note: The domain trace variables (DMDT*xxxx*) specified are different from the variables reported on the Trace Activity report. See "Domain table tracing in compatibility mode" on page 5-135.	
LSCTCNT	Current number of logically swapped users for terminal wait
LSCTMTE	Maximum think time allowed for logical swap candidate
MCVFRCNT	Number of pages needed to be stolen by force steal routine
MCVMGAGE	Expanded storage migration age
MCVSBLTF	Long term percentage of eligible storage that is actually fixed
MCVSIPR	Common page-in rate
MCVSTCRI	Highest system UIC
MCVTWSS	Common target working set size
OMDGAMRE	Maximum number of messages on the action message retention facility (AMRF) queue. If a large number of action messages are retained on the AMRF queue for a particular period, it may mean more operators are needed for that period.
OMDGCMDI	Number of commands issued per second.
OMDGOREB	Maximum number of operator reply entries (OREs) on the system reply queue. To eliminate thrashing, use this number to monitor and adjust the ORE buffer limit set at IPL time. To dynamically adjust this limit, use the CONTROL M command.
OMDGWQEB	Maximum number of WTO queue elements (WQEs) on the system output queue. To eliminate thrashing (excessive data movement which confines system to doing little useful work), use this number to monitor and adjust the WTO buffer time limit set at IPL time. To dynamically adjust this limit, use the CONTROL M command.
OMDGWTLI	Number of write-to-logs (WTLs) issued per second indicating the number of records going to SYSLOG within a time period. To control the number of data sets produced during the day, vary the number of records per SYSLOG data set.
OMDGWTOI	Total number of lines of messages, write-to-operators (WTOs) issued per second. Use it to determine the peak message rate period and the average message rate.
RAXESCT	Number of common storage pages on expanded storage
RAXFMCT	Number of frames allocated to common storage
RCEAEC	Total number of expanded storage E frames currently on the ESTE queue
RCEAECLO	Available expanded storage low threshold

Table 5-58. Variables in the Trace Activity Report (continued)

Variable	Value
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
RCEMVBEL	Number of pages moved from below 16 megabytes in central storage
RCENWSF	Total number of secondary and non-working set pages migrated to auxiliary storage.
RCEPAGMV	Number of times a frame was moved from one frame to another
RCEPOOL	Number of frames currently available to the system. 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

Table 5-58. Variables in the Trace Activity Report (continued)

Variable	Value
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
RMCARQSC	Requested swap count
RMCARSSC	Central storage shortage swap count
RMCATISC	Terminal input swap count
RMCATOSC	Terminal output swap count
RMCATSSC	Count of transition swaps
RMCAUSSC	Unilateral swap out count
RMCAXSSC	Auxiliary storage shortage swap count
RMCTTRPC	Number of pages used for transaction elements
SMCABFLS	Number of records lost because of a shortage of buffers
SMCABFWT	Number of buffers written
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

VSTOR - Virtual Storage Activity report

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 the report, specify

```
REPORTS(VSTOR(S))
REPORTS(VSTOR(D[,jobname]))
```

Purpose of the 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`
- `COMMON STORAGE DETAIL`
- `PRIVATE AREA SUMMARY`
- `PRIVATE AREA DETAIL`

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 under "Contents of the report" on page 5-141 later in this section.

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 via 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 multiples 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.

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.

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 of 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.

```
                                    V I R T U A L   S T O R A G E   A C T I V I T Y
                                                                                                                PAGE    1
            z/OS V1R4               SYSTEM ID   SYS1           DATE 03/14/2002          INTERVAL 14.59.946
                                    RPT VERSION V1R2 RMF       TIME 16.30.00            CYCLE 1.000 SECONDS

                                              COMMON STORAGE SUMMARY
NUMBER OF SAMPLES      90
      STATIC STORAGE MAP                                                 ALLOCATED CSA/SQA
  AREA        ADDRESS   SIZE             ------------ BELOW 16M --------------  ------- EXTENDED (ABOVE 16M) --------
EPVT         1DD00000  1571M                MIN             MAX             AVG    MIN             MAX             AVG
ECSA          B0A4000   300M         SQA  1176K 23.44.59  1176K 23.44.59  1176K  25.0M 23.48.09  25.0M 23.54.39  25.0M
EMLPA               0     0K         CSA   704K 23.44.59   708K 23.54.10   704K  32.7M 23.55.49  32.7M 23.54.10  32.7M
EFLPA         B0A1000    12K
EPLPA         42FE000   110M         ALLOCATED CSA BY KEY
ESQA          1853000  42.7M          0    292K 23.44.59   292K 23.44.59   292K  14.2M 23.44.59  14.2M 23.57.30  14.2M
ENUC          1000000  8524K          1     52K 23.44.59    52K 23.44.59    52K   908K 23.44.59   908K 23.44.59   908K
----- 16 MEG BOUNDARY ------          2     36K 23.44.59    36K 23.44.59    36K  1696K 23.44.59  1696K 23.44.59  1696K
NUCLEUS         FC0000   256K          3      0K 23.44.59     0K              0K     0K 23.44.59     0K              0K
SQA             DB0000  2112K          4     16K 23.44.59    16K 23.44.59    16K   196K 23.44.59   196K 23.44.59   196K
PLPA            BFC000  1744K          5      8K 23.44.59     8K 23.44.59     8K  1304K 23.44.59  1304K 23.44.59  1304K
FLPA            BFB000     4K          6     68K 23.44.59    68K 23.44.59    68K  10.8M 23.44.59  10.8M 23.44.59  10.8M
MLPA                0     0K          7    228K 23.44.59   232K 23.54.10   228K  3664K 23.55.49  3672K 23.54.10  3668K
CSA             A00000  2028K         8-F     4K 23.44.59     4K 23.44.59     4K   108K 23.44.59   108K 23.44.59   108K
PRIVATE           2000  10.0M
PSA                 0     8K         SQA EXPANSION INTO CSA
                                           0K 23.44.59     0K              0K     0K 23.44.59     0K              0K
PLPA INTERMODULE SPACE -    3K IN PLPA AND  185K IN EPLPA
PLPA SPACE REDUNDANT WITH MLPA/FLPA -    0K IN PLPA AND    9K IN EPLPA

                     ---------- BELOW 16M ----------------  --------------- ABOVE 16M -----------
                       MIN             MAX             AVG    MIN             MAX             AVG
CSA
 FREE PAGES (BYTES)  1320K 23.54.10  1324K 23.44.59  1323K   268M 23.47.00   268M 23.55.49   268M
 LARGEST FREE BLOCK  1320K 23.44.59  1320K 23.44.59  1320K   268M 23.44.59   268M 23.44.59   268M
 ALLOCATED AREA SIZE  708K 23.44.59   708K 23.44.59   708K   300M 23.44.59   300M 23.44.59   300M
SQA
 FREE PAGES (BYTES)   936K 23.44.59   936K 23.44.59   936K  17.6M 23.54.39  17.7M 23.48.09  17.7M
 LARGEST FREE BLOCK   844K 23.44.59   844K 23.44.59   844K  17.5M 23.50.19  17.5M 23.44.59  17.5M
 ALLOCATED AREA SIZE 1268K 23.44.59  1268K 23.44.59  1268K  42.7M 23.44.59  42.7M 23.44.59  42.7M
MAXIMUM POSSIBLE USER REGION - 10032K BELOW AND 1562M ABOVE

                     ---------- BELOW 16M ----------------  --------------- ABOVE 16M -----------
                       MIN             MAX             AVG    MIN             MAX             AVG
 CSA
  FREE PAGES (BYTES)  2976K 16.32.29  3056K 16.43.09  3030K  6312K 16.37.29  6376K 16.31.39  6338K
  LARGEST FREE BLOCK  2976K 16.32.29  3008K 16.29.59  2995K  6268K 16.29.59  6268K 16.29.59  6268K
  ALLOCATED AREA SIZE 1372K 16.29.59  1436K 16.32.29  1373K  21.4M 16.29.59  21.4M 16.29.59  21.4M
 SQA
  FREE PAGES (BYTES)     0K 16.29.59     4K 16.32.58     0K  54.0M 16.32.39  54.1M 16.43.09  54.0M
  LARGEST FREE BLOCK     0K 16.29.59     4K 16.32.58     0K  53.6M 16.29.59  53.6M 16.29.59  53.6M
  ALLOCATED AREA SIZE 1292K 16.29.59  1292K 16.29.59  1292K  97.4M 16.29.59  97.4M 16.29.59  97.4M
 MAXIMUM POSSIBLE USER REGION -  8084K BELOW AND 1869M ABOVE
```

Figure 5-42. VSTOR Report - Common Storage Summary Section

Table 5-59. 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.
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. **Note:** 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 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 5-60. Overview Names in the Virtual Storage Activity Report

Field Heading	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.

```
                                V I R T U A L   S T O R A G E    A C T I V I T Y
                                                                                                PAGE 2

           z/OS V1R4               SYSTEM ID   SYS4           DATE 03/16/2002          INTERVAL 28.08.260
                                   RPT VERSION V1R2 RMF       IME 05.31.51            CYCLE 0.250 SECONDS

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       96K 05.31.51     4K 05.31.51    124K 05.31.51   288K 05.31.51   226      68K 05.33.26   224K 05.31.51   76K
 1                        4K 05.31.51      0K 05.31.51    20K 05.31.51   239      88K 05.31.51    88K 05.31.51   88K
 2                                                                       245     168K 05.31.51   192K 05.59.56  168K
 3
 4
 5                                                        4K 05.31.51
 6       16K 05.31.51     0K 05.31.51     16K 05.31.51     8K 05.31.51
 7
8-F
ALL     112K 05.31.51     8K 05.31.51    140K 05.31.51   320K 05.31.51
-------------------------- MAXIMUM -----------------------------------
 0       96K 05.31.51     4K 05.31.51    124K 05.31.51   288K 05.31.51
 1                        4K 05.31.51     24K 05.53.29    24K 05.59.27
 2
 3
 4
 5                                                        4K 05.31.51
 6       16K 05.31.51     4K 05.41.39     16K 05.31.51     8K 05.31.51
 7
8-F
ALL     112K 05.31.51    12K 05.41.39    164K 05.53.29   324K 05.59.27
-------------------------- AVERAGE -----------------------------------
 0       96K              4K             124K            288K
 1                        4K               0K             20K
 2
 3
 4
 5                                                        4K
 6       16K              0K              16K             8K
 7
8-F
ALL     112K              8K             140K            320K
```

Figure 5-43. VSTOR Report - Common Storage Detail Section

Table 5-61. 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."

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.

```
                                V I R T U A L   S T O R A G E   A C T I V I T Y
                                                                                PAGE 2
   z/OS V1R4               SYSTEM ID   SYS1          DATE 03/14/2002         INTERVAL 14.59.946
                           RPT VERSION V1R2 RMF      TIME 16.30.00           CYCLE 1.000 SECONDS

                                      PRIVATE AREA SUMMARY
   JOB NAME -       RMF                                    REGION REQUESTED             32.0M
   STEP NAME -      RMF                                    REGION ASSIGNED (BELOW 16M)  9.90M
   PROGRAM NAME -   ERBMFMFC                               REGION ASSIGNED (ABOVE 16M)  32.0M
   NUMBER OF SAMPLES -    89

                                      PRIVATE STORAGE MAP
                    BELOW 16M                                EXTENDED (ABOVE 16M)
          9FFFFF ________________________                 ________________________ 7FFFFFFF
                 | LSQA/SWA              |                | LSQA/SWA              |
                 | 229/230        240K   |   BOTTOM OF    | 229/230        11.5M  |
          9C4000 |_____12.45.11_________|  ALLOCATED AREA |_____12.45.11_________| 7F482000
                 | UNUSED           0K   |                | UNUSED         1529M  |
          A00000 |_______________________|   GETMAIN LIMIT |_______________________| 1FC00000
                 | UNUSED         9468K  |                | UNUSED         16.3M  |
           85000 |_____12.45.11_________|      TOP OF     |_____12.45.11_________| 1EBBC000
           85000 |_____12.45.11_________|      TOP OF     |_____12.45.11_________| 1EBBC000
                 | USER                  |  ALLOCATED AREA | USER                  |
                 | REGION         508K   |                | REGION                |
            6000 |_______________________|                |                15.7M  |
                 | SYSTEM REGION   16K   |                |                       |
            2000 -------------------------                ------------------------- 1DC00000

                    ---------- BELOW 16M ----------------  --------------- ABOVE 16M -----------
                      MIN             MAX             AVG    MIN             MAX             AVG
  LSQA/SWA/229/230
   FREE PAGES (BYTES)     0K 12.45.11     0K             0K  1529M 12.45.11  1529M 12.45.11  1529M
   LARGEST FREE BLOCK     0K 12.45.11     0K             0K  1529M 12.45.11  1529M 12.45.11  1529M
   PAGES ALLOCATED
    (IN BYTES)          208K 12.45.11   208K 12.45.11   208K  11.5M 12.45.11  11.5M 12.45.11  11.5M
  USER REGION
   FREE PAGES (BYTES)  9508K 12.45.11  9508K 12.45.11  9508K  19.5M 12.45.11  19.5M 12.45.11  19.5M
   LARGEST FREE BLOCK
     IN GETMAIN LIMIT  9468K 12.45.11  9468K 12.45.11  9468K  16.3M 12.45.11  16.3M 12.45.11  16.3M
   PAGES ALLOCATED
     (IN BYTES)         500K 12.45.11   500K 12.45.11   500K  12.5M 12.45.11  12.5M 12.45.11  12.5M
```

Figure 5-44. VSTOR Report - Private Area Summary Section

Table 5-62. 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.

Table 5-62. Fields in the Virtual Storage Activity Report - Private Area Summary (continued)

Field Heading	Meaning
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 expanded storage (above the 16-megabyte line). Each map identifies:

Table 5-63. 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 5-64. 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.

Table 5-64. Fields in the Virtual Storage Activity Report - Bottom Half (continued)

Field Heading	Meaning
PAGES ALLOCATED (IN BYTES)	The MIN, MAX, and AVG values for the amount of allocated storage. **Note:** 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 long-running non-swappable jobs, such as JES2, JES3, VTAM, IMS. 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 contains additional information about the use of allocated storage in the private area below the 16-megabyte line. It includes MIN, MAX, and AVG values for the number of bytes of allocated blocks, broken down by subpool and by area (LSQA, SWA).

To measure the usage of virtual storage above the 2-gigabyte-line, the private area detail section provides MIN, MAX, AVG values for the interval as well as the PEAK value since the start of the address space.

```
                              V I R T U A L   S T O R A G E   A C T I V I T Y
                                                                                PAGE 4

       z/OS V1R4               SYSTEM ID   SYS4           DATE 04/16/2002         INTERVAL 28.08.260
                               RPT VERSION V1R2 RMF       TIME 05.31.51           CYCLE 0.250 SECONDS

JOB NAME -  CATALOG
NUMBER OF BYTES OF ALLOCATED BLOCKS BY AREA (BELOW 16 MEG)
SUBPOOL (AREA)      MIN             MAX             AVG
  229             240K 05.31.52   240K 05.31.52   240K
  230             488K 05.31.52   508K 05.42.24   488K
  236 (SWA)       100K 05.31.52   100K 05.31.52   100K
  237 (SWA)         4K 05.31.52     4K 05.31.52     4K
  255 (LSQA)       32K 05.31.52    36K 05.42.26    34K
USER REGION
    0                4K 05.31.52     4K 05.31.52     4K
  252 (REENTRANT)  732K 05.31.52   736K 05.58.51   732K
NUMBER OF BYTES ALLOCATED IN HIGH VIRTUAL MEMORY (ABOVE 2GB)
                   MIN             MAX             AVG   PEAK
  TOTAL              0 12.10.32 16.00M 12.10.12 546.1K 16.00E
```

Figure 5-45. VSTOR Report - Private Area Detail Section

WKLD - Workload Activity report - Compatibility Mode

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 a report, specify

```
REPORTS(WKLD(options))
```

The contents of the Workload Activity report depends significantly on the mode your system is running:

- **Compatibility Mode**

  When Monitor I is gathering data on a system that runs in compatibility mode, the Workload Activity report is a single-system report and can be created directly by Monitor I as real-time report or it can be created by the Postprocessor. The content of this report is described in this chapter.
- **Goal Mode**

  When Monitor I is gathering data on systems that run in goal mode, then a sysplex-wide report is available from the Postprocessor. Refer to "WLMGL - Workload Activity report - Goal Mode" on page 5-153 for a description.

Purpose of the report

The Workload Activity report provides information about the use of each performance group period (PGP) within a performance group. The PGPs are defined in the installation performance specification (IPS - IEAIPSxx Parmlib member) and the installation control specification (ICS - IEAICSxx Parmlib member). The definition in the ICS for a performance group can include a subsystem (SUBSYS), transaction class (TRXCLASS), account information (ACCTINFO), user identifier (USERID), service class (SRVCLASS) or transaction name (TRXNAME). To enable you to associate data for a performance group with the definition for that performance group in the ICS, RMF includes this information in the reports by performance group period and the reports by performance group.

Generating a report

To get the information in the format best suited to your needs, use one or a list of the following sub-report options:

- **By performance group period within performance group** - the most detailed report. Gives information on each performance group period. A summary for each performance group immediately follows the information for all performance group periods in the performance group; however, this summary is not produced if all performance group periods have zero interval service. Option:

  `WKLD(PERIOD)`
- **By performance group** - summarizes the data for all performance group periods within the performance group. Option:

  `WKLD(GROUP)`
- **By domain in a range of performance groups** - summarizes the data by domain number within a range of ten performance group numbers. Option:

  `WKLD(RANGE)`

- **By domain for entire system**- summarizes the data by domain for the entire system. Option:
  `WKLD(DOMAIN)`
- **For entire system**- summarizes the data for the entire system. Option:
  `WKLD(SYSTEM)`
- **By time slice group**- summarizes the data for all performance group periods within the time slice group. This report can be produced only when time slice groups have been defined in the IPS. Option:
  `WKLD(TIME)`

To use the performance group period report for individual workload reports, you can assign unique performance group numbers and domain numbers to a particular user or set of users, such as subsystems, monitor programs, groups of users, or batch initiator classes.

The system resources manager (SRM) uses the IPS and the ICS values to perform its workload balancing functions. If the SRM values change during a measurement interval, SRM notifies RMF. RMF then uses the old SRM values to generate SMF records and, if REPORT is in effect, reports for the part of the interval that had elapsed before the SRM values changed. RMF then begins a new interval using the new SRM values.

If the IPS or ICS value change during an interval when RMF is measuring workload activity, RMF will execute normal end-of-interval processing and generate a set of SMF records. If the REPORT option is in effect, RMF reports for the part of the interval that elapsed before the change. RMF then reinitializes all active measurements and continues normal processing using the new values. Although RMF automatically stops and restarts the session, no data is sampled during the brief period required to stop, modify, and restart the session.

Using the information given in the report

Capture ratio

The **APPL%** value tells you how much CPU time was captured for the workload in a given group or domain. If you have more than one CPU, divide this number by the number of CPUs in your system. You can use this number as input to capacity planning, to compute the capture ratio of the system, and to check for a single-CP constraint.

You calculate a capture ratio as follows:

```
                    APPL %
Capture ratio = ----------------
                AVG BUSY % * # CP
```

You can check the single and block **PAGE-IN RATES** for any workload that is sensitive to paging, such as CICS, to make sure no paging is taking place.

Contents of the report

Report performance groups

The ICS also enables you to define performance groups designed solely to **report on subsystem transaction activity**. These report performance groups are included only in the reports by performance group period and by performance group. Report performance groups have no service objective, domain, or time slice group.

Therefore, the DMN and TSG fields for a report performance group period in the PERIOD and GROUP subreports contain two asterisks(**).

The other fields in the report can contain, for report performance groups, only the data that a subsystem provides to the SRM. The system-defined subsystems (TSO/E and STC) and the job entry subsystems (such as JES2 or JES3) provide meaningful data for all fields in the report, except those noted earlier as not applying to report performance groups. For other subsystems (such as a user-written subsystem), only ENDED TRANSACTIONS and AVG TRANS TIME contain meaningful data and INTERVAL SERVICE might contain meaningful data. (See the current documentation for the subsystem to determine if it provides service information to the SRM.) A message will inform you if the data is all zero.

Note: If RMF is producing Workload Activity reports, the time is included in the AVG TIME OF ENDED TRANS (TSO/E response time). The time reported might be greater than the actual average time for normal transactions for the following reason: A TSO/E user can press the TEST REQ key on a 3277 terminal to suspend the TSO/E session and then resume the session by a LOGON reconnect command. For the time that elapses between these two actions, the system places the user in a wait state and the SRM continues to accumulate CPU and transaction time for the user.

Subsystem information

If your ICS associates the performance group with a single subsystem name, the reports by performance group period and the reports by performance group include two lines of information that precede the data for each performance group. These lines identify the ICS parameters for the performance group. The lines have the format:

```
SUBSYS=subsystem name TRXCLASS=transaction class ACCTINFO=account information
USERID=user identifier TRXNAME=transaction name SRVCLASS=service class name
```

See the *z/OS RMF User's Guide* for more details.

The report lists user identifier, transaction class, transaction name, or accounting information only when (1) the corresponding parameter appears in the ICS entry for the performance group and (2) the parameter is the only user identifier, transaction class, or transaction name associated with the performance group.

Execution velocity

You can use the execution velocity (**EX VEL %**) field to define performance goals when migrating to goal mode.

Report specific data fields

In addition to the report headers that are common to all reports, the Workload Activity report header contains:

- IPS identification number
- System mode indicator: MODE=COMPAT
- Sub-report title
- OPT and ICS Parmlib member names
- Indicator for I/O priority management: I/O PRTY MGMT = YES/NO

  This value indicates whether I/O delays have been included in the calculation of execution velocity. NO is an indication that the value has not been specified in the service policy or that no service policy has been defined. A value defined in the IPS specification will not be shown.

- Values for the IOC, CPU, MSO, and SRB service definition coefficients as defined in the IPS specification
- SU/SEC=xxx: the number of SRM processor service units per CPU second

Duration report

When you request a duration report, the message IPS CHANGED DURING DURATION INTERVAL will appear at the end of every sub-report if the IPS changed during the duration of time covered in the report. The message indicates a change to the SRM values in the IPS, ICS, or OPT. Because the SRM informs RMF when such a change occurs, the values reported are meaningful.

```
                                    W O R K L O A D   A C T I V I T Y
                                                                                                                PAGE    1
          z/OS V1R4               SYSTEM ID   SYS1           DATE 03/15/2002          INTERVAL 20.20.060      MODE=COMPAT
                                  RPT VERSION V1R2 RMF       TIME 11.30.00            IPS = IEAIPS60

        OPT = IEAOPT60   I/O PRTY MGMT = YES REPORT BY PERFORMANCE GROUP    SERVICE DEFINITION COEFFICIENTS     SU/SEC= 619.4
        ICS = IEAICS60                                PERIOD               IOC = 1.0   CPU = 1.0   SRB = 1.0   MSO = 1.0000

PGN  PGP DMN TSG  TRANSACTIONS   ----TRANS. TIME-----  --DASD I/O--  ---SERVICE---  -SERVICE RATES- PAGE-IN RATES  ----STORAGE----
                                        HHH.MM.SS.TTT

0001 1   001 **   AVG    33.00   ACTUAL        10.616  SSCHRT   0.0  IOC    18067   ABSRPTN  1,318  SINGLE   0.00  AVERAGE  1,237
                  MPL    33.00   EXECUT         8.316  RESP    15.2  CPU    358.5K  TRX SERV 1,318  BLOCK    0.00
                  ENDED      3   QUEUED         2.300  CONN     4.8  MSO    52636K  TCB      578.8  SHARED   0.00  TOTAL   40,841
                  END/S   0.00   R/S AFF            0  DISC     4.5  SRB    78612   SRB      126.9  HSP      0.00  CENTRAL 40,841
                  #SWAPS     6   INELIG             0  Q+PEND   5.9  TOT    53091K  RCT        0.0  HSP MISS 0.00  EXPAND    0.00
                  AVG ENC 0.32   CONV           8.388  IOSQ     0.0  /SEC  43,515   IIT        3.7  EXP SNGL 0.00  SHARED  4325.6
                  REM ENC 0.12   STDDEV         3.388                               HST        0.0  EXP BLK  0.00
                  MS ENC  0.00                                                       APPL%     58.1  EXP SHR  0.00
                                                                                     EX VEL%   24.1

0002 1   001 **   AVG     2.15   ACTUAL         1.761  SSCHRT   0.0  IOC    5,760   ABSRPTN  1,280  SINGLE   0.05  AVERAGE 360.43
                  MPL     1.48   EXECUT         1.755  RESP    28.4  CPU    168.5K  TRX SERV   881  BLOCK    0.00
                  ENDED  1,797   QUEUED             6  CONN    13.5  MSO     2139K  TCB      271.9  SHARED   0.00  TOTAL   534.26
                  END/S   1.47   R/S AFF            0  DISC     1.4  SRB    2,534   SRB        4.0  HSP      0.00  CENTRAL 534.26
                  #SWAPS 1,850   INELIG             0  Q+PEND  13.5  TOT     2316K  RCT        7.6  HSP MISS 0.00  EXPAND    0.00
                  AVG ENC 0.00   CONV              16  IOSQ     0.0  /SEC   1,898   IIT        1.2  EXP SNGL 0.00  SHARED    0.00
                  REM ENC 0.00   STDDEV        11.140                               HST        0.0  EXP BLK  0.00
                  MS ENC  0.00                                                       APPL%     23.3  EXP SHR  0.00
                                                                                     EX VEL%    8.3
```

Figure 5-46. WKLD Report

For the field descriptions of the Workload Activity report, see Table 5-65 on page 5-159.

WLMGL - Workload Activity report - Goal Mode

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))
```

The Workload Activity report in goal mode is different from the compatibility mode version, and only exists as a Postprocessor report.

The WLMGL option list

You can select for what level you want the report by specifying the SYSRPTS WLMGL suboptions:

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

WGROUP 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.

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 general appearance of a report is shown in the following sample:

```
Header                  +-->                                        W O R K L O A D   A C T I V I T Y
                        |
                        |
                        |              z/OS V1R4              SYSPLEX BBHULB            DATE  04/15/2002       INTERVAL  15
                        +-->                                  RPT VERSION V1R2 RMF      TIME  00.03.38

Service class header    --->                                      POLICY ACTIVATION DATE/TIME  04/14/2002 23.38.94
Separation/Divider      +--> ========================================//================================== WORKLOAD
 line(s)
                        |    WORKLOAD=CICSSTC
                        +-->                                   ALL CICS STC'S

Service class           +-->  TRANSACTIONS   TRANS.-TIME  HHH.MM. S.TTT   --DASD I/O -   ---SERVICE---   --SERVICE RATES--   -P
                        |     AVG       0.00  ACTUAL              0.040   SSCHRT  36.0   IOC     43255   ABSRPTN      5558   SI
                        |     MPL       0.00  EXECUTION           0.040   RESP    36.4   CPU    171896   TRX SERV     5286   BL
                        |     ENDED    1,040  QUEUED                0   CONN    13.6   MSO    520171   TCB           567   SH
                        |     END/SEC   0.02  R/S AFFINITY           0   DISC     1.4   SRB     35686   SRB           348   HS
                        |     #SWAPS   1,028  INELIGIBLE             0   Q+PEND  21.4   TOT    781008   RCT           0.1   HS
                        |     EXCTD      997  CONVERSION             0   IOSQ     0.0   /SEC       16   IIT           0.1   EX
                        |     AVG ENC   0.00  STD DEV             0.180                                  HST           0.1   EX
                        |     REM ENC   0.00                                                             APPL%       101.6   EX
                        |     MS ENC    0.00
                        +-->
Velocity migration      +--> VELOCITY MIGRATION:   I/O MGMT   1.8%    INIT MGMT   1.7%

Service classes being   +--> ------------------------------------------------SERVICE CLASSES BEING SERVED-----------------------
 served                 +--> ATMPLANT   ATMLABOR   POSMULTI

Service class period    +--> PER IMPORTANCE  PERF     --TRANSACTIONS--     ----------RESPONSE TIME----------   -EX VEL%-   -CPU
 summary in the WGPER   |                    INDX     --NUMBER--          ------GOAL------  ACTUAL    TOTAL   GOAL  ACT   USIN
 report                 |     1  1           0.91           574   72%           0.300 90%   93.0%    66.9%          73   65.0
                        |     2  2           1.09           109   14%           0.900 90%   89.2%
                        |     3  3           1.49            87   11%           2.000 90%   74.1%
                        |     4                              26    3%
                        +--> TOTAL                         796  100%
Subsystem               +-->            RESP ------------------------------STATE SAMPLES BREAKDOWN (%)-------------------------
                        |    SUB    P   TIME --ACTIVE-- READY  IDLE   ----------------------------WAITING FOR------------------
                        |    TYPE        (%)  SUB  APPL               CONV    I/O   LOCK   LOCL  SYSP  REMT TIME
                        |    CICS  BTE  93.4 10.9   0.0   0.0   0.0   89.2    0.0    0.0    0.0   0.0   0.0  0.0
                        +--> CICS  EXE  67.0 19.7   0.0  10.6   0.0    0.0   39.7   10.2    8.4   5.4   3.7  2.3

Service class period    +-->         ---RESPONSE TIME---  EX   PERF   AVG   --USING%--  ------------- EXECUTION DELAYS % -------
                        |            HH.MM.SS.TTT         VEL  INDX  ADRSP   CPU   I/O  TOTAL   CPU  I/O  AUX AUX  SWIN
                        |    GOAL    00.00.05.000  90.0%                                                    VIO PRIV
                        |    ACTUALS
                        |    *ALL                  52.6% 31.8%  4.0    5.7   3.6   2.6   13.1   8.5  4.3  0.2  0.1  0.1
                        |    D0                    57.5% 30.5%  4.0    2.1   3.5   2.3   13.3   9.0  4.2  0.1  0.0  0.1
                        |    D4                    53.0% 49.4%  4.0    1.9   3.3   3.1    6.5   1.8  4.3  0.1  0.0  0.0
                        +--> D6                    45.5% 24.1% ****    1.8   3.9   2.3   19.8  14.8  4.3  0.1  0.4  0.1

Response Time           +--> -----------------RESPONSE TIME DISTRIBUTION-------------------------------------------------------
 Distribution Chart     |       ----TIME----    --NUMBER OF TRANSACTIONS--    -------PERCENT-------  0    10   20   30   40   5
                        |       HH.MM.SS.TTT    CUM TOTAL        IN BUCKET    CUM TOTAL   IN BUCKET  |....|....|....|....|....|
                        |    <  00.00.02.500          692              692         36.8        36.8  >>>>>>>>>>>>>>>>>>
                        |    <= 00.00.03.000          759               67         40.4         3.6  >>>
                        |    <= 00.00.03.500          829               70         44.1         3.7  >>>
                        |    <= 00.00.04.000          898               69         47.8         3.7  >>>
                        |    <= 00.00.04.500          952               54         50.7         2.9  >>
                        |    <= 00.00.05.000          989               37         52.6         2.0  >>
                        +--> ...
```

Figure 5-47. WLMGL Report Sections

Policy change

For all reports, you have to consider that a policy change will interrupt the current interval.

The following sections show a sample of each of the reports, and at the end there is a description of the report fields.

Service Class Period report

To request this report, specify
`WLMGL(SCPER(option))`

For example, to produce a report for all service class periods of service class TSOCLASS, specify:
`WLMGL(SCPER(TSOCLASS))`

The example in Figure 5-48 shows a report for service class TSOCLASS and the first service class period.

```
                                        W O R K L O A D   A C T I V I T Y
                                                                                                                PAGE   7
      z/OS V1R4         SYSPLEX UTCPLXHD           DATE 06/30/2002          INTERVAL 30.01.919   MODE = GOAL
                        RPT VERSION V1R2 RMF       TIME 14.30.00

                                          POLICY ACTIVATION DATE/TIME 06/27/2002 14.52.04
REPORT BY: POLICY=DAILY      WORKLOAD=PRODUCTN   SERVICE CLASS=TSOCLASS   RESOURCE GROUP=REGTSO     PERIOD=2 IMPORTANCE=2
                                                 CRITICAL     =CPU + STORAGE
                                                 DESCRIPTION  =Normal 3-period TSO work

TRANSACTIONS    TRANS.-TIME  HHH.MM.SS.TTT   --DASD I/O--   ---SERVICE----   --SERVICE RATES--   PAGE-IN RATES    ----STORAGE----
AVG      4.42   ACTUAL              8.142   SSCHRT   6.5   IOC     56040    ABSRPTN       598   SINGLE     0.0   AVG      298.08
MPL      4.39   EXECUTION           8.142   RESP    53.6   CPU      1680K   TRX SERV      594   BLOCK      0.0   TOTAL   1309.04
ENDED    1879   QUEUED                 0   CONN    20.0   MSO      2938K   TCB          271.3   SHARED     0.0   CENTRAL 1309.04
END/SEC  1.04   R/S AFFINITY           0   DISC     3.2   SRB     56695    SRB            9.2   HSP        0.0   EXPAND     0.00
#SWAPS   3154   INELIGIBLE             0   Q+PEND  25.7   TOT      4731K   RCT            4.4   HSP MISS   0.0
EXCTD       0   CONVERSION             0   IOSQ     4.7   /SEC     2626    IIT            5.4   EXP SNGL   0.0   SHARED     0.00
AVG ENC  0.32   STD DEV             8.431                                  HST            0.0   EXP BLK    0.0
REM ENC  0.12                                                             APPL %        16.1   EXP SHR    0.0
MS ENC   0.01

VELOCITY MIGRATION:   I/O MGMT  31.8%    INIT MGMT 31.8%

        ---RESPONSE TIME---  EX   PERF   AVG   --USING%--  ------------- EXECUTION DELAYS % -------------  ---DLY%-- -CRYPTO%-    %
        HH.MM.SS.TTT         VEL  INDX  ADRSP   CPU   I/O  TOTAL   CPU  I/O  AUX AUX  SWIN                 UNKN IDLE  DLY  USG QUIE
GOAL    00.00.05.000  90.0%                                                  VIO PRIV
ACTUALS
*ALL                 52.6% 31.8%  4.0    5.7   3.6   2.6   13.1   8.5  4.3  0.2  0.1  0.1                 58.1 22.7  1.1  3.1  0.0
D0                   57.5% 30.5%  4.0    2.1   3.5   2.3   13.3   9.0  4.2  0.1  0.0  0.1                 60.3 20.5  0.2  1.1  0.0
D4                   53.0% 49.4%  4.0    1.9   3.3   3.1    6.5   1.8  4.3  0.1  0.0  0.0                 63.1 24.0  1.4  4.1  0.0
D6                   45.5% 24.1% ****    1.8   3.9   2.3   19.8  14.8  4.3  0.1  0.4  0.1                 50.3 23.7  0.3  0.3  0.0

                                    ----------RESPONSE TIME DISTRIBUTION----------
   ----TIME----    --NUMBER OF TRANSACTIONS--    -------PERCENT-------  0    10   20   30   40   50   60   70   80   90   100
   HH.MM.SS.TTT    CUM TOTAL        IN BUCKET    CUM TOTAL   IN BUCKET  |....|....|....|....|....|....|....|....|....|....|
<  00.00.02.500         692              692         36.8        36.8  >>>>>>>>>>>>>>>>>>>
<= 00.00.03.000         759               67         40.4         3.6  >>>
<= 00.00.03.500         829               70         44.1         3.7  >>>
<= 00.00.04.000         898               69         47.8         3.7  >>>
<= 00.00.04.500         952               54         50.7         2.9  >>
<= 00.00.05.000         989               37         52.6         2.0  >>
<= 00.00.05.500        1016               27         54.1         1.4  >
<= 00.00.06.000        1041               25         55.4         1.3  >
<= 00.00.06.500        1073               32         57.1         1.7  >>
<= 00.00.07.000        1101               28         58.6         1.5  >
<= 00.00.07.500        1101                0         58.6         0.0  >
<= 00.00.10.000        1233              132         65.6         7.0  >>>>
<= 00.00.20.000        1714              481         91.2        25.6  >>>>>>>>>>>>>>>
>  00.00.20.000        1879              165          100         8.8  >>>>>
```

Figure 5-48. WLMGL - Service Class Period Report

If a service class has more than one period, a summary is shown at the end of all service class period reports which is equivalent to the Service Class report (see Figure 5-49 on page 5-156).

Field descriptions

For a description of the fields, see "Field descriptions of all reports" on page 5-159.

Service Class report

To request this report, specify
`WLMGL(SCLASS(option))`

For example, to produce the report shown in Figure 5-49:
`WLMGL(SCLASS(TSOCLASS))`

The example in Figure 5-49 shows a report for service class TSOCLASS.

```
                                    W O R K L O A D   A C T I V I T Y
                                                                                                    PAGE   5
     z/OS V1R4              SYSPLEX UTCPLXHD          DATE 05/01/2002          INTERVAL 25.28.837   MODE = GOAL
                            RPT VERSION V1R2 RMF      TIME 16.00.00

                                      POLICY ACTIVATION DATE/TIME 05/01/2002 15.57.22

REPORT BY: POLICY=DAILY     WORKLOAD=PRODUCTN  SERVICE CLASS=TSOCLASS  RESOURCE GROUP=REGTSO
                                               CRITICAL     =NONE
                                               DESCRIPTION  =Normal 3-period TSO work

TRANSACTIONS    TRANS.-TIME  HHH.MM.SS.TTT   --DASD I/O--   ---SERVICE----   --SERVICE RATES--   PAGE-IN RATES    ----STORAGE----
AVG     33.72   ACTUAL              4.693   SSCHRT  4348   IOC     77945    ABSRPTN       278   SINGLE     0.0   AVG      265.69
MPL     33.70   EXECUTION           4.687   RESP    36.8   CPU      3953K   TRX SERV      278   BLOCK      0.0   TOTAL   8954.53
ENDED    5660   QUEUED                 2   CONN    13.8   MSO      7117K   TCB         638.2   SHARED     0.0   CENTRAL 8954.53
END/SEC  6.28   R/S AFFINITY           0   DISC     1.2   SRB    107721    SRB          17.4   HSP        0.0   EXPAND     0.00
#SWAPS   5871   INELIGIBLE             4   Q+PEND  21.8   TOT     11255K   RCT          21.8   HSP MISS   0.0
EXCTD       0   CONVERSION             0   IOSQ     0.0   /SEC    12750    IIT           2.6   EXP SNGL   0.0   SHARED     0.00
AVG ENC  0.32   STD DEV           37.097                                     HST           0.0   EXP BLK    0.0
REM ENC  0.12                                                                APPL %       56.7   EXP SHR    0.0
MS ENC   0.04
```

Figure 5-49. WLMGL - Service Class Report

The report can contain an additional part SERVICE CLASSES BEING SERVED if address spaces are doing work for transactions that were classified to a another service class in this report.

Workload Group report

To request this report, specify
`WLMGL(WGROUP(option))`

For example, to produce the report for a workload group called PRODUCTN, specify:
`WLMGL(WGROUP(PRODUCTN))`

The report has the same layout as a Service Class report, but all service classes associated with the workload PRODUCTN are combined in a workload summary.

Policy Summary report

To request this report, specify
`WLMGL(POLICY(option))`

For example, to produce a report for a policy named HOLIDAY, specify:
`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.

Workload Group and Service Class Period report

To request this report, specify
`WLMGL(WGPER(option))`

For example, to produce a report for a workload group called PRODUCTN, specify:
`WLMGL(WGPER(PRODUCTN))`

Each service class associated with the workload PRODUCTN is listed with its service class periods and the defined and achieved GOALS for the periods.

```
                                    W O R K L O A D   A C T I V I T Y
                                                                                                    PAGE   5
     z/OS V1R4              SYSPLEX UTCPLXHD          DATE 05/01/2002          INTERVAL 25.28.837   MODE = GOAL
                            RPT VERSION V1R2 RMF      TIME 16.00.00

                                       POLICY ACTIVATION DATE/TIME 05/01/2002 15.57.22

REPORT BY: POLICY=DAILY     WORKLOAD=PRODUCTN  SERVICE CLASS=TSOCLASS  RESOURCE GROUP=REGTSO
                                               CRITICAL     =NONE
                                               DESCRIPTION  =Normal 3-period TSO work

TRANSACTIONS    TRANS.-TIME  HHH.MM.SS.TTT   --DASD I/O--   ---SERVICE----   --SERVICE RATES--   PAGE-IN RATES    ----STORAGE----
AVG     33.72   ACTUAL              4.693   SSCHRT  4348   IOC     77945    ABSRPTN       278   SINGLE     0.0   AVG      265.69
MPL     33.70   EXECUTION           4.687   RESP    36.8   CPU      3953K   TRX SERV      278   BLOCK      0.0   TOTAL   8954.53
ENDED    5660   QUEUED                 2   CONN    13.8   MSO      7117K   TCB         638.2   SHARED     0.0   CENTRAL 8954.53
END/SEC  6.28   R/S AFFINITY           0   DISC     1.2   SRB    107721    SRB          17.4   HSP        0.0   EXPAND     0.00
#SWAPS   5871   INELIGIBLE             4   Q+PEND  21.8   TOT     11255K   RCT          21.8   HSP MISS   0.0
EXCTD       0   CONVERSION             0   IOSQ     0.0   /SEC    12750    IIT           2.6   EXP SNGL   0.0   SHARED     0.00
AVG ENC  0.32   STD DEV           37.097                                    HST           0.0   EXP BLK    0.0
REM ENC  0.12                                                               APPL %       56.7   EXP SHR    0.0
MS ENC   0.04

  PER IMPORTANCE  PERF     --TRANSACTIONS--     -------------RESPONSE TIME--------------   -EX VEL%-   -CPU--  -EXE--
                  INDX     -NUMBER-     -%-     ------GOAL------  ---ACTUAL---     TOTAL   GOAL  ACT   USING%  DELAY%
  1   2            1.1         4442      78     00.00.02.000  90%    90.0%          70.6%        1.7     0.0     0.0
  2   2            0.5          771      14     00.00.05.000  90%    96.2%                       1.8     0.0     0.0
  3   2            0.6          447       8     00.00.15.000  90%    95.3%                      18.5     0.0     0.0
TOTAL                          5660     100
```

Figure 5-50. WLMGL - Workload Group and Service Class Period Report

Report Class report

To request this report, specify
`WLMGL(RCLASS(option))`

For example, to produce a report for a report class called BATCH, specify:
`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
`WLMGL(RCPER(option))`

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 TSODEPTA, TSODEPTB, and TSODEPTC. 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 will be created automatically for each interval. This page contains the following information:

- Installed date
- Service definition coefficients
- System ID and SU/SEC including interval start time and interval length for each system
- Resource group definitions
- Resource group actuals for each service class

This page is always included at the end of an interval.

```
                                    W O R K L O A D   A C T I V I T Y
                                                                                                    PAGE   9
     z/OS V1R4              SYSPLEX UTCPLXHD          DATE 05/01/2002          INTERVAL 25.28.837   MODE = GOAL
                            RPT VERSION V1R2 RMF      TIME 16.00.00

                                      POLICY ACTIVATION DATE/TIME 05/01/2002 15.57.22
                                                  - SERVICE POLICY PAGE -

SERVICE DEFINITION: coeffs    Service coefficients                          -SERVICE DEFINITION COEFFICIENTS-
  INSTALL DATE: 05/01/2002 15.13.19  INSTALLED BY: SETUP                        IOC     CPU     SRB     MSO
  POLICY: DAILY     Default PRODUCTN policy with ResGrp
  I/O PRIORITY MANAGEMENT: YES                                                  5.0    10.0    10.0  3.0000
  DYNAMIC ALIAS MANAGEMENT: YES
SYSTEMS
  ---ID--- OPT SU/SEC  --TIME-- INTERVAL   ---ID--- OPT SU/SEC  --TIME-- INTERVAL   ---ID--- OPT SU/SEC  --TIME-- INTERVAL
  D0        42  619.46 16.00.00 00.25.28   D1        42  619.46 16.04.44 00.14.40   D2        42  619.46 16.04.15 00.21.08
  D3        42  619.46 16.03.58 00.14.39   D4        42  619.46 16.00.00 00.06.14   D6        42  619.46 16.04.01 00.21.26
  D7        42  619.46 16.03.54 00.14.41

RESOURCE GROUPS
  --NAME--    ----------DESCRIPTION-----------     -SERVICE-      ACTUAL       ---CAPACITY---
                                                    CLASS       CONSUMED       MIN      MAX
  HIGHPRTY    Very high priority work                                 82       5000   999999
                                                  IBMUSER             82
                                                  MCGRATH              0

  REGTSO      Non-priority TSO work                                  461       2000   500000
                                                  FASTPER3             0
                                                  TSOCLASS           461
```

Figure 5-51. WLMGL Report - Service Policy Page

Field descriptions of all reports

Table 5-65. Fields in the Workload Activity Report

Field Heading	Meaning
MODE	This field indicates the WLM operating mode.
POLICY ACTIVATION DATE/TIME	Time and date when this policy was activated
Report Header	These lines separate the reporting levels from each other. **POLICY** Policy name **WORKLOAD** Workload name **SERVICE CLASS** Service class name **REPORT CLASS** Report class name **RESOURCE GROUP** Resource group name The field contains *NONE if a resource group name was not specified for that service class. **PERIOD** Period number **IMPORTANCE** Importance of a service class period. **CRITICAL** Indicates whether a service class is specified as: **CPU** CPU critical **STORAGE** Storage critical **STORAGE + CPU** Both storage and CPU critical **NONE** Neither storage nor CPU critical **Note:** 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).
PGN	The performance group number that describes the rate at which service is to be allocated to the user. The report can show a range of PG numbers to eliminate printing of consecutive lines of zeros.
PGP	The performance group period that is a portion of the elapsed transaction time, such as the average service used by a transaction during the first 10 seconds of a total transaction time of 30 seconds. The report can show a range of PGP numbers to eliminate printing of consecutive lines of zeros.
DMN	The identification of the domain associated with the performance group during a performance group period. The range of values is 0 to 128. If no domain number was defined for the performance group period, two asterisks (**) appear in this field.
TSG	The identification of the time slice group associated with the performance group period. The range of values is from 1 to 16. Two asterisks (**) in this field indicate that no time slice group has been defined for the performance group period.

Table 5-65. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
TRANSACTIONS	Number of transactions. **AVG** 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. To get the average number of transactions related to active address spaces, subtract **AVG ENC** from **AVG**. **MPL** 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. To get the average number of transactions related to resident address spaces, subtract **AVG ENC** from **MPL**. **ENDED** 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. **END/SEC** The number of transactions that ended per second. **#SWAPS** 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. **EXCTD** Count of times a subsystem work manager reported that an execution phase has completed. A single transaction could have zero or more execution phases.
AVG ENC	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.
REM ENC	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.
MS ENC	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.

Table 5-65. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
TRANS.-TIME	Transaction time in HHH.MM.SS.TTT units. All times are reported in the period the transaction ended. The time a job was delayed due to TYPRUN=HOLD or TYPRUN=JCLHOLD is NOT included in any of the transaction times. **ACTUAL** 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 AFFINITY, and INELIGIBLE. **EXECUTION** The average execution time of ended transactions. **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. **R/S AFFINITY** 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. **INELIGIBLE** Average time the job was delayed due to operational delays or JES scheduling delays, examples are: • Job held by operator • Job class or job queue held • Duplicate jobname serialization • Job class execution limits **CONVERSION** Average time the job was delayed due to 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 included in the total response time (ACTUAL). **STD DEV** Standard deviation of ACTUAL. 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.

Table 5-65. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
DASD I/O	Information about DASD I/O activities. **SSCHRT** Number of start subchannels SSCH per second in the reported interval. **RESP** 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). The following three values are shown as reported by the channel measurement subsystem. **CONN** Average DASD connect time of the transactions in this group. **DISC** Average DASD disconnect time of the transactions in this group. **Q+PEND** Average DASD wait time (queue time + pending time) of the transactions in this group. This does not include IOSQ time. **IOSQ** Average time the transactions in this group spent on the IOS queue, based on sampled delays.
SERVICE	The service is calculated by multiplying the received service units with the appropriate service definition coefficient. **IOC** Total amount of input/output service received. **CPU** Total amount of task and preemptible-class SRB processor service received. **MSO** Total amount of main storage occupancy service received. **SRB** Total amount of non-preemptible SRB service received. **TOT** Sum of CPU, SRB, IOC, and MSO service. **/SEC** Rate at which service is provided in service units per second.

Table 5-65. Fields in the Workload Activity Report (continued)

Field Heading	Meaning

SERVICE RATES	This category is made up of the following:
	ABSRPTN
	Rate at which service is used while transactions are resident in main storage. This is the total service divided by the transaction residency time.
	TRX SERV
	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.
	TCB Task control block time used in seconds.
	SRB Service request block time in seconds.
	RCT Region control task time in seconds.
	IIT I/O interrupt time in seconds.
	HST Hiperspace service time in seconds.
	APPL%
	A percentage of CPU time used by transactions running in the service or report class period.
	The calculation is:
	`TCB + SRB + RCT + IIT + HST`
	`APPL % = -------------------------- * 100`
	`Interval`
	Where the interval length in a sysplex is the common interval length.
	Notes:
	1. APPL% 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 processors.
	2. In a sysplex, the values for seconds and CPU time percentage are meaningful only if all processors have the same speed.

EX VEL %	The execution velocity is a measure of the acceptable processor and storage delays while work is running. For details about the execution velocity, see “Common Monitor III report measurements” on page 2-6.

Table 5-65. Fields in the Workload Activity Report (continued)

Field Heading	Meaning	
PAGE-IN RATES	**SINGLE**	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.
	BLOCK	Rate of demand page-ins from DASD for blocked pages, expressed in pages per seconds.
	SHARED	Rate of shared storage page-ins
	HSP	Rate of standard hiperspace pages read into central storage from auxiliary storage.
	HSP MISS	Rate of expanded-storage-only (ESO) hiperspace reads due to lack of expanded storage.
	EXP SNGL	Rate of demand page-moves of single pages from expanded storage.
	EXP BLK	Rate of pages moved from expanded storage as a result of a page fault that was part of a block.
	EXP SHR	Rate of shared storage page-ins from expanded storage into central storage.
	Note: Enclave transaction residency time is not included in the calculation of these values because there is no paging on behalf of enclaves.	
STORAGE	Amount of storage frames.	
	AVG	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. **Note:** Enclave transaction residency or active time is not included in the calculation of this value.
	TOTAL	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.
	CENTRAL	Total number of central storage frames allocated to resident ASIDs. This value is the number of central storage frames weighted by the transaction residency time for each active ASID, divided by the RMF interval time.
	EXPAND	Total number of expanded storage frames allocated to resident ASIDs. This value is the number of expanded storage frames weighted by the transaction residency time for each active ASID, divided by the RMF interval time.
	SHARED	Total number of shared storage pages allocated to resident ASIDs.

Table 5-65. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
VELOCITY MIGRATION	The following two values assist you in better planning your velocity goals. **I/O MGMT** I/O Priority Management Value of achieved execution velocity including I/O using and delay samples. 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. You see your current definition in the Service Policy page (see Figure 5-51 on page 5-158). **INIT MGMT** Initiator Management Value of achieved execution velocity including batch initiator delay samples. 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. You find a description of these delays in this table for the field TRANS.-TIME.
Service Classes being Served	
SERVICE CLASSES BEING SERVED	This section is only available when the service class is a server. Each service class name being served by the reported service class (see name in the separation line) is displayed.
Work Manager/Resource Manager In *z/OS RMF Performance Management Guide*, in chapter ***Analyzing Sysplex Activity***, you can find some scenarios on how to use and interpret the WLMGL report.	
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 **BTE** indicates the representation of the states incurred in the begin-to-end phase of a transaction **EXE** indicates the representation of the states incurred in the execution phase of a transaction.
RESP TIME (%)	The transaction response time percentage in either the ***BEGIN-TO-END*** phase, or the ***EXECUTION*** phase.

Table 5-65. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
STATE SAMPLES BREAKDOWN (%)	Identifies the percentages of samples that a transaction has been detected in the reported states. **ACTIVE SUB** 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. **ACTIVE APPL** The active application state sample percentage in contrast to the active subsystem state sample percentage. This allows a subsystem to differentiate between work requests procressed by the subsystem itself (ACTIVE SUB) and work requests processed by an application invoked by the subsystem. **READY** 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. **IDLE** The idle state sample percentage. Idle indicates that no work request (or transaction) is allowed to run.
STATE SAMPLES BREAKDOWN (%) - WAITING FOR	STATE SAMPLES BREAKDOWN (%) - continued. This category presents the delay reason states with the fourteen highest non-zero values. These are sorted by the sum of the BTE and EXE rows in each column. The remaining less important values will be accumulated and presented as delay reason OTHR. Below is a list explaining the delay reasons. **LOCK** waiting for lock. **I/O** 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. **CONV** waiting for conversation could have been used in conjunction with IWMMSWCH to identify where the target is located. **DIST** 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. **LOCL** waiting for a session to be established locally, for example, on the current MVS image. **SYSP** waiting for a session to be established somewhere in the sysplex. **REMT** waiting for a session to be established somewhere in the network. **TIME** waiting for timer. **LTCH** waiting for a latch. **PROD** waiting for another product. **MISC** waiting for unidentified resource, possibly among another specific category, but which may not be readily determined. **SSLT** waiting for an SSL thread. **REGT** waiting for a regular thread. **WORK** waiting for registration to a work table.

Table 5-65. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
STATE SWITCHED SAMPL(%)	Subsystem state samples - continued **LOCAL** 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. **SYSPL** 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. **REMOT** 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.
Service or Report Class Period	
GOAL versus ACTUALS	The goal specified for a service class period is shown compared against the achieved values. As goal for a service class period either response time or execution velocity can be specified. In a report for a homogeneous report class period, this is the goal of the corresponding service class period. If measurement data for more than one system is available, the ACTUALS section starts with a *ALL line with the sysplex view. It is followed by one line for each system. **RESPONSE TIME** GOAL - if specified: the expected amount of time required to complete the work submitted. ***-AVG*** indicates that response time was specified as an average, while a percentage, for example ***90%***, indicates a target for a given percentile. A blank indicates that no specification was made. ACTUALS - represents the measured value in the same terms as the specified GOAL, either average or percentile. **EX VEL%** GOAL - if specified: the execution velocity that defines the acceptable amount of delay for work relative to the total execution time. A blank indicates that no specification was made. ACTUALS - represents the measured value. The execution velocity is a measure of the acceptable processor and storage delays while work is ready to run. For details about the execution velocity, see “Common Monitor III report measurements” on page 2-6. **PERF INDX** GOAL - not applicable ACTUALS - the performance index for a period represents how close a period came to reaching the goal (PI is 1.0), and how much this period suffered versus its goal. See Table 2-72 on page 2-194 for more details about the performance index.
AVG ADRSP	Average number of address spaces and enclaves that contributed delay and using samples to this class.
USING CPU	CPU using samples in percent.
USING I/O	I/O using samples in percent.

Table 5-65. Fields in the Workload Activity Report (continued)

Field Heading	Meaning	
EXECUTION DELAYS %	General execution delays included in TOTAL. Each dispatchable unit sampled can increase one of the CPU or paging delay samples. Besides the TOTAL value, only the nine highest values contributing to TOTAL will be shown. The remaining less important values will be accumulated and presented as OTHR.	
	TOTAL	Total delay used by SRM in its execution velocity calculation.
	CPU	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.
	I/O	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 page 5-162 for a description). **Note:** It depends on the definition in WLM whether this value is part of the TOTAL value or not, by default it is **not** contained in TOTAL.
	CAPP	CPU capping delay. A TCB or SRB is marked non-dispatchable because • a resource group maximum is being enforced • 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 *z/OS MVS Planning: Workload Management*). This value is NOT part of the CPU delay.
	SWIN	Swap-In delay. Swap-In has started but not completed.
	MPL	MPL delay. Ready but swap-in has not started.
	QMPL	Queue MPL - work is waiting for a server address space or batch initiator.
	SERV PRIV	Private area paging delay for a server address space.
	SERV VIO	VIO paging delay for a server address space.
	SERV SHSP	Hiperspace paging delay for a server address space.
	SERV SWIN	Swap-in delay for a server address space.
	SERV MPL	MPL delay for a server address space.
	AUX PRIV	Auxiliary paging from private.
	AUX COMM	Auxiliary paging from common.
	AUX XMEM	Auxiliary paging from cross memory.
	AUX VIO	Auxiliary paging from VIO.
	AUX SHSP	Auxiliary paging from standard hiperspaces.
	AUX EHSP	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).

Table 5-65. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
DLY%	The following states are NOT included in the TOTAL EXECUTION DELAYS. **UNKN** State is unknown. The address space or enclave was not found to be using or delayed for any sampled resource, but MVS has not been notified that it is idle. **IDLE** Idle state. Work is in STIMER wait, TSO terminal wait, APPC wait, OMVS input or output wait, or an initiator is waiting for work.
CRYPTO%	**DLY** Crypto delay state — a TCB or SRB was found to be waiting for a cryptographic asynchronous message processor (CAP), an adjunct processor (AP) or a processor feature queue. **USG** Crypto using state — a TCB or SRB was found to be using a CAP or AP.
% QUIE	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.
Response Time Distribution (only available for service and report class periods when a response time GOAL was specified)	
RESPONSE TIME DISTRIBUTION	WLM maintains counts of how many transactions were completed within a particular time. The ***response time goal*** defined for each service class period is split into 14 response time buckets where: • bucket 1 covers the gap from 0 to half the goal • buckets 2 to 11 cover the gap between half the goal to 1.5 times the goal evenly divided • bucket 12 covers two times the goal • bucket 13 covers four times the goal • bucket 14 covers the gap from four times the goal to infinity The chart presents the • number of total (ended) transactions, • response time, • number of total (ended) transactions in percent, • and a graphical illustration of the percentage. **TIME** Response time associated to this bucket. **NUMBER OF TRANSACTIONS** Number of transactions that completed for this period. **CUM TOTAL** Cumulative number of transactions so far **IN BUCKET** Number of transactions in this bucket **PERCENT** Percentage **CUM TOTAL** Cumulative percentage of transactions so far **IN BUCKET** Percentage of transactions associated to the bucket **Percent Scale** Graphical presentation of each bucket.

PP - WLMGL Report

Table 5-66. Fields in the WLMGL Report - POLICY

Field Heading	Meaning
	Service Policy Page
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 DEFINITIONS COEFFICIENTS	Service definitions coefficients as defined in the service policy to determine the: **IOC** Number of countable EXCP instructions. **CPU** Task processor time. **SRB** SRB processor time. **MSO** Approximate storage use for each service class period.
POLICY	Policy name and description.
I/O PRIORITY MANAGEMENT	**YES** 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. **NO** I/O priorities are the same as dispatching priorities.
DYNAMIC ALIAS MANAGEMENT	**YES** if WLM dynamic management of parallel access volumes is active, otherwise **NO**.
SYSTEMS	The system names contributing to the sysplex report. **ID** System name. **OPT** IEAOPTxx Parmlib member suffix (parameter description that control resource and workload management algorithms in the system resources manager). **SU/SEC** Service units per second per online CPU. **TIME** Begin time of the interval for this system. **INTERVAL** Interval length for this system in HH.MM.SS.

Table 5-66. Fields in the WLMGL Report - POLICY (continued)

Field Heading	Meaning
RESOURCE GROUPS	A resource group is an amount of processor capacity across one or more MVS systems. The report includes resource group name and description as: **NAME** Resource group name. **DESCRIPTION** Resource group name description. **SERVICE CLASS** Service class name associated with this RESOURCE GROUP. **ACTUAL CONSUMED** Unweighted CPU and SRB service units consumed all service classes in the resource group. Unweighted CPU and SRB service units consumed by each service class. **CAPACITY** Resource group limits in unweighted CPU and SRB service units. **MIN** Minimum amount of service across the sysplex that the resource group should receive if demand exists. **MAX** Maximum amount of service across the sysplex that the resource group should be allowed to consume.

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; they apply to the WKLD report (Exception condition in compatibility mode) and to the WLMGL report (Overview condition in goal mode).

Note: For all fields that belong exclusively to the WLMGL report, the following table entries are valid only in systems running in goal mode.

Table 5-67. Exception and Overview Names in the Workload Activity Report

Field Heading	Exception/Overview Name
TRANSACTIONS -	
- AVG	TRANSAVG
- MPL	TRANSMPL
- ENDED	TRANSTOT
- END/SEC	TRANS
- #SWAPS	SPERTRA
- AVG ENC	ENCAVG
- REM ENC	ENCREM
- MS ENC	ENCMS
TRANSACTION TIME -	
- ACTUAL	RTIMETOT
- EXECUTION	RTIME
- QUEUED	RTIMEQUE
- R/S AFFINITY	TRANSADT

Table 5-67. Exception and Overview Names in the Workload Activity Report (continued)

Field Heading	Exception/Overview Name
- INELIGIBLE	TRANSIQT
- CONVERSION	TRANSCVT
DASD I/O	
- SSCHRT	SSCHRT
- RESP	RESP
- CONN	CONN
- DISC	DISC
- Q+PEND	QPEND
- IOSQ	IOQS
INTERVAL SERVICE -	
- IOC	IOSRV
- CPU	CPUSRV
- MSO	MSOSRV
- SRB	SRBSRV
- TOT	TOTSRV
SERVICE RATES -	
- ABSRPTN	ABSRPTN
- TRX SERV	TRXSERV
- TCB	TCBSEC/TCBPER
- SRB	SRBSEC/SRBPER
- RCT	RCTSEC
- IIT	IITSEC
- HST	HSTSEC
- APPL%	APPLSEC/APPLPER
EX VEL %	EXVEL
Total number of EXCPs	EXCP
EXCP rate	EXCPRT
PAGE-IN RATES -	
- SINGLE	SINGLE
- BLOCK	BLOCK
- SHARED	SHARED
- HSP	HSP
- HSP MISS	HSPMISS
- EXP SNGL	EXPSNGL
- EXP BLK	EXPBLK
- EXP SHR	EXPSHR
STORAGE -	
- TOTAL	STOTOT
- CENTRAL	STOCEN
- EXPAND	STOEXP

Table 5-67. Exception and Overview Names in the Workload Activity Report (continued)

Field Heading	Exception/Overview Name
- SHARED	STOSHR
GOAL versus ACTUALS (Actuals) -	
- PERF INDEX	PI
USING CPU	CPUUSGP
USING I/O	IOUSGP
EXECUTION DELAYS % -	
- CPU	CPUDLYP
- I/O	IODLYP
- CAPP	CAPP
- SWIN	SWINP
- MPL	MPLP
- QMPL	QUEUEP
Total Server Delays	SERVP
Total Storage Delays	STOP
DLY% -	
- UNKN	UNKP
- IDLE	IDLEP
CRYPTO% -	
- DLY	CRYDLYP, CAPDLYP, APDLYP, FQDLYP
- USG	CRYUSGP, CAPUSGP, APUSGP

XCF - Cross-System Coupling Facility Activity report

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 a report, specify
`REPORTS(XCF)`

Contents of the report

The 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.

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 5-52 on page 5-175.

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 5-53 on page 5-176.

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 5-55 on page 5-177.

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
                                                                                                                PAGE    1
          z/OS V1R4              SYSTEM ID   SYSM          DATE 03/30/2002         INTERVAL 16.00.000
                                 RPT VERSION V1R2 RMF     TIME 09.15.00           CYCLE 1.000 SECONDS

                                                XCF USAGE BY SYSTEM

-------------------------------------------------------------------------------------------------------------------------------------------
                                                REMOTE SYSTEMS                                                                  LOCAL
-------------------------------------------------------------------------------------------------------------------------------  ----------
                                OUTBOUND FROM SYSM                                                    INBOUND TO SYSM           SYSM
-------------------------------------------------------------------------------------------------  ----------------------------  ----------

                                                 ----- BUFFER -----      ALL
TO          TRANSPORT  BUFFER        REQ   %    %    %    %     PATHS     REQ   FROM            REQ     REQ     TRANSPORT     REQ
SYSTEM      CLASS      LENGTH        OUT  SML  FIT  BIG  OVR  UNAVAIL  REJECT   SYSTEM           IN  REJECT     CLASS      REJECT

SYSN        BIG        40,892        114    0  100    0    0        0       1   SYSN            648       0     BIG             0
            DEFAULT       956        405    0  100    0    0      405       0                                   DEFAULT         0
            FEWFAST       956          0    0    0    0    0        0       0                                   FEWFAST         0
            TCCONS        956          0    0    0    0    0        0       0                                   TCCONS          0
            TCGRS         956      3,098    0  100    0    0        0       0                                   TCGRS           0
            TCRMF      62,464          3    0  100    0    0        0       0                                   TCRMF           0
            TCVLF         956          0    0    0    0    0        0       0                                   TCVLF           0
SYSO        BIG        40,892        151    0  100    0    0        0       0   SYSO          3,738       0
            DEFAULT       956        366    0  100    0    0      366       0
            FEWFAST       956          0    0    0    0    0        0       0
            TCCONS        956          0    0    0    0    0        0       0
            TCGRS         956          0    0    0    0    0        0       0
            TCRMF      62,464          0    0    0    0    0        0       0
            TCVLF         956          0    0    0    0    0        0       0
                                  ----------                                                  ----------
TOTAL                              4,137                                      TOTAL          4,386
```

Figure 5-52. XCF Activity Report - Usage by System

Table 5-68. 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 XCF uses for the message transfer. There is one line in the report for each transport class for each target system.
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. The value is reported as `'>99'` if the percentage is below 100 but rounded to 100.
% 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.
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.

Table 5-68. Fields in the XCF Activity Report - Usage by System (continued)

Field Heading	Meaning
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 XCF uses for the message transfer. There is one line in the report for each transport class for each target system.

```
                                        X C F  A C T I V I T Y
                                                                                                                    PAGE    2
          z/OS V1R4              SYSTEM ID   SYSM          DATE 03/30/2002          INTERVAL 16.00.000
                                 RPT VERSION V1R2 RMF      TIME 09.15.00            CYCLE 1.000 SECONDS

                                        XCF USAGE BY MEMBER
-------------------------------------------------------------------------------------------------------------------------------

             MEMBERS COMMUNICATING WITH SYSM                                          MEMBERS ON SYSM
-------------------------------------------------------------------                 ---------------------------------------------

                                                 REQ          REQ
                                                FROM           TO                                                 REQ          REQ
GROUP       MEMBER              SYSTEM          SYSM         SYSM          GROUP       MEMBER                     OUT           IN
                                                                          A000000F    RECEIVER               *NOW ACTIVE
BOERMF22    BOERMF22$SYSN       SYSN               0            0          BOERMF22    BOERMF22$SYSM                0            0
            BOERMF22$SYSO       SYSO               0            0                                               ----------   ----------
                                              ----------   ----------      TOTAL                                        0            0
TOTAL                                            0            0

B000000C    MEMBER1             SYSN        *NOW ACTIVE
B0000001    RECEIVER            SYSO        *NOW INACTIVE
            SENDER              SYSN        *NOW INACTIVE
COFVLFNO    SYSN                SYSN               0            6          COFVLFNO    SYSM                         0            6
            SYSO                SYSO               0            0                                               ----------   ----------
                                              ----------   ----------      TOTAL                                        0            6
TOTAL                                            0            6

CTTXGRP     CTTX####SYSN        SYSN              88          106          CTTXGRP     CTTX####SYSM               350          418
            CTTX####SYSO        SYSO             109          159                                               ----------   ----------
                                              ----------   ----------      TOTAL                                      350          418
TOTAL                                          197          265

C0000003    MEMBERA             SYSO        *NOW ACTIVE
            MEMBERB             SYSO        *NOW ACTIVE
C0000004    RECEIVER            SYSN        *NOW ACTIVE                  C0000004    SENDER                 *NOW ACTIVE
IXCLO00A    M112                SYSN               0            0          IXCLO00A    M111                         0            0
            M113                SYSO               0            0                                               ----------   ----------
                                              ----------   ----------      TOTAL                                        0            0
TOTAL                                            0            0

IXCLO012    M81                 SYSN             344          350          IXCLO012    M80                        890          887
            M82                 SYSO             302          293                                               ----------   ----------
                                              ----------   ----------      TOTAL                                      890          887
TOTAL                                          646          643
```

Figure 5-53. XCF Activity Report - Usage by Member

Table 5-69. 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.	There is one line per member in the report.
REQ IN	The number of messages received by the member on the local system.	

```
                                        X C F   A C T I V I T Y                                                 PAGE    3
          z/OS V1R4               SYSTEM ID   SYSM          DATE 03/30/2002          INTERVAL 16.00.000
                                  RPT VERSION V1R2 RMF      TIME 09.15.00            CYCLE 1.000 SECONDS

TOTAL SAMPLES = 768                                  XCF PATH STATISTICS
-------------------------------------------------------------------------------------------------------------------------
                           OUTBOUND FROM SYSM                                          INBOUND TO SYSM
-------------------------------------------------------------------------------   -----------------------------------------
       T FROM/TO                                                                        T FROM/TO
TO     Y DEVICE, OR      TRANSPORT       REQ     AVG Q                          FROM    Y DEVICE, OR       REQ BUFFERS
SYSTEM P STRUCTURE       CLASS           OUT     LNGTH    AVAIL    BUSY   RETRY  SYSTEM  P STRUCTURE         IN UNAVAIL

SYSN   S IXCGRS          TCGRS           581      0.72      557       2      0   SYSN    S IXCGRS           418       0
       S IXCPLEX_PATH1   BIG           2,674      0.72    1,655     1,01     0           S IXCPLEX_PATH1  2,783       0
       S IXCPLEX_PATH2   TCGRS         2,522      0.72    2,509       1      0           S IXCPLEX_PATH2    426       0
       S IXCPLEX_PATH3   TCCONS            4      0.72        4              0           S IXCPLEX_PATH3  2,099       0
       S IXCPLEX_PATH4   TCRMF             3      0.72        2              0           S IXCPLEX_PATH4    419       0
       S IXCVLF          TCVLF             0      0.00        0              0           S IXCVLF           431       0
```

Figure 5-54. XCF Activity Report - Path Statistics (Coupling Facility Structure)

```
                                        X C F   A C T I V I T Y                                                 PAGE    3
          z/OS V1R4               SYSTEM ID   SYS1          DATE 03/14/2002          INTERVAL 15.00.000
                                  RPT VERSION V1R2 RMF      TIME 16.30.00            CYCLE 1.000 SECONDS

TOTAL SAMPLES = 900                                  XCF PATH STATISTICS
-------------------------------------------------------------------------------------------------------------------------
                           OUTBOUND FROM SYS1                                          INBOUND TO SYS1
-------------------------------------------------------------------------------   -----------------------------------------
       T FROM/TO                                                                        T FROM/TO
TO     Y DEVICE, OR      TRANSPORT       REQ     AVG Q                          FROM    Y DEVICE, OR       REQ BUFFERS
SYSTEM P STRUCTURE       CLASS           OUT     LNGTH    AVAIL    BUSY   RETRY  SYSTEM  P STRUCTURE         IN UNAVAIL

SYS3   C 0CA0 TO 0CA0    DEFAULT      91,864      0.04   90,789   1,075      0   SYS3    C 0CA1 TO 0CA1   84,966       0
       C 0C80 TO 0C80    DEFAULT      86,081      0.02   84,990   1,091      0           C 0C81 TO 0C81   78,088       0
       C 0EA0 TO 0EA0    DEFAULT      87,490      0.03   86,423   1,067      0           C 0EA1 TO 0EA1   82,012       0
       C 0E80 TO 0E80    DEFAULT      77,021      0.02   75,955   1,066      0           C 0E81 TO 0E81   94,293       0
                                    ---------                                                              -------
TOTAL                                342,456                                    TOTAL                   339,359
```

Figure 5-55. XCF Activity Report - Path Statistics (Channel-to-Channel)

PP - XCF Report

Table 5-70. 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	**CTC:** 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. **Coupling facility:** The coupling facility structure name. UNK* TO appears in this field when one or more paths are unknown at begin 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 LNGTH	The average number of messages queued for data transfer over each outbound signalling path. The calculation is: `AVG Q LNGTH = # Message Entries for this Device / # 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.
FROM SYSTEM	The name of the system from which the messages are sent.
FROM/TO DEVICE, OR STRUCTURE	**CTC:** 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. **Coupling facility:** The coupling facility structure name. UNK* TO appears in this field when one or more paths are unknown at begin of the interval. The counter `AVG Q LNGTH` contains data, all other fields are blank.
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.

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*.

Exception report

How to request this report

The scope of data that can be reported in the Exception report depends on the data being gathered by Monitor I gatherer sessions.

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

An Exception report presents a summary of the values that exceeded installation-defined thresholds over a specific period of time.

Generating a 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. For a description of the exception control statements see the *z/OS RMF User's Guide*.

Report description

Figure 5-56 on page 5-180 and Figure 5-57 on page 5-181 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.

PP - Exception Report

```
                                    R M F   E X C E P T I O N   R E P O R T
                                                                                                    PAGE  001
           z/OS V1R4                SYSTEM ID   MVS1           START 03/04/2002-00.02.25  INTERVAL 00.29.59
                                    RPT VERSION V1R2 RMF       END   03/04/2002-08.02.26  CYCLE 1.000 SECONDS

NUMBER OF INTERVALS 16            TOTAL LENGTH OF INTERVALS 07.59.52
CONTROL   DATE    TIME    INT  EXCEPTION       ACTUAL
STMT NAME MM/DD HH.MM.SS MM.SS THRESHOLD       VALUE  EXCEPTION DESCRIPTION    NAME VALUE
--------- ----- -------- ----- ------------ --------- ------------------------ ----------

CPU1      03/04 03.02.26 29.59 LE 3             2.453 PERCENT CPU BUSY        -1
CPU2      03/04 03.02.26 29.59 LE 3             2.315 PERCENT CPU BUSY        -2
CPU3      03/04 03.02.26 29.59 LE 3             2.301 PERCENT CPU BUSY        -3
CPU4      03/04 03.02.26 29.59 LE 3             2.391 PERCENT CPU BUSY        -4
CPU5      03/04 03.02.26 29.59 LE 3             2.858 PERCENT CPU BUSY        -5
CPU1      03/04 03.32.26 29.59 LE 3             2.637 PERCENT CPU BUSY        -1
CPU2      03/04 03.32.26 29.59 LE 3             2.502 PERCENT CPU BUSY        -2
CPU3      03/04 03.32.26 29.59 LE 3             2.165 PERCENT CPU BUSY        -3
CPU4      03/04 03.32.26 29.59 LE 3             2.279 PERCENT CPU BUSY        -4
CPU1      03/04 04.02.26 30.00 LE 3             2.690 PERCENT CPU BUSY        -1
CPU2      03/04 04.02.26 30.00 LE 3             2.233 PERCENT CPU BUSY        -2
CPU3      03/04 04.02.26 30.00 LE 3             2.350 PERCENT CPU BUSY        -3
CPU4      03/04 04.02.26 30.00 LE 3             2.191 PERCENT CPU BUSY        -4
CPU1      03/04 04.32.26 29.59 LE 3             2.830 PERCENT CPU BUSY        -1
CPU2      03/04 04.32.26 29.59 LE 3             2.345 PERCENT CPU BUSY        -2
CPU3      03/04 04.32.26 29.59 LE 3             2.373 PERCENT CPU BUSY        -3
CPU4      03/04 04.32.26 29.59 LE 3             2.274 PERCENT CPU BUSY        -4
CPU1      03/04 05.02.26 29.59 LE 3             2.783 PERCENT CPU BUSY        -1
CPU2      03/04 05.02.26 29.59 LE 3             2.505 PERCENT CPU BUSY        -2
CPU3      03/04 05.02.26 29.59 LE 3             2.257 PERCENT CPU BUSY        -3
CPU4      03/04 05.02.26 29.59 LE 3             2.405 PERCENT CPU BUSY        -4
CPU4      03/04 05.32.26 29.59 LE 3             2.548 PERCENT CPU BUSY        -4
CPU3      03/04 06.32.26 30.00 LE 3             2.811 PERCENT CPU BUSY        -3
CPU4      03/04 06.32.26 30.00 LE 3             2.861 PERCENT CPU BUSY        -4
CPU4      03/04 07.02.26 29.59 LE 3             2.988 PERCENT CPU BUSY        -4
```

Figure 5-56. Exception Report - Low CPU Utilization

Example

If you want to produce an Exception report for selected performance data (that is, average TSO response times) during the prime shift (between 8 a.m. and 6 p.m.), you can use the following report option statements:

```
ETOD(0800,1800)
EXCEPT(TSO(RTIME(0020),GE,0))
```

Note: The sample report assumes that TSO is running in performance group 2. Also, note that a response time greater than zero is not a *real exception*, but it can be used to create the desired report.

```
                                    R M F   E X C E P T I O N   R E P O R T
                                                                                                    PAGE  001
           z/OS V1R4                SYSTEM ID   MVS1           START 03/04/2002-08.02.26  INTERVAL 00.29.59
                                    RPT VERSION V1R2 RMF       END   03/04/2002-18.02.26  CYCLE 1.000 SECONDS

NUMBER OF INTERVALS 20            TOTAL LENGTH OF INTERVALS 09.59.51
CONTROL   DATE    TIME    INT  EXCEPTION       ACTUAL
STMT NAME MM/DD HH.MM.SS MM.SS THRESHOLD      VALUE  EXCEPTION DESCRIPTION    NAME VALUE
--------- ----- -------- ----- ------------ --------- ------------------------ ----------
TSO       03/04 08.02.26 29.59 GE 0            1.472 RESPONSE TIME           -00020
TSO       03/04 08.32.26 30.00 GE 0            1.685 RESPONSE TIME           -00020
TSO       03/04 09.02.26 29.59 GE 0            1.268 RESPONSE TIME           -00020
TSO       03/04 09.32.26 30.00 GE 0            0.753 RESPONSE TIME           -00020
TSO       03/04 10.02.26 29.59 GE 0            1.161 RESPONSE TIME           -00020
TSO       03/04 10.32.26 29.59 GE 0            0.900 RESPONSE TIME           -00020
TSO       03/04 11.02.26 30.00 GE 0            1.884 RESPONSE TIME           -00020
TSO       03/04 11.32.26 30.00 GE 0            0.501 RESPONSE TIME           -00020
TSO       03/04 12.02.26 29.59 GE 0            0.410 RESPONSE TIME           -00020
TSO       03/04 12.32.26 30.00 GE 0            0.820 RESPONSE TIME           -00020
TSO       03/04 13.02.26 30.00 GE 0            1.016 RESPONSE TIME           -00020
TSO       03/04 13.32.26 29.59 GE 0            0.610 RESPONSE TIME           -00020
TSO       03/04 14.02.26 30.00 GE 0            0.671 RESPONSE TIME           -00020
TSO       03/04 14.32.26 29.59 GE 0            0.594 RESPONSE TIME           -00020
TSO       03/04 15.02.26 29.59 GE 0            1.449 RESPONSE TIME           -00020
TSO       03/04 15.32.26 30.00 GE 0            2.716 RESPONSE TIME           -00020
TSO       03/04 16.02.26 30.00 GE 0            0.822 RESPONSE TIME           -00020
TSO       03/04 16.32.26 30.00 GE 0            1.793 RESPONSE TIME           -00020
TSO       03/04 17.02.26 29.59 GE 0            0.572 RESPONSE TIME           -00020
TSO       03/04 17.32.26 30.00 GE 0            0.404 RESPONSE TIME           -00020
```

Figure 5-57. Exception Report - TSO Response Time

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 identifier associated with 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 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.

PP - Exception Report

Table 5-71. 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. **Note:** 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(TSO001IO(DART('TSO001'),GE,3))
EXCEPT(TSO002IO(DART('TSO002'),GE,3))
EXCEPT(TSO004IO(DART('TSO004'),GE,3))
EXCEPT(TSO009IO(DART('TSO009'),GE,3))
EXCEPT(TSO013IO(DART('TSO013'),GE,3))
```

The exception-condition name DRTAVG specifies the average response time:

```
EXCEPT(TSO001RT(DRTAVG('TSO001'),GE,25))
EXCEPT(TSO002RT(DRTAVG('TSO002'),GE,25))
EXCEPT(TSO004RT(DRTAVG('TSO004'),GE,25))
EXCEPT(TSO009RT(DRTAVG('TSO009'),GE,25))
EXCEPT(TSO013RT(DRTAVG('TSO013'),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:

```
                                   R M F   O V E R V I E W   R E P O R T
                                                                                                                PAGE  001
           z/OS V1R4               SYSTEM ID   SYS1           START 03/14/2002-08.02.26  INTERVAL 00.29.59
                                   RPT VERSION V1R2 RMF       END   03/14/2002-18.02.26  CYCLE 1.000 SECONDS

NUMBER OF INTERVALS 20          TOTAL LENGTH OF INTERVALS 09.59.51

DATE    TIME    INT   TSO001IO  TSO002IO  TSO004IO  TSO009IO  TSO013IO  TSO001RT  TSO002RT  TSO004RT  TSO009RT  TSO013RT
MM/DD HH.MM.SS MM.SS

03/14 08.02.26 29.59                                   4.309                                 28.329
03/14 09.02.26 29.59                       18.592                            26.368
03/14 11.32.26 30.00                                                                       27.843
03/14 12.02.26 29.59                                 6.458     9.962                  33.949
03/14 12.32.26 30.00                                                                       26.709
03/14 13.02.26 30.00                                                                                 29.427
03/14 14.02.26 30.00                                                           32.073                26.733
03/14 14.32.26 29.59     4.826                                                                 34.448
03/14 15.02.26 29.59               3.134
03/14 15.32.26 30.00                                                                       28.152
03/14 16.02.26 30.00             13.403
03/14 16.32.26 30.00                                                                                 26.077
```

Figure 5-58. Overview Report - Exception Version

The reporting range covers 20 intervals, but you see in Figure 5-58 that only intervals with at least one exception value are listed in the report.

PP - Exception Report

Table 5-72. 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.
FFFFFFFF	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

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 printable report, specify

```
OVERVIEW(REPORT)
```

To create data records either for spreadsheet processing or other applications, specify:

```
OVERVIEW(RECORD)
```

Report description

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

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 5-67 on page 5-171.

Please, note that you cannot specify generic exception classes like DASD without a qualifier. This is valid for exception reporting, but not for overview reporting. Here, you have to specify explicitly a qualifier, which could be for DASD 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 5-202). 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 TOTSRV specifies the total service units, the qualifier S.TSOSERV refers to service class TSOSERV, and suboption NOSYSTEMS defines sysplex reporting:

```
OVW(SERVUNIT(TOTSRV(S.TSOSERV)),NOSYSTEMS)
```

The exception-condition name RTIMETOT specifies the average response time:

```
OVW(RTIMEP1(RTIMETOT(S.TSOSERV.1)),NOSYSTEMS)
OVW(RTIMEP2(RTIMETOT(S.TSOSERV.2)),NOSYSTEMS)
OVW(RTIMEP3(RTIMETOT(S.TSOSERV.3)),NOSYSTEMS)
```

With the exception-condition name PI, you specify the performance index:

```
OVW(PIP1(PI(S.TSOSERV.1)),NOSYSTEMS)
OVW(PIP2(PI(S.TSOSERV.2)),NOSYSTEMS)
```

The exception-condition name TRANS specifies the transaction rate:

```
OVW(TRXP1(TRANS(S.TSOSERV.1)),NOSYSTEMS)
OVW(TRXP2(TRANS(S.TSOSERV.2)),NOSYSTEMS)
OVW(TRXP3(TRANS(S.TSOSERV.3)),NOSYSTEMS)
```

The *summary version* of the Overview report looks like this:

```
                                   R M F   O V E R V I E W   R E P O R T
                                                                                                  PAGE  001
           z/OS V1R4               SYSPLEX ID UTCPLXHD        START 04/29/2002-10.00.00  INTERVAL 00.10.00
                                   RPT VERSION V1R2 RMF       END   04/29/2002-14.00.02

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 MM.SS

04/29 10.00.00 10.01     28406     0.229    14.491    18.916     0.763     2.415     12.40      0.82      0.10
04/29 10.10.00 10.00     32696     0.264     6.071    20.415     0.880     1.012     15.52      1.00      0.09
04/29 10.20.00 10.02     34245     0.261     4.425    17.877     0.870     0.738     15.66      1.02      0.11
04/29 10.30.00 10.00     34943     0.270     5.695    19.504     0.900     0.949     15.03      1.02      0.12
04/29 10.40.00 10.00     34830     0.313    12.344    20.806     1.043     2.057     15.29      1.02      0.11
04/29 10.50.00 10.01     18360     0.373     6.619    23.352     1.243     1.103      9.34      0.54      0.06
04/29 11.00.00 10.03     29893     3.928    25.872    57.726    13.093     4.312     12.47      0.90      0.10
04/29 11.10.00 10.00     35164     0.251     8.943    15.304     0.837     1.491     15.48      1.00      0.13
04/29 11.20.00 10.00     33544     0.236     4.563    17.671     0.787     0.761     15.50      1.03      0.11
04/29 11.30.00 10.00     35057     0.257     5.829    13.389     0.857     0.972     15.73      1.01      0.11
04/29 11.40.00 10.01     35811     0.252    10.049    15.257     0.840     1.675     15.40      1.07      0.13
04/29 11.50.00 10.00     35419     0.248     4.507    17.461     0.827     0.751     15.24      1.03      0.13
04/29 12.00.00 10.01     35902     0.437     6.883    18.944     1.457     1.147     15.12      1.02      0.13
04/29 12.10.00 10.00     36967     0.247     9.635    16.407     0.823     1.606     16.10      1.06      0.12
04/29 12.20.00 10.02     36024     0.260     4.552    18.229     0.867     0.759     15.15      1.02      0.13
04/29 12.30.00 10.00     36296     0.263     5.072    20.555     0.877     0.845     15.29      1.06      0.12
04/29 12.40.00 10.00     35129     0.262    10.237    16.135     0.873     1.706     15.37      1.02      0.12
04/29 12.50.00 10.00     35355     0.274     8.098    15.106     0.913     1.350     15.31      1.06      0.12
04/29 13.00.00 10.00     36936     0.213     3.833    10.036     0.710     0.639     16.00      1.09      0.12
04/29 13.10.00 10.00     36919     0.182     5.205     9.323     0.607     0.868     16.17      1.08      0.13
04/29 13.20.00 10.00     31098     0.208     5.420     9.599     0.693     0.903     14.14      0.87      0.12
04/29 13.29.59 10.01     34909     1.184    13.976    25.991     3.947     2.329     14.20      1.00      0.13
04/29 13.40.00 10.00     34807     0.234     3.865    11.235     0.780     0.644     15.75      1.03      0.11
04/29 13.50.00 10.02     31773     0.227     8.372    12.955     0.757     1.395     14.00      0.93      0.11
```

Figure 5-59. 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 a very usable support for converting Postprocessor listings and Overview records into spreadsheets. You can use it to create and submit Postprocessor jobs directly on the workstation without an explicit 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*, and you find many report samples in the *z/OS RMF Performance Management 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 you the capability to check whether you really get the data that you expect.

PP - Overview Report

The Postprocessor creates for each report one Overview Header record and several (one for each interval) Overview Data records. If the input data for the Postprocessor consist 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 5-73. 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 5-74. 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. **Bit** **Meaning When Set** **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: ZVvvrrmm (z/OS) or VEvvrrmm (OS/390)
28	1C	EBCDIC	4	OVRSID	System identification
32	20	EBCDIC	6	OVRITIME	TOD monitor interval start: hhmmss
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	OVRECLGT	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

Table 5-74. Overview Header Record - Header Section (continued). One per record.

Offsets		Format	Length	Name	Description
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	OVRSPLID	Sysplex Id (for sysplex records)

Table 5-75. Overview Header Record - Report Column Names Section. One per interval.

Offsets		Format	Length	Name	Description
0	0	EBCDIC	8	OVRRNAME	Report column name
8	8	EBCDIC	8	OVRCOND	OVERVIEW condition name
16	10	EBCDIC	53	OVRQUAL	OVERVIEW condition qualifier
69	45	EBCDIC	1	*	Reserved
70	46	EBCDIC	2	OVROPER	Exception operator
72	48	EBCDIC	9	OVRTHV	Exception threshold value
81	51	EBCDIC	1	*	Reserved

Overview data record

Table 5-76. 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 5-77. 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	5	OVRRIINT	Reporting interval length: mm.ss
36	24	EBCDIC	4	*	Reserved

Table 5-78. Overview Data Record - Report Data Section. One per exception.

Offsets		Format	Length	Name	Description
0	0	EBCDIC	1	*	Reserved

Table 5-78. Overview Data Record - Report Data Section (continued). One per exception.

Offsets		Format	Length	Name	Description
1	1	EBCDIC	9	OVRRVAL	Report actual data

Plot report

How to request this report

The scope of data that can be reported in the Plot report depends on the data being gathered by Monitor I gatherer sessions.

You define the contents by one or more PLOTS statements (see the *z/OS RMF User's Guide* for details), as follows:

```
PLOTS(option)
```

Contents of the report

Plot reports present a graphic summary of system activity over a specified time range. Each page of plot report output consists of headings and space for two plot reports.

Generating a report

Plot reports are generated from data contained in SMF records built during a Monitor I session. The Postprocessor selects the SMF records to be used according to the type of plots that you request and the reporting period that you specify. You describe the type of plot reporting you require on control statements for the Postprocessor. For a description of Postprocessor control statements, see the *z/OS RMF User's Guide*.

Report description

Each plot report consists of a heading that identifies the plot report, the measurement scale for that report, and a line of plot data for each plot interval within the reporting period.

The **heading fields** identify the type of operating system, the release number and level of the operating system, the four-character SMF identifier associated with the system at system generation, and the RMF report level. The START field shows the date, in the form mm/dd/yy, and the time, in the form hh.mm.ss, when the first RMF measurement interval in the plot reporting period began. The END field shows the date and time when the last interval ended. The INTERVAL field shows the length of the plot interval: the value specified on the PINTV control statement or, when the PINTV statement was omitted, the average length of the RMF Monitor I measurement interval during the reporting period, in the form hh.mm.ss. The CYCLE field shows the average length of the cycle at which data was sampled during the plot reporting period.

The **measurement scale** can be either a percentage range or a range of fixed numbers. It is possible that the measurement scale provided for each plot cannot reflect the activity of your particular system accurately; perhaps the range is too small or too large. When the value recorded for a plot line exceeds the upper limit of the measurement scale, RMF prints plus signs (+) across the entire line. A field within the line of plus signs indicates the value recorded. Should your reports indicate that a measurement scale is not adequate for your particular installation, you can modify the measurement scale for any type of plot. An example of how to modify the measurement scale for a plot report appears under "Modifying the measurement scale" on page 5-199.

```
                                        R M F   P L O T   R E P O R T
          z/OS V1R4              SYSTEM ID   SYS1           START 03/03/2002 23.30.00  INTERVAL 00.27.29
                                 RPT VERSION V1R2 RMF       END   03/04/2002 10.00.00  CYCLE 1.000 SECONDS

           CPU BUSY PERCENTAGE - CPU 0                                  CPU BUSY PERCENTAGE - CPU 1
          0      20.0      40.0      60.0      80.0     100.0           0      20.0      40.0      60.0      80.0     100.0
          |---------|---------|---------|---------|---------|           |---------|---------|---------|---------|---------|
03/03 23:30|******** .         .         .         .         . 03/03 23:30|*********.         .         .         .         .
03/04 00:00|*****    .         .         .         .         . 03/04 00:00|******   .         .         .         .         .
      00:30|*****    .         .         .         .         .       00:30|******   .         .         .         .         .
      01:00|*****    .         .         .         .         .       01:00|******   .         .         .         .         .
      01:30|******   .         .         .         .         .       01:30|******** .         .         .         .         .
      02:00|**       .         .         .         .         .       02:00|***      .         .         .         .         .
      02:21|**********         .         .         .         .       02:21|**********         .         .         .         .
      02:30|*********.         .         .         .         .       02:30|******** .         .         .         .         .
      03:00|***************    .         .         .         .       03:00|****************   .         .         .         .
      03:30|**************     .         .         .         .       03:30|**************     .         .         .         .
      04:00|**       .         .         .         .         .       04:00|**       .         .         .         .         .
      04:30|**       .         .         .         .         .       04:30|**       .         .         .         .         .
      05:00|**       .         .         .         .         .       05:00|**       .         .         .         .         .
      05:30|**       .         .         .         .         .       05:30|**       .         .         .         .         .
      06:00|**       .         .         .         .         .       06:00|***      .         .         .         .         .
      06:30|**       .         .         .         .         .       06:30|***      .         .         .         .         .
      07:00|***********        .         .         .         .       07:00|************       .         .         .         .
      07:30|****************   .         .         .         .       07:30|****************** .         .         .         .
      08:00|**********************       .         .         .       08:00|**********************     .         .         .
      08:30|********************************       .         .       08:30|********************************     .         .
      09:00|********************************       .         .       09:00|********************************     .         .
      09:30|***********************************    .         .       09:30|***********************************  .         .
```

Figure 5-60. Plot Report

Figure 5-60 shows a sample of a plot report. Each line in the plot report consists of a line of asterisks that indicate the point on the measurement scale recorded for each plot interval. Each line in the plot is preceded by the time when the RMF measurement interval that collected the plotted data began; the first plot line in each calendar day included in the report is indicated by the date, in the form mm/dd. When more than one RMF interval is used for each plot line, the data from the intervals used is combined. The time shown is the time for the first measurement interval.

Types of plot reports

This section describes the types of plot reports you can request to measure various aspects of system activity.

The plot reports are described in alphabetical order.

ASCH Address Spaces (ASCH)

plots the maximum number of APPC/MVS transaction scheduler (ASCH) address spaces. The value reported corresponds to the MAX number of ASCH address spaces in the processor (CPU) activity interval report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 MAXIMUM ASCH UTILIZATION
0        60       120       180       240       300
|---------|---------|---------|---------|---------|
```

The number of ASCH address spaces can be plotted only when processor (CPU) activity was measured and ASCH was active during the reporting period.

Batch Users (BATCH)

plots the maximum number of batch users. The value reported corresponds to the MAX number of BATCH address spaces in the processor (CPU) activity interval report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 MAXIMUM BATCH UTILIZATION
0        20        40        60        80       100
|---------|---------|---------|---------|---------|
```

The maximum number of batch users can be plotted only when processor (CPU) activity was measured during the reporting period.

Channel Path Busy (CH)

plots the percentage of time when a channel path was busy. The value reported corresponds to the TOTAL UTILIZATION (%) field in the channel path activity interval report.

The measurement scale consists of percentages. The heading and measurement scale for the report are:

```
 CHANNEL PATH nn         BUSY PERCENTAGE
0      10.0      20.0      30.0      40.0      50.0
|---------|---------|---------|---------|---------|
```

CHANNEL PATH nn is the hexadecimal channel path identifier. The channel path identifier uniquely identifies the channel path and must be specified if this plot report is requested.

The total utilization percentage of the channel can be plotted only when channel path activity was measured during the reporting period.

CSA Allocated (CSAA)

plots the maximum amount of allocated CSA below the 16-megabyte line. The value reported corresponds to the MAX amount for ALLOCATED CSA in the virtual storage activity summary report for common storage.

The measurement scale consists of kilobytes (K bytes). The heading and measurement scale for the report are:

```
MAXIMUM ALLOCATED CSA (IN K BYTES)
0        500      1000      1500      2000     2500
|---------|---------|---------|---------|---------|
```

Allocated CSA can be plotted only when virtual storage activity was measured during the reporting period.

CSA Free (CSAFP)

plots the minimum number of free CSA pages below the 16-megabyte line. The value reported corresponds to the MIN amount for CSA FREE PAGES (BYTES) in the virtual storage activity summary report for common storage.

The measurement scale consists of kilobytes (K bytes). The heading and measurement scale for the report are:

```
MINIMUM CSA FREE PAGES (IN K BYTES)
0        500      1000      1500      2000      2500
|---------|---------|---------|---------|---------|
```

Free CSA can be plotted only when virtual storage activity was measured during the reporting period.

Demand Paging Rate (PAGE)

plots the demand paging rate from DASD.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 DEMAND PAGING RATE
0      40.0      80.0     120.0     160.0     200.0
|---------|---------|---------|---------|---------|
```

The demand paging rate can be plotted only when paging activity was measured during the reporting period.

Device Activity Rate (DAC)

plots the device activity rate. It is the number of successful SSCH instructions per second issued to the device (for a specific device). For a multiple exposure device, the activity rate is the sum of the activity rates for all exposures. The rate corresponds to the DEVICE ACTIVITY RATE field in the I/O device activity interval report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 DEVICE nnnn               ACTIVITY RATE
0     3.000     6.000     9.000    12.000    15.000
|---------|---------|---------|---------|---------|
```

DEVICE nnnn identifies the device included in the report; nnnn is a four-digit hexadecimal device number that uniquely identifies the device and must be specified of the plot report is requested.

An asterisk preceding the time of a plot data line indicates that records were ignored because the sampling cycle changed during the reporting period.

The device activity rate can be plotted only when I/O device activity for the appropriate device class was measured during the reporting period.

Device Busy Delay Time (DBDL)

plots the average number of milliseconds when any I/O requests for the device were delayed because the device was busy. The value corresponds to the AVG DB DELAY field in the Monitor I Device Activity report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 DEVICE nnnn   AVERAGE DEVICE BUSY DELAY TIME
0      20.0      40.0      60.0      80.0     100.0
|---------|---------|---------|---------|---------|
```

DEVICE nnnn identifies the device included in the report; nnnn is a four-digit hexadecimal number that uniquely identifies the device and must be specified if the plot report is requested.

Device Connect Time (DCON)

plots the average number of milliseconds per request the device was connected to a channel path.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 DEVICE nnnn             AVERAGE CONNECT TIME
0       20.0      40.0      60.0      80.0     100.0
|---------|---------|---------|---------|---------|
```

DEVICE nnnn identifies the device included in the report; nnnn is a four-digit hexadecimal device number that uniquely identifies the device and must be specified if the plot report is requested.

An asterisk preceding the time of a plot data line indicates that records were ignored because the sampling cycle changed during the reporting period.

The device connect time can be plotted only when I/O device activity for the device was measured during the reporting period.

Device Response Time (DRT)

plots the average number of milliseconds required to complete an I/O request.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 DEVICE nnnn          AVERAGE RESPONSE TIME
0      20.0      40.0      60.0      80.0     100.0
|---------|---------|---------|---------|---------|
```

DEVICE nnnn identifies the device included in the report; nnnn is a four-digit hexadecimal device number that uniquely identifies the device and must be specified if the plot report is requested.

An asterisk preceding the time of a plot data line indicates that records were ignored because the sampling cycle changed during the reporting period.

The device response time can be plotted only when I/O device activity for the appropriate device class was measured during the reporting period.

Expanded Storage Migration Rate (ESMR)

plots the average rate of page migration from expanded to auxiliary storage. The rate corresponds to the MIGRATION RATE for TOTAL PAGES in the Monitor I paging activity report. RMF can plot the migration rate only when paging activity was measured during the reporting period and the system has installed expanded storage.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 MIGRATION RATE FROM EXPANDED TO AUXILIARY STORAGE
0      40.0      80.0     120.0     160.0     200.0
|---------|---------|---------|---------|---------|
```

Average HUIC Count (HUIC)

Plots the average high unreferenced interval count. The value corresponds to the AVG HIGH UIC field in the Monitor I paging activity report. RMF can plot the average high UIC only when paging activity was measured during the reporting period.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 AVERAGE HIGH UIC COUNT
0       60.0     120.0     180.0     240.0     300.0
|---------|---------|---------|---------|---------|
```

I/O Activity Rate (IOAC)

Plots the I/O activity rate. It is the number of requests that were successfully selected from the logical control unit queue per second.

The activity rate is calculated by adding the path taken count for each path of the control unit per second. There is no corresponding field in the Monitor I I/O Queuing Activity report. In that report, the field CHPID TAKEN shows the separate rates for each channel path.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
   LCU nnnn          ACTIVITY RATE
0     50.000     100.000   150.000   200.000  250.000
|---------|---------|---------|---------|---------|
```

LCU nnnn identifies the logical control unit included in the report; nnnn is a one, two, three, or four-digit hexadecimal number that uniquely identifies the logical control unit and must be specified if this plot report is requested. The LCU number can be specified as a one to four-digit number.

An asterisk preceding the time of a plot data line indicates that records were ignored because the sampling cycle changed during the reporting period.

The activity rate can be plotted only when I/O queueing activity was measured during the reporting period.

OMVS Address Spaces

Plots the maximum number of OMVS address spaces. The value reported corresponds to the MAX number of OMVS address spaces in the processor (CPU) activity interval report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 MAXIMUM OMVS UTILIZATION
0        60       120       180       240       300
|---------|---------|---------|---------|---------|
```

The number of OMVS address spaces can be plotted only when processor (CPU) activity was measured and OMVS was active during the reporting period.

Page Placement on Expanded Storage (PTES)

plots the average paging transfer rate to expanded storage. The rate corresponds to the RATE OF PAGE MOVEMENT for TOTAL PAGES in the Monitor I paging activity report. RMF can plot the page transfer rate to expanded storage only when paging activity was measured during the reporting period and the system has installed expanded storage.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 PAGE PLACEMENT ON EXPANDED STORAGE
0      80.0     160.0     240.0     320.0     400.0
|---------|---------|---------|---------|---------|
```

LCU nnnn identifies the logical control unit included in the report; nnnn is a four-digit hexadecimal number that uniquely identifies the logical control unit and must be specified if this plot report is requested. The LCU number can be specified as a one to four-digit number.

Processor Busy (CPUID)

plots the percentage of time when the processor was busy. The percentage comes from the same source as the WAIT TIME PERCENTAGE field in the processor (CPU) activity interval report.

The value reported is calculated by subtracting the wait time percentage for a processor from 100 percent.

The measurement scale consists of percentages. The heading and measurement scale for the report are:

```
 CPU BUSY PERCENTAGE - CPU n
0      20.0      40.0      60.0      80.0     100.0
|---------|---------|---------|---------|---------|
```

CPU n identifies the processor included in the report, where n is the processor identifier.

The percentage of time a processor was busy can be plotted only when processor (CPU) activity was measured during the reporting period.

Service Rate (SEV)

plots the service rate. It is the number of service units provided per second. The rate corresponds to the INTERVAL SERVICE/PER SEC= field in the Workload Activity report. This data is available only in compatibility mode.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 SERVICE RATE
0      2000      4000      6000      8000     10000
|---------|---------|---------|---------|---------|
```

An asterisk preceding the time of a plot data line indicates that the IPS changed during the reporting period.

The service rate can be plotted only when workload activity was measured during the reporting period.

SQA Allocated (SQAA)

plots the maximum amount of allocated SQA below the 16-megabyte line. The value reported corresponds to the MAX amount for ALLOCATED SQA in the virtual storage activity summary report for common storage.

The measurement scale consists of kilobytes (K bytes). The heading and measurement scale for the report are:

```
MAXIMUM ALLOCATED SQA (IN K BYTES)
0        500      1000      1500      2000      2500
|---------|---------|---------|---------|---------|
```

Allocated SQA can be plotted only when virtual storage activity was measured during the reporting period.

SQA Expansion (SQAE)

plots the amount of SQA expansion into CSA below the 16-megabyte line. The value reported corresponds to the SQA EXPANSION INTO CSA field in the virtual storage activity summary report for common storage.

The measurement scale consists of kilobytes (K bytes). The heading and measurement scale for the report are:

```
AMOUNT OF SQA EXPANSION INTO CSA (IN K BYTES)
0        500      1000      1500      2000      2500
|---------|---------|---------|---------|---------|
```

SQA expansion into CSA can be plotted only when virtual storage activity was measured during the reporting period.

SQA Free (SQAFP)

plots the minimum number of free SQA pages below the 16-megabyte line. The value reported corresponds to the MIN amount for SQA FREE PAGES (BYTES) in the virtual storage activity summary report for common storage.

The measurement scale consists of kilobytes (K bytes). The heading and measurement scale for the report are:

```
MINIMUM SQA FREE PAGES (IN K BYTES)
0        500      1000      1500      2000      2500
|---------|---------|---------|---------|---------|
```

Free SQA can be plotted only when virtual storage activity was measured during the reporting period.

Started Task and Mount Task Users (STC)

plots the maximum number of started task and mount task users. The value reported corresponds to the MAX number of STC address spaces in the processor (CPU) activity interval report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 MAXIMUM STC UTILIZATION
0        10        20        30        40        50
|---------|---------|---------|---------|---------|
```

The number of started task and mount task users can be plotted only when the processor (CPU) activity was measured during the reporting period.

Swap Rate (SWA)

plots the swap sequence rate. It is the number of times per second that storage was swapped out and then swapped in. The rate 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 interval report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 SWAP RATE
0       2.000     4.000     6.000     8.000    10.000
|---------|---------|---------|---------|---------|
```

The swap rate can be plotted only when paging activity was measured during the reporting period.

Total Paging Rate (TPAG)

plots the total paging rate. To calculate the rate, the Postprocessor determines the total number of page-ins and page-outs that occurred for any reason during each plot interval; this total is divided by the number of seconds in the plot interval. The rate reported thus represents the I/O activity performed because of paging. The rate corresponds to the sum of the SUM/TOTAL RATE/PAGE IN and SUM/TOTAL RATE/PAGE OUT fields in the Monitor I paging activity report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 TOTAL PAGING RATE
0     100.0     200.0     300.0     400.0     500.0
|---------|---------|---------|---------|---------|
```

The total paging rate can be plotted only when paging activity was measured during the reporting period.

Transaction Rate (TRA)

plots the number of ended transactions per second. The rate reported comes from the same source as the ENDED TRANSACTIONS field in the Workload Activity report. This data is available only in compatibility mode.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 TRANSACTION RATE
0       4.000     8.000    12.000    16.000    20.000
|---------|---------|---------|---------|---------|
```

An asterisk preceding the time of a plot data line indicates that the IPS changed during the reporting period.

The transaction rate can be plotted only when workload activity was measured during the reporting period.

TSO Sessions (TSO)

plots the maximum number of TSO/E sessions. The value reported corresponds to the MAX number of TSO/E address spaces in the processor (CPU) activity interval report.

The measurement scale consists of fixed numbers. The heading and measurement scale for the report are:

```
 MAXIMUM TSO UTILIZATION
0        60       120       180       240       300
|---------|---------|---------|---------|---------|
```

The number of TSO/E sessions can be plotted only when the processor (CPU) activity was measured and TSO/E was active during the reporting period.

Modifying the measurement scale

Programming Interface information

Measurement scales are provided for each plot report; these scales reflect average system use. However, you might find that some or all of the measurement scales provided are not suitable for reflecting actual system use at your installation. Such a condition is indicated when plot lines consistently overflow the measurement scale or when plot lines consistently cluster on the low end of the measurement

scale. When either condition occurs, you can modify the upper limit of the measurement scale for any plot that requires a change. When you modify the upper limit, the major break points within the scale are modified automatically.

The upper limits of each measurement scale appear in the data control section ERBMFPLD within the ERBRMFPX load module. Figure 5-61 shows each plot option, the offset into control section ERBMFPLD required to access the upper limit for each plot, the contents of each entry in ERBMFPLD, and the upper limit of each measurement scale.

Each entry consists of two words. The first word contains the maximum value. The second word contains the scale factor, that is, the power of ten by which the maximum value is multiplied to arrive at the upper limit. For example, the upper limit for the swap rate plot (SWA) is 10.000; this is computed by multiplying the maximum value (10000) by the scale factor (10^{-3} or .001). Note that the scale factor must not be modified; you change the upper limit by modifying the maximum value only.

To change the scale for a particular report, perform the following steps:

1. Determine the upper limit that you require for the scale you are modifying. The values RMF prints when a plot line exceeds the upper limit can help you to determine a new upper limit that is suitable for your installation. The limit you choose must be a multiple of five.
2. Use Figure 5-61 to locate the offset to the maximum value for the scale you are modifying. Determine the maximum value that, multiplied by the scale factor, will yield the upper limit that you want. Convert that number to hexadecimal.
3. Run AMASPZAP to update the maximum value for the entry in ERBMFPLD that describes the plot whose measurement scale you are modifying.

		ERBMFPLD Entry			
		Maximum Value			
PLOTS Option	Offset (hex)	DEC	HEX	Scale Factor	Upper Unit
CPUID	0	1000	(3E8)	-1	100.0
STC	8	50	(32)	0	50
TSO	10	300	(12C)	0	300
BATCH	18	100	(64)	0	100
CH	20	500	(1F4)	-1	50.0
DCON	28	1000	(3E8)	-1	100.0
DAC	30	15000	(3A98)	-3	15.000
DRT	38	100	(64)	0	100
PAGE	40	2000	(7D0)	-1	200.0
TPAG	48	5000	(1388)	-1	500.0
SWA	50	10000	(2710)	-3	10.000
SEV	58	10000	(2710)	0	10000
TRA	60	20000	(4E20)	-3	20.000
IOAC	68	250000	(3D090)	-3	250.000
SQAA	70	2500	(9C4)	0	2500
CSAA	78	2500	(9C4)	0	2500
SQAE	80	2500	(9C4)	0	2500
SQAFP	88	2500	(9C4)	0	2500
CSAFP	90	2500	(9C4)	0	2500
DBDL	98	1000	(3E8)	-1	100.0
CUBDL	A0	1000	(3E8)	-1	100.0
PTES	A8	4000	(FA0)	-1	400.0
ESMR	B0	2000	(7D0)	-1	400.0
HUIC	B8	3000	(BB8)	-1	300.0
DPBDL	C0	1000	(3E8)	-1	100.0
APPC	C8	300	(12C)	0	300
OMVS	D0	300	(12C)	0	300

Figure 5-61. Plot Report - Measurement Scale Entries

Example

The measurement scale for the swap rate plot report runs from 0 to 10.000. Assume that you want to raise the upper limit to extend the scale from 0 to 20.000. The maximum value for SWA, located at offset X'50' from the beginning of ERBMFPLD, is 10000 (X'2710'). To achieve an upper limit of 20.000, the maximum value for SWA must be changed to 20000 (X'4E20'). The superzap job required to change the measurement scale is:

```
//ZAPLIMIT  JOB     accounting information,MSGLEVEL=1
//S1        EXEC    PGM=AMASPZAP
//SYSPRINT  DD      SYSOUT=A
//SYSLIB    DD      DSN=SYS1.SERBLINK,DISP=OLD
//SYSIN     DD      *
  NAME    ERBRMFPX    ERBMFPLD
  VER 50      00002710
  REP 50      00004E20
/*
```

End of Programming Interface information

Summary report

How to request this report

To produce this report, specify

```
SUMMARY(INT | TOT)
```

Contents of the 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 5-62 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.

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 5-4 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.

The columns for SERVICE RATE and TRANS RATE are available only if workload activity data has been gathered in compatibility mode.

```
                                    R M F   S U M M A R Y   R E P O R T
                                                                                                    PAGE  001
           z/OS V1R4              SYSTEM ID   MVS1           START 03/09/2002-08.02.26  INTERVAL 00.29.59
                                  RPT VERSION V1R2 RMF       END   03/09/2002-12.02.26  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

03/09 08.02.26 29.59  22.6    12 465.9     4     2    64    51    76    73     0     0     0     0 0.00   0.00
03/09 08.32.26 30.00  28.9    17 642.2     7     4    89    77    74    71     0     0     0     0 0.00   0.00
03/09 09.02.26 29.59  26.0    19 566.7     8     5   103    97    73    71     0     0     0     0 0.00   0.00
03/09 09.32.26 30.00  27.8    12 583.5     7     4   103   100    75    70     0     0     0     0 0.00   0.00

03/09 10.02.26 29.59  25.5    20 542.6     8     6   107   102    70    68     0     0     0     0 0.00   0.00
03/09 10.32.26 29.59  43.7    11 545.3     7     4   106    99    73    70     0     0     0     0 0.00   0.00
03/09 11.02.26 30.00  48.6    13 470.1    10     6    98    94    72    69     0     0     0     0 0.00   0.00
03/09 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 5-62. 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 5-79. Fields in the Summary Report

Field Heading	Meaning	
NUMBER OF INTERVALS	The number of RMF measurement intervals included in the reporting period.	When the report is more than one page, this field reflects the contents of the page on which it appears.
TOTAL LENGTH OF INTERVALS	The total length of the reporting period in the form hh.mm.ss.	
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 processors. 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.	

Table 5-79. Fields in the Summary Report (continued)

Field Heading	Meaning
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).
The following two columns are omitted for a system running in goal mode.	
SERVICE RATE	The number of service units (input/output service, processor service, and central storage service) absorbed during each second of the interval.
TRANS RATE	The average number of transactions that ended during each second of the interval.

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

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully. The major accessibility features in z/OS™ enable users to:

- Use assistive technologies such as screen-readers and screen magnifier software
- Operate specific or equivalent features using only the keyboard
- Customize display attributes such as color, contrast, and font size

Using assistive technologies

Assistive technology products, such as screen-readers, function with the user interfaces found in z/OS. Consult the assistive technology documentation for specific information when using it to access z/OS interfaces.

Keyboard navigation of the user interface

Users can access z/OS user interfaces using TSO/E or ISPF. Refer to *z/OS TSO/E Primer*, *z/OS TSO/E User's Guide*, and *z/OS ISPF User's Guide Volume I* for information about accessing TSO/E and ISPF interfaces. These guides describe how to use TSO/E and ISPF, including the use of keyboard shortcuts or function keys (PF keys). Each guide includes the default settings for the PF keys and explains how to modify their functions.

© Copyright IBM Corp. 1990, 2003

Glossary

This glossary contains chiefly definitions of terms used in this book, but some more general RMF and MVS terms are also defined.

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

A

APPC/MVS. Advanced program-to-program communication

ASCH address space. APPC transaction scheduler address space

AS. *Address space*

address space. That part of MVS main storage that is allocated to a job.

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

B

background session. In RMF, a monitor session that is started and controlled from the operator console. Contrast with *interactive session*

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

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

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

C

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

captured storage. See shared page group.

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

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

central processor complex (CPC). A physical collection of hardware that consists of central storage, one or more central processors, timers, and channels.

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

CICS. Customer Information Control System

compatibility mode. The implicit state of an MVS system when no workload manager service policies are in effect. Contrast with *goal mode*.

contention. Two or more incompatible requests for the same resource. For example, contention occurs if a user requests a resource and specifies exclusive use, and another user requests the same resource, but specifies shared use.

coupling facility. See *Cross-system Extended Services/Coupling Facility*.

CP. *Central processor*

criteria. Performance criteria set in the WFEX report options. You can set criteria for all report classes (PROC, SYSTEM, TSO, and so on).

CPU speed. Measurement of how much work your CPU can do in a certain amount of time.

cross-system coupling facility (XCF). A component of MVS that provides functions to support cooperation between authorized programs running within a *sysplex*.

Cross-system Extended Services/Coupling Facility (XES/CF). Provides services for MVS systems in a sysplex to share data on a coupling facility (CF).

CS. Central storage

Customer Information Control System (CICS). An IBM licensed program that enables transactions entered at remote terminals to be processed concurrently by user-written application programs. It includes facilities for building, using, and maintaining data bases.

cycle. In RMF, the time at the end of which one sample is taken. Varies between 50 ms and 9999 ms. See also *sample*.

© Copyright IBM Corp. 1990, 2003

D

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

DMN. Domain

domain. In compatibility mode, an optional method for setting bounds for the amount of service to be granted to a particular service class.

DP. Dispatching priority

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.

E

EMIF. ESCON multiple image facility

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

EPDM. Enterprise Performance Data Manager/MVS

ES. Expanded storage

ESCON multiple image facility (EMIF). A facility that allows channels to be shared among PR/SM logical partitions in an ESCON environment.

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

exception reporting. In RMF, the reporting of performance measurements that do not meet user-defined criteria. Shows potential performance problems explicitly, thus avoiding the need for constant monitoring.

expanded storage (ES). (1) An extension of processor storage. (2) Optional high-speed storage that transfers 4KB pages to and from central storage.

G

generalized trace facility (GTF). A service program that records significant system events, such as supervisor calls and start I/O operations, for the purpose of problem determination.

GO mode. In RMF, the Monitor III mode in which the screen is updated with the interval you specified in your session options. The terminal cannot be used for anything else when it is in GO mode. See also *mode*.

goal mode. The implicit mode of an MVS system that has active service policies and performance goals defined by the workload manager. Contrast with *compatibility mode*.

graphic mode. In RMF Monitor III, the mode which presents the performance data from the system in graphic format using the GDDM product. Contrast with *tabular mode*.

GTF. generalized trace facility

H

high-speed buffer (HSB). A cache or a set of logically partitioned blocks that provides significantly faster access to instructions and data than provided by central storage.

HS. hiperspace

HSB. High-speed buffer

HSM. Hierarchical Storage Manager

I

IMS. Information Management System

Information Management System (IMS). A database/data communication (DB/DC) system that can manage complex databases and networks. Synonymous with IMS/VS.

installation performance specification (IPS). In MVS, a set of installation-supplied control information used by the system workload manager. An IPS includes performance group definitions, performance objectives, and coefficients used to establish the service rate. See also service rate.

interactive session. In RMF, a monitor display-session that is controlled from the display terminal. Contrast with *background session*.

J

JES. Job Entry Subsystem

L

LCU. Logical control unit

License Manager. The IBM License Manager is the base for a new software pricing model. It allows vendors to enable their products for licensed software management by customers and is the basic tool IBM will use to implement the Workload License Charges pricing model on z900 servers.

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.

M

migration rate. The rate (pages/second) of pages being moved from expanded storage through central storage to auxiliary storage.

mintime. The smallest unit of sampling in Monitor III. Specifies a time interval during which the system is sampled. The data gatherer combines all samples gathered into a set of samples. The set of samples can be summarized and reported by the reporter.

mode. Monitor III can run in various modes: GO mode (see *GO mode*) and STOP mode, which is the default mode. See also *graphic mode* and *tabular mode*.

MPL. Multiprogramming level

O

OMVS. Reference to z/OS UNIX System Services

P

partitioned data set (PDS). A data set in direct access storage that is divided into partitions, called members, each of which can contain a program, part of a program, or data.

PDS. partitioned data set

performance management. (1) The activity which monitors and allocates data processing resources to applications according to goals defined in a service level agreement or other objectives. (2) The discipline that encompasses collection of performance data and tuning of resources.

performance group. Group of work with the same performance objectives managed by the SRM.

PG. Performance group

PGN. Performance group number

PR/SM. Processor Resource/Systems Manager

Processor Resource/Systems Manager (PR/SM). The feature that allows the processor to run several operating systems environments simultaneously and provides logical partitioning capability. See also *LPAR*.

R

range. The time interval you choose for your report.

Resident time. The time the address space was swapped in, in units of seconds.

S

sample. Once in every cycle, the number of jobs waiting for a resource, and what job is using the resource at that moment, are gathered for all resources of a system by Monitor III. These numbers constitute one sample.

SCP. System control program

seek. The DASD arm movement to a cylinder. A seek can range from the minimum to the maximum seek time of a device. In addition, some I/O operations involve multiple imbedded seeks where the total seek time can be more than the maximum device seek time.

service class. In Workload Manager, a subdivision of a *workload*. Performance goals and capacity boundaries are assigned to service classes.

service level agreement (SLA). A written agreement of the information systems (I/S) service to be provided to the users of a computing installation.

Service Level Reporter (SLR). An IBM licensed program that provides the user with a coordinated set of tools and techniques and consistent information to help manage the data processing installation. For example, SLR extracts information from SMF, IMS, and CICS logs, formats selected information into tabular or graphic reports, and gives assistance in maintaining database tables.

service rate. In the system resources manager, a measure of the rate at which system resources (services) are provided to individual jobs. It is used by the installation to specify performance objectives, and used by the workload manager to track the progress of individual jobs. Service is a linear combination of processing unit, I/O, and main storage measures that can be adjusted by the installation.

shared page groups. An address space can decide to share its storage with other address spaces using a function of RSM. As soon as other address spaces use these storage areas, they can no longer be tied to only one address space. These storage areas then reside as *shared page groups* in the system. The pages of shared page groups can reside in central, expanded, or auxiliary storage.

SLA. service level agreement

SLIP. serviceability level indication processing

SLR. Service Level Reporter

SMF. System management facility

SMF buffer. A wrap-around buffer area in storage, to which RMF data gatherers write performance data, and from which the Postprocessor extracts data for reports.

speed. See *workflow*

SRB. Service request block

SRM. System resource manager

SSCH. Start subchannel

system control program (SCP). Programming that is fundamental to the operation of the system. SCPs include MVS, VM, and VSE operating systems and any other programming that is used to operate and maintain the system. Synonymous with *operating system*.

sysplex. A complex consisting of a number of coupled MVS systems.

T

tabular mode. In RMF, the mode in which Monitor III displays performance data in the form of lists. Contrast with *graphic mode*.

TCB. Task control block

threshold. The exception criteria defined on the report options screen.

throughput. A measure of the amount of work performed by a computer system over a period of time, for example, number of jobs per day.

TPNS. Teleprocessing network simulator

TSO. Time Sharing Option, see *Time Sharing Option/Extensions*

Time Sharing Option Extensions (TSO/E). In MVS, a time-sharing system accessed from a terminal that allows user access to MVS system services and interactive facilities.

U

UIC. Unreferenced interval count

uncaptured time. CPU time not allocated to a specific address space.

using. Jobs getting service from hardware resources (PROC or DEV) are *using* these resources.

V

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

VTOC. Volume table of contents

W

workflow. (1) The workflow of an address space represents how a job uses system resources and the speed at which the job moves through the system in relation to the maximum average speed at which the job could move through the system. (2) The workflow of resources indicates how efficiently users are being served.

workload. A logical group of work to be tracked, managed, and reported as a unit. Also, a logical group of service classes.

WLM. Workload Manager

WSM. Working Set Manager

X

XCF. Cross-system coupling facility

XES/CF. See *Cross-system Extended Services/Coupling Facility*.

Notices

This information was developed for products and services offered in the U.S.A.

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.

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

IBM Corporation
Mail Station P300
2455 South Road
Poughkeepsie New York 12601-5400
U.S.A.

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

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

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

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, 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"

© Copyright IBM Corp. 1990, 2003

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.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Programming Interface Information

This book documents intended Programming Interfaces that allow the customer to write programs to obtain the services of RMF.

This book is intended to help the customer to use RMF reports, and contains a detailed description of all reports.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States and/or other countries:

- CICS
- DB2
- ESCON
- GDDM
- Hiperspace
- IBM
- IMS
- OS/390
- PR/SM
- Processor Resource/Systems Manager
- RACF
- Resource Measurement Facility
- RMF
- VTAM
- z/OS
- zSeries

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

Domino and Lotus are trademarks of Lotus Development Corporation.

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

Index

Special characters

% ANY ALLOC (allocated %) field
  in Monitor I DEVICE report 5-67
% CU BUSY (control unit busy %) field
  in Monitor I IOQ report 5-98
  in Monitor II IOQUEUE report 3-43
  in Monitor III IOQUEUE report 2-120
% Delayed for (percent delayed for) field
  in Monitor III DELAY report 2-65
% DEV CONN (device connect %) field
  in Monitor I DEVICE report 5-66
% DEV RESV (device reserved %) field
  in Monitor I DEVICE report 5-67
% DEV UTIL (device utilization %) field
  in Monitor I DEVICE report 5-67
% DP BUSY (director port busy %) field
  in Monitor I IOQ report 5-98
  in Monitor II IOQUEUE report 3-43
  in Monitor III IOQUEUE report 2-119
% Frames (percentage frames) field
  in Monitor III STORR report 2-169
% I/O interrupts handled VIA TPI (I/O interrupts) field
  in Monitor I CPU report 5-45
% of TRX (percentage of transactions) field
  in Monitor III SYSRTD report 2-188
% slots in use field
  in Monitor II PGSP report 3-49
%D RV (reserved device %) field
  in Monitor II DEVV report 3-31

A

abnormally ended transactions field
  in Monitor III SYSWKM report 2-208
accessibility A-1
account number (ACCT INFO) field
  in Monitor II TRX report 3-66
action field
  in Monitor III STORC report (Act) 2-160
  in Monitor III WFEX report options action panel 2-222
activated at field
  in Monitor III SYSSUM report 2-194
active frames field
  in Monitor III STORF report 2-166
  in Monitor III STORJ report 2-132
active percentage field
  in Monitor III DEVJ report (Active) 2-126
  in Monitor III DEVR report (ACT %) 2-78
  in Monitor III STORR report (ACT %) 2-170
active policy field
  in Monitor III SYSSUM report 2-194
activity rate field
  in Monitor II DEVV report 3-29
  in Monitor III DEVR report (Act Rate) 2-78
  in the Monitor III DEVN report (Rate) 2-75
actual field
  in Monitor III SYSWKM report 2-208
actual value field
  in Postprocessor EXRPTS report 5-182
address space activity
  reports 3-8
address space analysis
  in Monitor I CPU report 5-41
address space DELAY (%)
  definition of 2-9
address space identifier (ASID) field
  in Monitor III STORC report 2-160
address space resource data
  Monitor II report (ARD)
    field descriptions 3-10, 3-11
    option 3-8
    sample 3-8
address space resource data by jobname
  Monitor II report (ARDJ)
    field descriptions 3-10, 3-11
    option 3-8
address space SRM data
  Monitor II report (ASRM)
    field descriptions 3-18, 3-19
    option 3-17
    sample 3-17
address space SRM data by jobname
  Monitor II report (ASRMJ)
    field descriptions 3-18, 3-19
    option 3-17
    sample 3-17
address space state data
  Monitor II report (ASD)
    field descriptions 3-14, 3-16
    option 3-13
    sample 3-13
address space state data by jobname
  Monitor II report (ASDJ)
    field descriptions 3-14, 3-16
    option 3-13
address space Using (%) 2-8
address space Workflow (%) 2-8
  definition of 2-7
address spaces field
  in Monitor I PAGING report 5-109
address spaces serving this service class field
  in Monitor III SYSWKM report 2-211
AFC (available frames) field
  in Monitor II SPAG report 3-61
AFC (average frame queue) field
  in Monitor II SRCS report 3-63
alert field
  in Monitor III WFEX report definition and criteria panel 2-225
allocated (% ANY ALLOC) field
  in Monitor I DEVICE report 5-67
allocated area size field
  in Monitor I VSTOR (common storage summary) report 5-143
allocated CSA by subpool field
  in Monitor I VSTOR (common storage detail) report 5-144
allocated CSA/SQA field
  in Monitor I VSTOR common storage summary report 5-143
allocated SQA by subpool field
  in Monitor I VSTOR (common storage detail) report 5-144
allocated storage
  definition 5-140
AMASPZAP program
  changing sample to cycles ratio 5-141
amount of common storage not released field
  in Monitor III STORCR report 2-164
amount used field
  in Monitor III STORC report 2-161
AP 5-169
APPC/MVS service
  monitoring with RMF 5-122
ARD (address space resource data)
  Monitor II report
    field descriptions 3-10, 3-11
    option 3-8
    sample 3-8
ARDJ (address space resource data by jobname)
  Monitor II report
    field descriptions 3-10, 3-11
    option 3-8
ASCH (average) field
  in Postprocessor SUMMARY report 5-204
ASCH (maximum) field
  in Postprocessor SUMMARY report 5-204
ASCH address spaces
  monitoring with RMF 5-45
ASD (address space state data)
  Monitor II report
    field descriptions 3-14, 3-16
    option 3-13
    sample 3-13
ASDJ (address space state data by jobname)
  Monitor II report
    field descriptions 3-14, 3-16
    option 3-13
ASID (address space identifier) field
  in Monitor II SENQ report 3-56
  in Monitor II SENQR report 3-59
  in Monitor III STORC report 2-160

© Copyright IBM Corp. 1990, 2003

ASID (address space identifier) field *(continued)*
in Monitor III SYSWKM report 2-211
ASRM (address space SRM data)
Monitor II report
field descriptions 3-18, 3-19
option 3-17
sample 3-17
ASRMJ (address space SRM data by jobname)
Monitor II report
field descriptions 3-18, 3-19
option 3-17
sample 3-17
automatic customization 2-231
where to specify 2-231
workflow/exception (WFEX) options 2-214
auxiliary slots field
in Monitor III STORF report 2-166
in Monitor III STORJ report 2-132
auxiliary storage (AUX STOR) direct field
in Monitor I PAGING report 5-123
auxiliary storage (AUX STOR) total field
in Monitor I PAGING report 5-123
auxiliary storage (AUX STOR) via transition field
in Monitor I PAGING report 5-123
auxiliary storage/swap in field
in Monitor I PAGING report 5-123
auxiliary storage/swap out field
in Monitor I PAGING report 5-123
available at end of range field
in Monitor III STORC report 2-160
available frames (AFC) field
in Monitor II SPAG report 3-61
available slots field
in Monitor I PAGING report 5-118
AVE HI UIC (average high unreferenced interval count) field
in Monitor III STORR report 2-169
average absorption field
in Monitor I WKLD report 5-163
average active users field
in Monitor III STORR report 2-172
average connect (AVG CONN) time field
in Monitor I DEVICE report 5-66
average CPU utilization field
in Monitor III SYSINFO report (average CPU UTIL) 2-180
in Monitor III WFEX report (Average CPU Util) 2-217
average CPU Utilization field
in the Monitor III SYSTREND report (CPU %) 2-202
average CSA to SQA conversion field
in Monitor III STORC report 2-159
average device busy (AVG DB) delay field
in Monitor I DEVICE report 5-66
average disconnect (AVG DISC) time field
in Monitor I DEVICE report 5-66
average DPB delay (AVG DPB DLY) field
in Monitor I DEVICE report 5-66, 5-129
average execution time (Avg. Exec. Time) field
in Monitor III SYSWKM report 2-208
average field
in Monitor I TRACE report 5-136
average frame queue (AFC) field
in Monitor II SRCS report 3-63
average frames field
in Monitor III STORJ report 2-132
in Monitor III STORS report 2-174
average high unreferenced interval count (AVG HI UIC) field
in Monitor III STORR report 2-169
average IOSQ time field
in Monitor I DEVICE report 5-65
average mount time field
in Monitor I DEVICE report 5-68
average number allocation (AVG NUMBER ALLOC) field
in Monitor I DEVICE report 5-67
average number delayed for field
in Monitor III STORS report 2-174
in Monitor III SYSINFO report 2-182
average pages field
in Monitor I PAGING report 5-111
in Monitor II PGSP report (AVG PAGES PER I/O) 3-49
average pend time field
in Monitor I DEVICE report 5-66
average queue length field
in Monitor I ENQ report (AVG Q LNGTH) 5-78
in Postprocessor XCF report 5-178
average ready users (RUA) field
in Monitor II DDMN report 3-25
average response (AVG RESP) time field
in Monitor I DEVICE report 5-65
average response time (Avg. Resp. Time) field
in Monitor III SYSRTD report 2-188
in Monitor III SYSSUM report (active time) 2-198
in Monitor III SYSSUM report (queue time) 2-198
in Monitor III SYSSUM report (total time) 2-199
in Monitor III SYSWKM report 2-208
average service rate/transaction (AVG TRX SERV RATE) field
in Monitor I WKLD report 5-163
average SRB time field
in the Monitor III SYSTREND report 2-202
average standard deviation (STD DEVIATION) field
in Monitor I WKLD report 5-161
average TCB time field
in the Monitor III SYSTREND report 2-202
average transaction time (AVG TRANS TIME) field
in Monitor II TRX report 3-66
average use summary field
in Monitor III STORC report 2-160
average users delayed field
in Monitor III DEVJ report 2-127
Average Users line
in Monitor III GROUP report 2-110
average using (AVG USG) field
in Monitor III SYSINFO report 2-181

B

bad slots field
in Monitor I PAGESP report 5-105
in Monitor I PAGING report 5-119
bad swap sets field
in Monitor I PAGESP report 5-105
Basic mode
CPU utilization 5-42
BASIC mode
channel path activity report 5-36
batch jobs (average) field
in Postprocessor SUMMARY report 5-203
batch jobs (maximum) field
in Postprocessor SUMMARY report 5-203
blocks per second field
in Monitor I PAGING report 5-111
bottom of allocated area field
in Monitor I VSTOR (private area summary) report 5-146
buffer (bigger) % field
in Postprocessor XCF report 5-175
buffer (fit) % field
in Postprocessor XCF report 5-175
buffer (over) % field
in Postprocessor XCF report 5-175
buffer (smaller) % field
in Postprocessor XCF report 5-175
buffer length field
in Postprocessor XCF report 5-175
buffers unavailable field
in Postprocessor XCF report 5-178
busy field
in Postprocessor XCF report 5-178
busy time percentage field
in Monitor I CPU report 5-44

C

C (class) field
in Monitor III DELAY report 2-63
in Monitor III DEV report 2-72
in Monitor III DEVR report 2-80
in Monitor III PROC report 2-144
in Monitor III STOR report 2-156
in Monitor III STORC report 2-160
in Monitor III STORF report 2-165
in Monitor III WFEX report definition and criteria panel 2-224
in Monitor III WFEX report options action panel 2-222
in Monitor III XCF report 2-233
CACHDET (cache activity details)
Monitor III report
command 2-33
field descriptions 2-35
sample 2-33

CACHDET (cache details)
Monitor III reporter session command summary 2-17
CACHE (cache subsystem activity)
Postprocessor report
Subsystem Device Overview report 5-14
Subsystem Device report 5-17
Subsystem Overview report 5-9
Subsystem Status report 5-9
Subsystem Summary report 5-7
summary 5-6
cache activity details
Monitor III report (CACHDET)
command 2-33
field descriptions 2-35
sample 2-33
cache activity summary
Monitor III report (CACHSUM)
command 2-38
field descriptions 2-39
sample 2-38
CACHE structure in coupling facility 5-23
CACHSUM (cache activity summary)
Monitor III report
command 2-38
field descriptions 2-39
sample 2-38
CACHSUM (cache summary)
Monitor III reporter session command summary 2-17
CAP 5-169
capacity field
in Monitor III RG report (actual) 2-147
in Monitor III RG report (max) 2-147
in Monitor III RG report (min) 2-147
in Monitor III SYSSUM report (actual) 2-199
in Monitor III SYSSUM report (max) 2-199
in Monitor III SYSSUM report (min) 2-199
Capacity report
contents 2-60
field descriptions 2-60
capping delay (Capp) field
in Monitor III SYSWKM report 2-211
capping field
in Monitor I CPU report 5-49
category field
in Monitor I PAGING report 5-109
CCF 5-55
central processing complex (CPC) 2-57, 3-22, 5-35, 5-36
in Monitor I CHAN report 5-37
central processor unit (CPU) activity
Monitor I report
field descriptions 5-43, 5-44, 5-45, 5-50
field descriptions (partition data report section) 5-47
sample (partition data report section) 5-47
central storage frames (CS F) field
in Monitor II ASD report 3-14
central storage frames (CS F) field *(continued)*
in Monitor II ASDJ report 3-14
central storage target (CS TAR) field
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
central storage/processor/SRM activity
Monitor II report (SRCS)
field descriptions 3-63
option 3-62
sample 3-62
CFACT (coupling facility activity)
Monitor III report
command 2-42, 2-48
field descriptions 2-45
report options panel 2-48
sample 2-43
summary 2-42
Monitor III reporter session command summary 2-17
CFOVER (coupling facility overview)
Monitor III report
command 2-49
field descriptions 2-49
sample 2-49
summary 2-49
Monitor III reporter session command summary 2-17
CFSYS (coupling facility systems)
Monitor III report
command 2-51
field descriptions 2-52
sample 2-51
summary 2-51
Monitor III reporter session command summary 2-17
CHAN (channel path activity)
Monitor I report
field descriptions 5-37, 5-38, 5-39, 5-40
sample 5-37
CHANNEL (channel path activity)
Monitor II report
field descriptions 3-21, 3-22
option 3-20
sample 3-20
Monitor III report
field descriptions 2-56, 2-58
option 2-55
Monitor III reporter session command summary 2-17
Channel Path
generation 2-56
channel path activity
Monitor I report
field descriptions 5-37, 5-38, 5-39, 5-40
sample 5-37
Monitor II report
field descriptions 3-21, 3-22
option 3-20
sample 3-20
Monitor III report
field descriptions 2-56, 2-58
option 2-55
channel path ID field
in Monitor I CHAN report 5-38
channel path ID field *(continued)*
in Monitor II CHANNEL report 3-21
in Monitor III CHANNEL report 2-56
channel path measurement facility (CPMF) 5-35
in Monitor I CHAN report 5-38
channel path measurement table 5-35
channel path type field
in Monitor I CHAN report 5-38
in Monitor II CHANNEL report 3-21
in Monitor III CHANNEL report 2-56
channel path utilization field
in Monitor I CHAN report 5-39
in Monitor II CHANNEL report 3-22
in Monitor III CHANNEL report 2-57
channel paths field
in Monitor I IOQ report 5-98
in Monitor III IOQUEUE report 2-118
choosing graphic and tabular reports 2-6
CHPID taken field
in Monitor I IOQ report 5-98
in Monitor II IOQUEUE report 3-43
in Monitor III IOQUEUE report 2-119
CIDX (contention index) field
in Monitor II DDMN report 3-25
CL (current location) field
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
class (C) field
in Monitor III DELAY report 2-63
in Monitor III DEV report 2-72
in Monitor III DEVR report 2-80
in Monitor III PROC report 2-144
in Monitor III STOR report 2-156
in Monitor III STORC report 2-160
in Monitor III STORF report 2-165
in Monitor III WFEX report definition and criteria panel 2-224
in Monitor III WFEX report options action panel 2-222
in Monitor III XCF report 2-233
class field
in Monitor III GROUP report 2-107
in Monitor III SYSWKM report 2-208
CMPL field
in Monitor II DDMN report 3-25
command and selection lines, tailoring 2-15
commands
Monitor III Reporter session report and report options command summary 2-19
commands you can use in monitor III 2-12
common monitor III report measurements 2-6
common storage
detail report 5-144
Monitor III report (STORC)
command 2-158

common storage *(continued)*
Monitor III report (STORC) *(continued)*
field descriptions - job information 2-160, 2-161
field descriptions - system information 2-159, 2-160
options panel 2-162
sample 2-159
summary 2-158
common storage remaining
Monitor III report (STORCR)
command 2-163
field descriptions 2-164
sample 2-163
summary 2-163
Concepts, Monitor III
mintime 2-10
range 2-10
report interval 2-10
shortened intervals, reason for 2-12
connect time field
in Monitor II DEVV report 3-30
connect time percentage field
in Monitor III DEV report (CON %) 2-72
in Monitor III DEVJ report 2-126
in Monitor III DEVR report (CON %) 2-78
contention
detected by the system 5-134
contention index (CIDX) field
in Monitor II DDMN report 3-25
contention rate field
in Monitor I IOQ report 5-99
in Monitor II IOQUEUE report 3-43
in Monitor III IOQUEUE report 2-119
contention time field
in Monitor I ENQ report 5-77
control statement name field
in Postprocessor EXRPTS report 5-182
control unit busy field
in Monitor I IOQ report (% CU BUSY) 5-98
in Monitor II IOQUEUE report (% CU BUSY) 3-43
in Monitor III IOQUEUE report (% CU BUSY) 2-120
control unit field
in Monitor I IOQ report 5-97
in Monitor III IOQUEUE report 2-118
in Monitor III STORR report 2-170
Coupling facility
Monitor III report
CFACT (coupling facility activity) 2-42
CFOVER (coupling facility overview) 2-49
CFSYS (coupling facility systems) 2-51
Postprocessor report 5-22
CACHE structure 5-23
LIST structure 5-23
LOCK structure 5-23
sections 5-22
Coupling facility *(continued)*
Postprocessor report *(continued)*
structure of 5-23
subchannel activity section, field description 5-31
usage summary section 5-23
usage summary section, field description 5-24
XES structure activity 5-22
XES structure activity section, field description 5-28
CPC (CPC capacity)
Monitor III report 2-59
contents 2-60
field descriptions 2-60
Monitor III reporter session command summary 2-17
CPMF (channel path measurement facility) field
in Monitor I CHAN report 5-38
CPU (central processor unit) activity
Monitor I report
field descriptions 5-43, 5-44, 5-45, 5-47, 5-50
sample (partition data report section) 5-47
CPU busy field
in Postprocessor SUMMARY report 5-203
CPU model field
in Monitor I CPU report 5-43
CPU number field
in Monitor I CPU report 5-43
CPU protection xxi
CPU serial number field
in Monitor I CPU report 5-44
CPU time field
in Monitor II ARD report 3-11
in Monitor II ARDJ report 3-11
CPU utilization
LPAR view 5-42
MVS view 5-42
CPU utilization field
in Monitor II SRCS report 3-63
in Monitor III SYSINFO report (average CPU UTIL) 2-180
CPU Utilization field
in Monitor III WFEX report (Average CPU Util) 2-217
in the Monitor III SYSTREND report (CPU %) 2-202
criteria names selection panel 2-227
criteria set
in Monitor III WFEX report definition and criteria panel 2-226
critical value field
in Monitor III WFEX report 2-220
cross memory (XM) field
in Monitor II ARD report 3-11
in Monitor II ARDJ report 3-11
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
cross system coupling facility job delays (XCFJ)
Monitor III report
command 2-133
field descriptions 2-133
cross system coupling facility job delays (XCFJ) *(continued)*
Monitor III report *(continued)*
summary 2-133
cross-system coupling facility (XCF)
Monitor III report
command 2-232
field descriptions 2-233
report options panel 2-233
sample 2-232
summary 2-232
Monitor III reporter session
summary 2-19
Postprocessor report
field descriptions 5-175, 5-176, 5-177, 5-178
sample 5-174
sample (path statistics) 5-177
sample (usage by member) 5-176
CRYPTO (CRYPTO hardware activity)
Postprocessor report
summary 5-54
Cryptographic Coprocessor Facility 5-55
cryptographic processors
adjunct processor 5-169
asynchronous message processor 5-169
CS F (central storage frames) field
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
CS TAR (central storage target) field
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
CSA frames field
in SRCS report 3-63
CSA in field
in Monitor II SPAG report 3-61
CSA rate field
in Monitor II ARD report 3-11
in Monitor II ARDJ report 3-11
CU (control unit) field
in the Monitor III DEVT report 2-83
current location (CL) field
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
cursor-sensitive control, using 2-5
cycles/sample field
in Monitor I TRACE report 5-136

D

DASD rate field
in Postprocessor SUMMARY report 5-203
DASD response time field
in Postprocessor SUMMARY report 5-203
data
gathering 1-2
Data Encryption Standard 5-55
data field
in Monitor III SYSRTD report 2-188
data fields
address space state data 3-14
in Postprocessor EXRPTS report 5-181

data index (DI)
Monitor III report
command 2-26
field descriptions 2-29
report options panel 2-31
sample 2-26
summary 2-26
data set delays
DSND (Monitor III report)
field descriptions 2-86
sample 2-86
DSNJ (Monitor III report) 2-89
DSNV (Monitor III report) 2-91
data set name field
in Monitor I PAGESP report 5-106
in Monitor II PGSP report 3-50
data-in-virtual (DIV) frames field
in Monitor III STORJ report 2-133
Data, Gathering and Reporting, Introducing 1-1
date field
in Postprocessor EXRPTS report 5-182
in Postprocessor SUMMARY report 5-203
DB (device busy) delay field
in Monitor II DEVV report 3-30
DCM
description xx
Monitor II
CHANNEL report 3-20
IOQUEUE report 3-40
Monitor III
CHANNEL report 2-55
IOQUEUE report 2-117
Postprocessor
CHANNEL report 5-35
IOQ report 5-92
DDMN (domain activity)
Monitor II report
field descriptions 3-25
option 3-24
sample 3-24
DDNAME
in Monitor III DI report 2-29
DEL Q LNGTH (delayed queue length) field
in Monitor II IOQUEUE report 3-43
in Monitor III IOQUEUE report 2-119
DELAY
Monitor III report
command 2-62
command summary 2-17
field descriptions 2-63, 2-67
job selection panel 2-69
options panel 2-68
sample 2-56, 2-62
summary 2-62
Delay (%) of an address space
definition of 2-9
Delay DB% (device busy delay percentage) field
in Monitor III DEVJ report 2-127
delay information
job groups measured 2-2
delay monitoring 1-5
delay percentage field
in Monitor III DELAY report (DLY %) 2-63
in Monitor III DEV report (DLY %) 2-72
in Monitor III DEVR report (DLY %) 2-80
in Monitor III ENQ report (DLY %) 2-100
in Monitor III ENQR report (Delayed) 2-104
in Monitor III HSM report (DLY %) 2-116
in Monitor III HSMJ report 2-129
in Monitor III JES report (DLY %) 2-122
in Monitor III JESJ report 2-129
in Monitor III PROC report (DLY %) 2-144
in Monitor III STOR report (DLY %) 2-156
in Monitor III XCF report (DLY %) 2-233
in Monitor III XCFJ report 2-133
DELAY Q LNGTH (delayed queue length) field
in Monitor I IOQ report 5-99
delay samples
for workflow of a resource 2-7
delayed name field
in Monitor III ENQR report 2-104
delayed queue length field
in Monitor I IOQ report (DELAY Q LNGTH) 5-99
in Monitor II IOQUEUE report (DEL Q LNGTH) 3-43
delayed status field
in Monitor III ENQR report 2-104
DELAYJ (delays for specified job) command
Monitor III reporter session
summary 2-17
demand paging field
in Postprocessor SUMMARY report 5-204
DES 5-55
description field
in Monitor III GROUP report 2-107
detailed storage delays
Monitor III Utility report 2-84
DEV (device delays)
Monitor III report
command 2-71
field descriptions 2-72
options panel 2-73
sample 2-71
summary 2-71
Monitor III reporter session command
summary 2-17
DEV (device type) field
in Monitor II SENQR report 3-59
DEV NUM (device number) field
in Monitor I DEVICE report 5-64
in Monitor I PAGESP report 5-105
in Monitor II DEVV report 3-29
in Monitor II PGSP report 3-49
DEV TYPE (device type) field
in Monitor II PGSP report 3-49
DEV UTIL (device utilization) % field
in Monitor II DEVV report 3-31
DEV/CU (device/control unit) field
in Monitor III DEVR report 2-79
DEVICE (device activity)
Monitor I report
field descriptions 5-64, 5-65, 5-66, 5-67, 5-68
option 5-60
sample 5-63
sample (communication equipment) 5-63
sample (graphics) 5-63
sample (magnetic tape) 5-63
sample (unit record) 5-63
device activity
Monitor I report
field descriptions 5-64, 5-65, 5-66, 5-67, 5-68
option 5-60
sample 5-63
sample (communication equipment) 5-63
sample (graphics) 5-63
sample (magnetic tape) 5-63
sample (unit record) 5-63
device activity (DEVN)
Monitor III report
command 2-74
field descriptions 2-75
sample 2-75
summary 2-74
device activity rate field
in Monitor I DEVICE report 5-65
device activity trend (DEVT)
Monitor III report
command 2-82
field descriptions 2-83
sample 2-83
summary 2-82
device busy delay field
in Monitor II DEVV report (DB) 3-30
in Monitor III DEVJ report (Delay DB%) 2-127
device connect % (% DEV CONN) field
in Monitor I DEVICE report 5-66
device connect time field
in Monitor II ARD report 3-10
in Monitor II ARDJ report 3-10
device delays
Monitor III report (DEV)
command 2-71
field descriptions 2-72
options panel 2-73
sample 2-71
summary 2-71
device delays job (DEVJ)
Monitor III report
command 2-125
field descriptions 2-126, 2-127
summary 2-125
device field
in Monitor III DEVJ report 2-126

device number field
in Monitor I DEVICE report (DEV NUM) 5-64
in Monitor I PAGESP report (DEV NUM) 5-105
in Monitor II DEVV report (DEV NUM) 3-29
in Monitor II PGSP report (DEV NUM) 3-49
in Monitor III DEVR report 2-78
in Postprocessor XCF report 5-178
in Postprocessor XCF report (FROM/TO DEVICE OR STRUCTURE) 5-178
in the Monitor III DEVT report 2-83
device reserved % (% DEV RESV) field
in Monitor I DEVICE report 5-67
device resource delays (DEVR)
Monitor III report
command 2-77
field descriptions 2-78, 2-80
report options panel 2-81
sample 2-78
summary 2-77
device type field
in Monitor I DEVICE report 5-64
in Monitor I PAGESP report 5-105
in Monitor II PGSP report (DEV TYPE) 3-49
in Monitor II SENQR report (DEV) 3-59
in Monitor III STORR report 2-170
device using percentage field
in Monitor III PROCJ report 2-131
device utilization % field
in Monitor I DEVICE report (% DEV UTIL) 5-67
in Monitor II DEVV report (% DEV UTIL) 3-31
device/control unit (DEV/CU) field
in Monitor III DEVR report 2-79
DEVJ (device delays job)
JOB report 2-125
Monitor III report
command 2-125
field descriptions 2-126, 2-127
summary 2-125
Monitor III reporter session command
summary 2-18
DEVN (device activity)
Monitor III report
command 2-74
field descriptions 2-75
sample 2-75
summary 2-74
DEVR (device resource delays)
Monitor III report
command 2-77
field descriptions 2-78, 2-80
report options panel 2-81
sample 2-78
summary 2-77
DEVR (resource-oriented device delays)
Monitor III reporter session command
summary 2-18
DEVT (device activity trend)
Monitor III report
command 2-82
field descriptions 2-83
sample 2-83
summary 2-82
DI (data index)
Monitor III report
command 2-26
field descriptions 2-29
report options panel 2-31
sample 2-26
summary 2-26
DI (data index) command
Monitor III reporter session
summary 2-18
director port busy field
in Monitor I IOQ report (% DP BUSY) 5-98
in Monitor II IOQUEUE report (% DP BUSY) 3-43
in Monitor III IOQUEUE report (% DP BUSY) 2-119
disability A-1
disconnect time field
in Monitor II DEVV report 3-30
in Monitor III DEVJ report 2-126
in Monitor III DEVR report (DSC %) 2-79
in Monitor III STORR report (DSC %) 2-171
dispatch interval field
in Monitor I CPU report 5-48
dispatch priority (DP PR) field
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
display report
Monitor II
reports 3-66
Monitor III 2-26, 2-233
reports 3-1
distribution of queue lengths (%) field
in Monitor I CPU report 5-45
DIV (data-in-virtual) frames field
in Monitor III STORJ report 2-133
DLY % (delay percentage) field
in Monitor III DELAY report 2-63
in Monitor III DEV report 2-72
in Monitor III DEVR report 2-80
in Monitor III ENQ report 2-100
in Monitor III HSM report 2-116
in Monitor III JES report 2-122
in Monitor III PROC report 2-144
in Monitor III STOR report 2-156
in Monitor III XCF report 2-233
DMN (domain) field
in Monitor I WKLD report 5-159
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
in Monitor II DDMN report 3-25
in Monitor III DEVR report 2-80
in Monitor III STOR report 2-156
in Monitor III STORC report 2-160
in Monitor III STORS report 2-174
in Monitor III XCF report 2-233
domain (DMN) field
in Monitor I WKLD report 5-159
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
in Monitor II DDMN report 3-25
in Monitor III DEVR report 2-80
in Monitor III STOR report 2-156
in Monitor III STORC report 2-160
in Monitor III STORS report 2-174
in Monitor III XCF report 2-233
domain activity (DDMN)
Monitor II report
field descriptions 3-25
option 3-24
sample 3-24
domain table tracing
expanded storage 5-134
DOMINO (Lotus Domino server activity)
Postprocessor report
summary 5-71
DP PR (dispatch priority) field
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
DSC % (disconnect time percentage) field
in Monitor III DEVR report 2-79
in Monitor III STORR report 2-171
DSD (detailed storage delays)
Monitor III report
description 2-84
summary 2-84
DSND (data set delays)
Monitor III report
command 2-86
field descriptions 2-87
report options panel 2-87
sample 2-87
summary 2-86
Monitor III reporter session command
summary 2-18
DSNJ (data set delays - Job report)
Monitor III report
command 2-89
field descriptions 2-90
sample 2-89
summary 2-89
Monitor III reporter session command
summary 2-18
DSNV (data set delays - Volume report)
Monitor III report
command 2-91
field descriptions 2-92
report options panel 2-92
sample 2-91
summary 2-91
Monitor III reporter session command
summary 2-18
duration field
in Monitor III RG report 2-147
in Monitor III SYSSUM report 2-199
duration report
description and format 5-2, 5-53
dynamic channel path management
description xx
Monitor II
CHANNEL report 3-20

dynamic channel path management *(continued)*
Monitor II *(continued)*
IOQUEUE report 3-40
Monitor III
CHANNEL report 2-55
IOQUEUE report 2-117
Postprocessor
CHANNEL report 5-35
IOQ report 5-92

E

earliest field
in the Monitor III DEVT report 2-83
in the Monitor III SYSTREND report 2-202
effective logical processor utilization field
in Monitor I CPU report 5-49
effective physical processor utilization field
in Monitor I CPU report 5-50
elapsed time (ELAP time) field
in Monitor III STORC report 2-161
ENCLAVE (enclave activities)
Monitor III report
command 2-94
field descriptions 2-95
sample 2-94
ENCLAVE (enclave activity)
Monitor III reporter session command
summary 2-18
enclave activities
Monitor III report (ENCLAVE)
command 2-94
field descriptions 2-95
sample 2-94
enclave support
DELAY report (Monitor III) 2-67
ENCLAVE report (Monitor III) 2-94
PROC report (Monitor III) 2-145
SYSINFO report (Monitor III) 2-180
end field
in Monitor I TRACE report 5-136
ENQ (enqueue activity)
Monitor I report
field descriptions 5-77, 5-78
sample (detail report) 5-76
ENQ (enqueue delays)
Monitor III report
command 2-99
field descriptions 2-99, 2-102
report options panel 2-102
sample 2-99
summary 2-99
Monitor III reporter session command
summary 2-18
ENQJ (enqueue delays job)
Monitor III report
command 2-128
field descriptions 2-128
summary 2-128
Monitor III reporter session command
summary 2-18
ENQR (enqueue resource delays)
Monitor III report
command 2-103
ENQR (enqueue resource delays) *(continued)*
Monitor III report *(continued)*
field descriptions 2-103, 2-104
report options panel 2-105
sample 2-103
summary 2-103
ENQR (resource-oriented enqueue delays)
Monitor III reporter session command
summary 2-18
enqueue activity
Monitor I report
field descriptions 5-77, 5-78
sample (detail report) 5-76
Monitor II report (SENQ)
field descriptions 3-55, 3-56
option 3-53
sample (detail) 3-55
sample (global resource) 3-54
sample (summary) 3-53
enqueue delays
Monitor III report (ENQ)
command 2-99
field descriptions 2-99, 2-102
report options panel 2-102
sample 2-99
summary 2-99
Monitor III report (SYSENQ)
command 2-177
field descriptions 2-178
sample 2-177
enqueue delays job
Monitor III report (ENQJ)
command 2-128
field descriptions 2-128
summary 2-128
enqueue resource delays (ENQR)
Monitor III report
command 2-103
field descriptions 2-103, 2-104
report options panel 2-105
sample 2-103
summary 2-103
enterprise systems connection (ESCON) channel 5-37
ES Move Rate (expanded storage move rate) field
in Monitor III STORJ report 2-133
ES RATE (expanded storage page-in rate) field
in Monitor III STORF report 2-167
ES RTE (expanded storage rate) field
in Monitor II SPAG report 3-61
ESCON channel 5-37
ESF (expanded storage frame) field
in Monitor I PAGING report 5-114
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
ESF AVL (expanded storage frames available) field
in Monitor II SPAG report 3-61
exception description field
in Postprocessor EXRPTS report 5-182
exception rate field
in Monitor II ARD report 3-11
in Monitor II ARDJ report 3-11
exception report
definition 5-1
description and format 5-179, 5-180
field descriptions 5-182
heading fields 5-181
Postprocessor report 5-179
sample 5-179, 5-180
exception threshold field
in Postprocessor EXRPTS report 5-182
EXCP rate field
in Monitor III GROUP report 2-109
execution data field
in Monitor III SYSRTD report 2-188
execution velocity (actual) field
in Monitor III SYSSUM report 2-195
execution velocity (EX VEL %) field
in Monitor I WKLD report 5-163
execution velocity (Ex Vel) field
in Monitor III SYSRTD report 2-188
execution velocity (goal) field
in Monitor III SYSSUM report 2-195
exiting Monitor III reporter session 2-14
EXP STOR (expanded storage) effective field
in Monitor I PAGING report 5-123
expanded storage (direct) field
in Monitor I PAGING report 5-123
expanded storage (EXP STOR) effective field
in Monitor I PAGING report 5-123
expanded storage (total) field
in Monitor I PAGING report 5-123
expanded storage frame (ESF) configuration field
in Monitor I PAGING report 5-114
expanded storage frames field
in Monitor I PAGING report 5-115
in Monitor II ASD report (ESF) 3-14
in Monitor II ASDJ report (ESF) 3-14
in Monitor II SPAG report (ESF AVL) 3-61
expanded storage move rate (ES Move Rate) field
in Monitor III STORJ report 2-133
expanded storage page-in rate (ES RATE) field
in Monitor III STORF report 2-167
expanded storage rate field
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
in Monitor II SPAG report (ES RTE) 3-61
exposures field
in Monitor III DEVJ report 2-126

F

FCD (FICON director activity)
Postprocessor report
summary 5-80
FF 16M (fixed pages below 16M) field
in Monitor II ARD report 3-10
in Monitor II ARDJ report 3-10
fixed frames field
in Monitor I PAGING report 5-118
in Monitor III STORJ report 2-133

fixed pages below 16M (FF 16M) field
in Monitor II ARD report 3-10
in Monitor II ARDJ report 3-10
frame occupancy field
in Monitor III STORF report 2-166
Frames ACT% (percentage of storage frames) field
in Monitor III GROUP report 2-108
frames online field
in Monitor III STORR report 2-169
free pages (bytes) field
in Monitor I VSTOR (common storage summary) report 5-143
in Monitor I VSTOR (private area summary) report 5-146
free storage
definition 5-140
freed without migration field
in Monitor I PAGING report 5-115

G

gathering data 1-2
getmain limit field
in Monitor I VSTOR (private area summary) report 5-146
goal field
in Monitor III SYSWKM report 2-208
goal type field
in Monitor III SYSSUM report 2-199
goals versus actual field
in Monitor III SYSSUM report 2-195
graphic reports, choosing 2-6
GROUP (group response time)
Monitor III report
command 2-106
field descriptions 2-107, 2-113
options panel 2-114
sample 2-106
summary 2-106
Monitor III reporter session command
summary 2-18
group field
in Monitor III GROUP report 2-107
in Postprocessor XCF report 5-177
group response time (GROUP)
Monitor III report
command 2-106
field descriptions 2-107, 2-113
options panel 2-114
sample 2-106
summary 2-106

H

H (hardcopy mode) command
output 3-6
heading fields
Monitor I session reports 5-4
help panels 2-15
heterogeneous report class period 2-192
HFS (HFS statistics)
Monitor II report
field descriptions 3-34
option 3-36
sample 3-34
HFS (HFS statistics) *(continued)*
Postprocessor report
field descriptions 5-83, 5-84, 5-85
sample 5-83, 5-84
summary 5-82
HFS statistics 3-34
HI UIC (highest unreferenced interval count) field
in Monitor II SPAG report 3-61
in Monitor II SRCS report 3-63
hierarchical storage manager delays (HSM)
Monitor III report
command 2-115
field descriptions 2-115, 2-116
report options panel 2-116
sample 2-115
summary 2-115
hierarchical storage manager delays job (HSMJ)
Monitor III report
command 2-129
field descriptions 2-129
summary 2-129
highest unreferenced interval count field
in Monitor I PAGING report (HIGH UIC) 5-114, 5-116
in Monitor II SPAG report (HI UIC) 3-61
in Monitor II SRCS report (HI UIC) 3-63
hiperspace and VIO (V&H I+O) field
in Monitor II SPAG report 3-61
holding field
in Monitor III ENQJ report 2-128
in Monitor III PROCJ report 2-131
holding job field
in Monitor III PROC report 2-145
holding name field
in Monitor III ENQR report 2-104
holding name/SYS field
in Monitor III ENQ report 2-102
holding percentage field
in Monitor III ENQ report 2-102
in Monitor III ENQR report 2-104
holding status field
in Monitor III ENQ report 2-102
in Monitor III ENQR report 2-104
homogeneous report class period 2-192
HSM (hierarchical storage manager delays)
Monitor III report
command 2-115
field descriptions 2-115, 2-116
report options panel 2-116
sample 2-115
summary 2-115
Monitor III reporter session command
summary 2-18
HSMJ (hierarchical storage manager delays job)
Monitor III report
command 2-129
field descriptions 2-129
summary 2-129
Monitor III reporter session command
summary 2-18
HTTP (HTTP server activity)
Postprocessor report
summary 5-87

I

I/O Device Activity
Monitor II report
field descriptions 3-29, 3-31
option 3-26
sample (table) 3-28
I/O interrupts (% I/O interrupts handled VIA TPI) field
in Monitor I CPU report 5-45
I/O queuing activity
Monitor I report
field descriptions 5-95, 5-96, 5-97, 5-98, 5-99
sample 5-94
Monitor II report
field descriptions 2-120, 3-42, 3-43, 3-44
option 3-40
sample 3-42
Monitor III report
field descriptions 2-118
option 2-117
sample 2-118
I/O requests rate (I/O REQ RATE) field
in Monitor II PGSP report 3-49
I/O total interrupt rate field
in Monitor I CPU report 5-44
I% (interval %) field
in Monitor II DEVV report 3-29
ICS (installation control specification) field
in Monitor III SYSINFO report 2-180
ID field
in Monitor III STORCR report 2-164
idle frames field
in Monitor III STORJ report 2-132
idling percentage (IDL %) field
in Monitor III DELAY report 2-64
ILOCK (IRLM long lock detection)
Monitor II report
field descriptions 3-38
sample 3-38
IMPLT field
in Monitor II DDMN report 3-25
importance (I) field
in Monitor III SYSSUM report 2-194
in (Q) queue field
in Monitor II SRCS report 3-63
in use (%) field
in Monitor I PAGESP report 5-106
in-storage current (INC) field
in Monitor II DDMN report 3-25
indicator field
in Monitor III WFEX report definition and criteria panel 2-225
in Monitor III WFEX report options action panel 2-222
installation control specification (ICS) field
in Monitor III SYSINFO report 2-180

installation performance specification (IPS) field
in Monitor III SYSINFO report 2-180
installed at field
in Monitor III SYSSUM report 2-194
INT MM.SS (interval length) field
in Postprocessor EXRPTS report 5-182
in Postprocessor SUMMARY report 5-203
Intelligent Resource Director (IRD) xix
interval % (I%) field
in Monitor II DEVV report 3-29
interval length (INT MM.SS) field
in Postprocessor EXRPTS report 5-182
in Postprocessor SUMMARY report 5-203
interval report
description 5-3
Introducing RMF Data Gathering and Reporting 1-1
IOCDS (I/O configuration data set) 5-92
IOQ (I/O queuing activity)
Monitor I report
field descriptions 5-95, 5-96, 5-97, 5-98, 5-99
sample 5-94
Monitor II report
field descriptions 2-120, 3-42, 3-43, 3-44
option 3-40
sample 3-42
Monitor III report
field descriptions 2-118
option 2-117
sample 2-118
Postprocessor report
summary 5-92
IOQUEUE (I/O queuing activity)
Monitor III reporter session command summary 2-18
IOS queue time field
in the Monitor III DEVN report (IosQ) 2-75
IOSQ time field
in Monitor II DEVV report 3-30
IPL definitions field
in Monitor III STORC report 2-159
IPS (installation performance specification) field
in Monitor III SYSINFO report 2-180
IRD xix
IRLM long lock detection 3-37

J

JES (job entry subsystem)
in workload reporting 5-151
JES delays
Monitor III report
command 2-121
field descriptions 2-122, 2-123
report options panel 2-123
sample 2-121
summary 2-121
JES delays *(continued)*
Monitor III reporter session command summary 2-18
JESJ (JES delays for specified job)
Monitor III reporter session command summary 2-18
JESJ (job entry subsystem delays job)
Monitor III report
command 2-129
field descriptions 2-129
summary 2-129
JOB command 2-124
Monitor III reporter session summary 2-18
JOB delays report
cross system coupling facility (XCF) delays variation 2-133
device (DEVJ) delays variation 2-125
enqueue (ENQJ) delays variation 2-128
hierarchical storage manager (HSM) delays variation 2-129
job entry subsystem (JES) delays variation 2-129
operator message (MSG) delays variation 2-130
operator mount (MNT) delays variation 2-130
processor (PROC) delays variation 2-130
quiesce (QSCJ) delays variation 2-131
storage (STOR) delays variation 2-132
job ended date field
in Monitor III STORCR report 2-164
job ended time field
in Monitor III STORCR report 2-164
job entry subsystem delays (JES)
Monitor III report
command 2-121
field descriptions 2-122, 2-123
report options panel 2-123
sample 2-121
summary 2-121
job entry subsystem delays job (JESJ)
Monitor III report
command 2-129
field descriptions 2-129
summary 2-129
job field
in Monitor III ENQJ report 2-128
in Monitor III HSMJ report 2-129
in Monitor III JESJ report 2-129
in Monitor III PROCJ report 2-131
job report menu 2-23
job report options panel 2-134
jobname field
in Monitor I VSTOR (private area summary) report 5-145
in Monitor II ARD report 3-10
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
in Monitor II ASRM report 3-18
in Monitor II ASRMJ report 3-18
in Monitor II SENQ report 3-56
in Monitor II SENQR report 3-58
in Monitor III DEV report 2-72
jobname field *(continued)*
in Monitor III DEVJ report (JOB) 2-125
in Monitor III DEVR report 2-79
in Monitor III ENQ report 2-100
in Monitor III HSM report 2-116
in Monitor III JES report 2-122
in Monitor III MSGJ report 2-130
in Monitor III PROC report 2-144
in Monitor III STOR report 2-156
in Monitor III STORCR report 2-164
in Monitor III STORF report 2-165
in Monitor III SYSWKM report 2-211
in Monitor III XCF report 2-233
Jobs (using/delay) field
in the Monitor III DEVN report 2-75
jobs at maximum contention field
in Monitor I ENQ report 5-78

K

keyboard A-1

L

L+C FF (LPA and CSA fixed frames) field
in Monitor II SRCS report 3-63
label field
in Monitor III WFEX report definition and criteria panel 2-225
in Monitor III WFEX report options action panel 2-222
largest free block field
in Monitor I VSTOR (common storage summary) report 5-143
in Monitor private area summary report 5-146
latest field
in the Monitor III DEVT report 2-83
in the Monitor III SYSTREND report 2-202
LCU field
in Monitor I DEVICE report 5-65
in Monitor I IOQ report 5-97
in Monitor II DEVV report 3-29
in Monitor II IOQUEUE report 3-43
in Monitor III IOQUEUE report 2-119
Licence manager
description xx
LIST structure in coupling facility 5-23
LLI (link library list)
Monitor II report
field descriptions 3-47
sample 3-45
LOCK structure in coupling facility 5-23
log swap effective field
in Monitor I PAGING report 5-123
log swap field
in Monitor I PAGING report 5-123
log swap/expanded storage field
in Monitor I PAGING report 5-123
logically swapped out (OUT LOG) field
in Monitor II SRCS report 3-63
LPA and CSA fixed frames (L+C FF) field
in Monitor II SRCS report 3-63

LPA fixed frames (LPA FF) field
in Monitor II SRCS report 3-63
LPA frames (LPA F) field
in Monitor II SRCS report 3-63
LPA in field
in Monitor II SPAG report 3-61
LPA rate field
in Monitor II ARD report 3-11
in Monitor II ARDJ report 3-11
LPAR cluster
description xx, 5-50
Postprocessor report 5-50
LPAR management physical processor utilization field
in Monitor I CPU report 5-50
LPAR mode
channel path activity report 5-36
CPU utilization 5-42
LSQA CSF field
in Monitor II ARD report 3-10
in Monitor II ARDJ report 3-10
LSQA ESF field
in Monitor II ARD report 3-10
in Monitor II ARDJ report 3-10

M

main delay path field
in Monitor III XCF report 2-233
main delay reason(s) field
in Monitor III HSM report 2-116
in Monitor III JES report 2-122
main delay volume(s) field
in Monitor III DEV report 2-72
major/minor names field
in Monitor III ENQ report 2-100
majorname field
in Monitor II SENQ report 3-56
in Monitor II SENQR report 3-59
maximum field
in Monitor I TRACE report 5-136
in Monitor II DDMN report (MAX) 3-25
maximum possible user region field
in Monitor I VSTOR (common storage summary) report 5-143
measurement scale for plot reports
modifying 5-199
measurements, common monitor III report 2-6
member name field
in Postprocessor XCF report 5-177
migrated from expanded storage field
in Monitor I PAGING report 5-115, 5-123
migration age field
in Monitor I PAGING report (MIGR AGE) 5-114
in Monitor II SPAG report (MIG AGE) 3-61
in Monitor III STORR report (MIGR AGE) 2-169
migration rate (MIG RTE) field
in Monitor II SPAG report 3-61
migration system information
Monitor III Utility report 2-136
minimum field
in Monitor I TRACE report 5-136
in Monitor II DDMN report (MIN) 3-25
minorname field
in Monitor II SENQ report 3-56
in Monitor II SENQR report 3-59
mintime 2-10
MNTJ (operator mount job delays)
Monitor III report
command 2-130
field descriptions 2-130
summary 2-130
Monitor III reporter session command
summary 2-18
mode
NOW BASIC 5-36
NOW LPAR 5-36
model number field
in Monitor III SYSINFO report 2-180
Monitor I session
report heading fields 5-4
reports 4-1, 5-2
system activity measured 1-4
Monitor II
reports 3-1, 3-66
Monitor II background session
format of printed output 3-7
system activity measured 1-4
Monitor II display session
commands
input area 3-4
format of printed output 3-6
header area 3-3
mode area 3-3
output area 3-3
panel format
header area 3-4
input area 3-4
mode area 3-4
output area 3-4
report name area 3-4
status/message area 3-4
report name area 3-3
status/message area 3-3
system activity measured 1-4
Monitor III
common report measurements 2-6
delay information 2-2
mintime 2-10
range 2-10
report interval 2-10
reports
description 2-26
example of using 2-5
how to use 2-1
shortened intervals, reason for 2-12
Monitor III data gatherer session
delay information 1-5
system activity measured 1-4
Monitor III reporter session
delay information 1-5
starting 2-14
system activity measured 1-4
Monitor III Reporter session commands
report and report options command summary 2-19
monitor III, commands you can use in 2-12
mount (MT) pend % field
in Monitor I DEVICE report 5-67
MSGJ (operator message job delays)
Monitor III report
command 2-130
field descriptions 2-130
summary 2-130
Monitor III reporter session command
summary 2-18
MSI (migration SYSINFO)
Monitor III report 2-136
description 2-136
field descriptions 2-137
sample 2-137
MT (mount) pend % field
in Monitor I DEVICE report 5-67
MVS partition name field
in Monitor I CPU report 5-48

N

name field
in Monitor I CPU report 5-48
in Monitor I ENQ report 5-77
in Monitor III DELAY report 2-63
in Monitor III PROC report 2-145
in Monitor III STORC report 2-160
in Monitor III SYSINFO report 2-181
in Monitor III SYSSUM report 2-194
in Monitor III WFEX report 2-217, 2-218, 2-220
in Monitor III WFEX report definition and criteria panel 2-226
in Postprocessor EXRPTS report 5-182
No. of completed transactions (For nnnn TRX) field
in Monitor III SYSWKM report 2-208
non-swappable in-storage (NSW) field
in Monitor II DDMN report 3-25
Non-VIO (NVI) rate field
in Monitor II ARD report 3-11
in Monitor II ARDJ report 3-11
NON-VIO slots field
in Monitor I PAGING report 5-118
not ready % field
in Monitor I DEVICE report 5-68
NOW BASIC mode 5-36
NOW LPAR mode 5-36
NSW (non-swappable in-storage) field
in Monitor II DDMN report 3-25
number field
in Monitor III DEVJ report 2-126
NUMBER IO REQ (number IO requests) field
in Monitor I PAGESP report 5-106
number IO requests (NUMBER IO REQ) field
in Monitor I PAGESP report 5-106
number of asids field
in Monitor I CPU report 5-45

number of configured partitions field
in Monitor I CPU report 5-48
number of intervals field
in Postprocessor SUMMARY report 5-203
number of log prcrs (logical processors) field
in Monitor I CPU report 5-49
number of mounts field
in Monitor I DEVICE report 5-68
number of physical processors field
in Monitor I CPU report 5-48
number of samples field
in Monitor I TRACE report 5-136
in Monitor I VSTOR (private area summary) report 5-145
number of sets field
in Monitor I TRACE report 5-136
NVI (Non-VIO) rate field
in Monitor II ARD report 3-11
in Monitor II ARDJ report 3-11

O

occurrences of output wait field
in Monitor I PAGING report 5-123
OMPLT field
in Monitor II DDMN report 3-25
OMVS (average) field
in Postprocessor SUMMARY report 5-204
OMVS (maximum) field
in Postprocessor SUMMARY report 5-204
OMVS (OMVS kernel activity)
Postprocessor report
command 5-101
field descriptions 5-102, 5-103
sample 5-102
summary 5-101
OMVS kernel activity report
Postprocessor report
command 5-101
field descriptions 5-102, 5-103
sample 5-102
summary 5-101
OPD (OMVS process data)
Monitor III reporter session command
summary 2-18
operator message job delays (MSGJ)
Monitor III report
command 2-130
field descriptions 2-130
summary 2-130
operator mount job delays (MNTJ)
Monitor III report
command 2-130
field descriptions 2-130
summary 2-130
option (OPT) field
in Monitor III SYSINFO report 2-180
option list, WLMGL 5-153
OUT LOG (logically swapped out) field
in Monitor II SRCS report 3-63
out ready queue (OUT RQ) field
in Monitor III SRCS report 3-63
out users queue (OUTO) field
in Monitor II DDMN report 3-25
out wait queue (OUT WQ) field
in Monitor II SRCS report 3-63
OVERVIEW
Postprocessor report
command 5-185
summary 5-185
overview report
definition 5-1
Overview report
Postprocessor report
command 5-185
summary 5-185

P

page in field
in Monitor I PAGING report 5-110
page movement time % field
in Monitor I PAGING report 5-111
page movement within central storage field
in Monitor I PAGING report 5-111
page out field
in Monitor I PAGING report 5-111
page space type field
in Monitor I PAGESP report 5-105
page transfer time field
in Monitor I PAGESP report 5-106
in Monitor II PGSP report (TRAN TIME PAGE) 3-49
page-in events field
in Monitor I PAGING report 5-111
page-in rate field
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
in Monitor III GROUP report (PgIn Rate) 2-109
in Monitor III STORF report (PGIN RATE) 2-167
in Monitor III STORJ report (Page In Rate) 2-133
in Monitor III STORS report (PGIN RATE) 2-174
page/swap data set activity
Monitor I report
field descriptions 5-105, 5-106
sample (page data section) 5-105
Monitor II report (PGSP) 3-48
field descriptions 3-49, 3-50
option 3-48
pageable system areas (non-VIO) field
in Monitor I PAGING report 5-109
pages (XFER'D) transferred field
in Monitor I PAGESP report 5-106
pages allocated (in bytes) field
in Monitor I VSTOR (private area summary) report 5-147
pages swapped (PGS-SWPD) in field
in Monitor II SPAG repor 3-61
pages swapped (PGS-SWPD) out field
in Monitor II SPAG report 3-61
PAGESP (page/swap data set activity)
Monitor I report
field descriptions 5-105, 5-106
sample (page data section) 5-105
paging
Monitor II report (SPAG) 3-60
field descriptions 3-61
option 3-60
sample 3-60
PAGING (paging activity)
Monitor I report
field descriptions 5-109, 5-110, 5-111, 5-124
sample (central storage) 5-107
sample (expanded storage/frame and slot) 5-113
sample (swap placement) 5-120
paging activity
Monitor I report
field descriptions 5-109, 5-110, 5-111, 5-124
sample (central storage) 5-107
sample (expanded storage/frame and slot) 5-113
sample (swap placement) 5-120
partition effective dispatch time field
in Monitor I CPU report 5-49
partition total dispatch time field
in Monitor I CPU report 5-49
partition utilization field
in Monitor I CHAN report 5-39
in Monitor II channel path activity (CHANNEL) report 2-57, 3-22
path number field
in Monitor III XCFJ report 2-133
paths unavailable field
in Postprocessor XCF report 5-175
PAV (parallel access volume)
in Monitor I DEVICE report 5-64
peak allocation values field
in Monitor III STORC report 2-159
pend reasons field
in Monitor III STORR report 2-171
pend time field
in Monitor II DEVV report 3-30
pend time percentage field
in Monitor III DEVJ report 2-127
in Monitor III DEVR report (PND % Reasons) 2-79
in Monitor III STORR report (PND %) 2-171
in the Monitor III DEVN report (PND % Reasons) 2-75
Pending (pend time percentage) field
in the Monitor III DEVN report 2-75
percent delayed for field
in Monitor III DELAY report (% Delayed for) 2-65
in Monitor III STOR report 2-156
percent used field
in Monitor III STORC report 2-161
percentage frames (% Frames) field
in Monitor III STORR report 2-169
percentage of storage frames field
in Monitor III GROUP report 2-108
performance group number field
in Monitor I WKLD report (PGN) 5-159
in Monitor II ASD report (PG) 3-14
in Monitor II ASDJ report (PG) 3-14
in Monitor II ASRM report (PG) 3-19

performance group number field
(continued)
in Monitor II ASRMJ report (PG) 3-19
in Monitor II TRX report (PERF GRP) 3-66
in Monitor III DEVR report (PG) 2-80
in Monitor III STOR report (PG) 2-156
in Monitor III STORC report (PG) 2-160
in Monitor III STORS report (PG) 2-174
in Monitor III XCF report (PG) 2-233
performance group period (PGP) field
in Monitor I WKLD report 5-159
performance group period field
in Monitor II ASD report (PP) 3-14
in Monitor II ASDJ report (PP) 3-14
in Monitor II ASRM report (PP) 3-19
in Monitor II ASRM report (SP) 3-19
in Monitor II ASRMJ report (PP) 3-19
in Monitor II ASRMJ report (SP) 3-19
in Monitor II TRX report (PERF PER) 3-66
performance index 2-196
period field
in Monitor III GROUP report 2-107
in Monitor III SYSWKM report 2-208
PGS-SWPD (pages swapped) in field
in Monitor II SPAG report 3-61
PGS-SWPD (pages swapped) out field
in Monitor II SPAG report 3-61
phase (P) field
in Monitor III SYSWKM report 2-209
plot
Postprocessor report 5-191
description and format 5-191, 5-199
heading fields 5-191
modifying the measurement scale 5-199
types 5-192
sample 5-191
plot report
definition 5-1
PLPLA intermodule space field
in Monitor I VSTOR common storage summary report 5-143
PLPLA space redundant field
in Monitor I VSTOR (common storage summary) report 5-143
PND % (pend time percentage) field
in Monitor III DEVR report 2-79
in Monitor III STORR report 2-171
policy report
Postprocessor report 5-156
possible cause or action field
in Monitor III WFEX report 2-220
possible causes field
in Monitor III XCF report 2-232
Postprocessor
reports 5-1, 5-180
system activity measured 1-4
preallocated datasets
in DSINDEX report 2-27
PRI FF (private fixed frames) field
in Monitor II SRCS report 3-63
primary delay category field
in Monitor III HSMJ report 2-129
in Monitor III JESJ report 2-129
primary delay field
in Monitor III DEVJ report 2-125
in operator message job delays (MSGJ) report 2-130
primary delay percentage field
in Monitor III HSMJ report 2-129
in Monitor III JESJ report 2-129
primary delay reason field
in Monitor III HSMJ report 2-129
in Monitor III JESJ report 2-129
primary menu for Monitor III 2-15
primary reason field
in Monitor III DELAY report 2-66
primary response time component field
in Monitor III GROUP report 2-108
printed report 2-26, 2-233
Monitor I 5-3
Monitor II 3-1, 3-66
Postprocessor 5-1, 5-180
PRIV FF (private fixed pages) field
in Monitor II ARD report 3-10
in Monitor II ARDJ report 3-10
PRIV_IN BLK (private area blocked in) field
in Monitor II SPAG report 3-61
PRIV_IN NBK (private area non-blocked in) field
in Monitor II SPAG report 3-61
private area
detail report 5-147
Monitor I VSTOR (summary) report 5-145
partial private area data 5-147
private area blocked in (PRIV_IN BLK) field
in Monitor II SPAG report 3-61
private area non-blocked in (PRIV_IN NBK) field
in Monitor II SPAG report 3-61
private area out (PRV OUT) field
in Monitor II SPAG report 3-61
private fixed frames (PRI FF) field
in Monitor II SRCS report 3-63
private fixed pages (PRIV FF) field
in Monitor II ARD report 3-10
in Monitor II ARDJ report 3-10
private frames/ central storage field
in Monitor II SRCS report 3-63
private frames/ expanded storage field
in Monitor II SRCS report 3-63
probable cause field
in Monitor III DEVJ report 2-125
PROC (processor delays)
Monitor III report
command 2-143
field descriptions 2-143, 2-145
report options panel 2-145
sample 2-143
summary 2-143
Monitor III reporter session command
summary 2-18
processor activity
Monitor I report 5-41
processor delays
Monitor III report (PROC)
command 2-143
field descriptions 2-143, 2-145
report options panel 2-145
sample 2-143
summary 2-143
processor job delays (PROCJ)
Monitor III report
command 2-130
field descriptions 2-131
summary 2-130
processor using percentage field
in Monitor III PROCJ report 2-131
in Monitor III SYSWKM report 2-211
PROCJ (processor job delays)
Monitor III report
command 2-130
field descriptions 2-131
summary 2-130
Monitor III reporter session command
summary 2-18
program name field
in private area summary report 5-145
PRV OUT (private area out) field
in Monitor II SPAG report 3-61

Q

QLEN (queue length) distribution % field
in Monitor I ENQ report 5-78
QSCJ (quiescent jobs)
Monitor III reporter session command
summary 2-18
qualifier field
in Monitor III WFEX report definition and criteria panel 2-224
in Monitor III WFEX report options action panel 2-222
queue length (QLEN) distribution % field
in Monitor I ENQ report 5-78
quiesce delay (Quies) field
in Monitor III SYSWKM report 2-211

R

R LS (reason-last swap out) field
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
range 2-10
Range field
in the Monitor III DEVT report 2-83
in the Monitor III SYSTREND report 2-202
read from EXP STOR (expanded storage) field
in Monitor I PAGING report 5-114
read from REAL STOR (central storage) field
in Monitor I PAGING report 5-117
real storage frames field
in Monitor I PAGING report 5-117

reason field
in Monitor III WFEX report 2-218
reason-last swap out (R LS) field
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
Reasons field
in Monitor III DEVR report 2-79
in the Monitor III DEVN report 2-75
refresh field
in Monitor III SYSSUM report 2-194
region assigned field
in Monitor I VSTOR (private area summary) report 5-146
region requested field
in Monitor I VSTOR (private area summary) report 5-146
report class period
heterogeneous 2-192
homogeneous 2-192
Report Class Period report 5-157
Report Class report 5-157
Report Class Report 5-157
report commands
DEVN (device activity) 2-74
DEVT (device activity trend) 2-82
PROC (processor delays report) 2-143
STOR (storage delays report) 2-155
STORC (common storage) 2-158
STORCR (common storage remaining) 2-163
STORF (storage frames report) 2-165
STORR (storage resource delays report) 2-168
STORS (storage delay summary) 2-173
SYSINFO (system information) 2-179
SYSRTD (response time distribution) 2-184
SYSSUM (sysplex summary) 2-189
SYSTREND (system trend) 2-201
SYSWKM (work manager delays) 2-204
WFEX (workflow/exceptions) 2-213
XCF (cross-system coupling facility) 2-232
report commands, Monitor III
DEV (device) 2-71
DEVJ (device delays job report) 2-125
DEVR (device resource delays) 2-77
DI 2-26
ENQJ (enqueue delays job report) 2-128
ENQR (enqueue resource delays) 2-103
enqueue (ENQ) 2-99
GROUP (group response time report) 2-106
HSM (hierarchical storage manager) 2-115
HSMJ (hierarchical storage manager delays job report) 2-129
JES (job entry subsystem) 2-121
JESJ (job entry subsystem delays job report) 2-129
JOB 2-124
report commands, Monitor III *(continued)*
MNTJ (operator mount job delays report) 2-130
MSGJ (operator message job delays report) 2-130
PROCJ (processor job delays report) 2-130
QSCJ (quiesce delays variation of job report) 2-131
STORJ (storage job delays report) 2-132
XCFJ (cross system coupling facility delays job report) 2-133
report interval 2-10
processing 5-150
report options (Monitor III) 2-12
report performance groups 5-150
request (REQ) field
in Monitor II SENQ report 3-56
in Monitor II SENQR report 3-59
request from field
in Postprocessor XCF report 5-177
request in field
in Postprocessor XCF report 5-176, 5-177, 5-178
request out field
in Postprocessor XCF report 5-175, 5-177, 5-178
request reject field
in Postprocessor XCF report 5-176
request to field
in Postprocessor XCF report 5-177
request type field
in Monitor I ENQ report 5-78
requested delay field
in Monitor III DEVJ report 2-125
in operator message job delays (MSGJ) report 2-130
Reserve Activity
Monitor II report (SENQR)
field descriptions 3-58
option 3-58
sample 3-58
reserve status (RSV) field
in Monitor II SENQR report 3-59
reserved device % (%D RV) field
in Monitor II DEVV report 3-31
resource group data
Monitor III report 2-146
Monitor III Utility report 2-146
resource group field
in Monitor III RG report 2-147
in Monitor III SYSSUM report 2-199
resource name field
in Monitor III ENQR report 2-104
resource waiting percentage field
in Monitor III ENQ report 2-100
resource waiting status field
in Monitor III ENQ report 2-100
resource workflow (%) 2-8
resource-oriented reports
device activity report 2-74
device activity trend report 2-82
resource-oriented storage delays (STORR)
Monitor III report
command 2-168
field descriptions 2-168, 2-172
resource-oriented storage delays (STORR) *(continued)*
Monitor III report *(continued)*
sample 2-169
summary 2-168
response time (actual) field
in Monitor III SYSSUM report 2-195
response time (goal) field
in Monitor III SYSSUM report 2-195
Response Time ACT line
in Monitor III GROUP report 2-112
response time breakdown field
in Monitor III SYSWKM report 2-209
response time breakdown field (delayed by)
in Monitor III SYSWKM report 2-210
response time distribution (SYSRTD)
Monitor III report
command 2-184
field descriptions 2-188
option panel 2-188
sample 2-185
response time field
in Monitor II DDMN report (RTOB) 3-25
in Monitor II DEVV report 3-29
in Monitor III DEVR report (Resp Time) 2-78
in Monitor III GROUP report 2-109
in Monitor III SYSINFO report (RESP Time) 2-181
in Monitor III SYSRTD report 2-188
in the Monitor III DEVN report (RspT) 2-75
retry field
in Postprocessor XCF report 5-178
RLSDS (VSAM RLS activity by data set)
Monitor III reporter session command summary 2-18
RLSLRU (VSAM RLS activity LRU statistics)
Monitor III reporter session command summary 2-18
RLSSC (VSAM RLS activity by storage class)
Monitor III reporter session command summary 2-18
RMF (Resource Measurement Facility)
help panels 2-15
tutorial 2-15
RMF command 2-14
row and position field
in Monitor III WFEX report options action panel 2-222
RSV (reserve status) field
in Monitor II SENQR report 3-59
RTOB (response time) field
in Monitor II DDMN report 3-25
RUA (average ready users) field
in Monitor II DDMN report 3-25

S

S (shared) field
in Monitor III DEVR report 2-78
S T (space type) field
in Monitor II PGSP report 3-49

samples field
in Monitor I PAGING report 5-118
samples/set field
in Monitor I TRACE report 5-136
SCPER 5-155
SDEVICE (shared device activity)
Postprocessor report
command 5-125
field descriptions 5-127
sample 5-125, 5-126
summary 5-125
SDS (sysplex data server activity)
Monitor II report
field descriptions 3-52
option 3-51
sample 3-51
seconds/set field
in Monitor I TRACE report 5-136
selecting a report 2-15
selection criteria
in the Monitor III DEVN report 2-75
selection line, tailoring 2-15
SENQ (enqueue activity)
Monitor II report
field descriptions 3-55, 3-56
option 3-53
sample (detail) 3-55
sample (global resource) 3-54
sample (summary) 3-53
SENQR (reserve activity)
Monitor II report
field descriptions 3-58
option 3-58
sample 3-58
server field
in Monitor III ENQJ report 2-128
service class (SRVCLASS) field
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
service class field
in Monitor III SYSWKM report 2-211
Service Class Period report 5-155
Service Class Report 5-155
service definition field
in Monitor III SYSSUM report 2-194
service field
in Monitor III SYSWKM report 2-211
Service Policy Page in WLMGL report 5-158
service rate
in Postprocessor SUMMARY report 5-204
SESS (SRM service total) field
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
session options xv
shared device report
Postprocessor report (SDEVICE)
command 5-125
field descriptions 5-127, 5-132
sample 5-125, 5-126
summary 5-125
shared field
in Monitor III DEVJ report 2-126
in Monitor III DEVR report (S) 2-78
SHARED SLOTS AND FRAMES
in Monitor I PAGING report 5-119
shortcut keys A-1
shortened intervals, reason for 2-12
SHR field
in Monitor I CHAN report 5-38
in Monitor II CHANNEL report 3-21
in Monitor III CHANNEL report 2-57
signalling path busy (SUSP) field
in Postprocessor XCF report 5-178
single state case
defined for Using 2-9
defined for Workflow 2-7
single-system environment, using MIII reports in 2-3
single-system reports 5-3
slots allocated field
in Monitor I PAGESP report 5-105
slots used field
in Monitor I PAGESP report 5-105
software pricing
Postprocessor CPU report 5-46
sort criteria
in the Monitor III DEVN report 2-75
space type field
in Monitor II PGSP report (S T) 3-49
in Monitor III STORR report 2-171
speed field
in Monitor III WFEX report 2-217
speedometer
in Monitor III WFEX report 2-221
SQA + ESQA (system queue area) overflow field
in Monitor III STORR report 2-170
SQA expansion into CSA field
in Monitor I VSTOR common storage summary report 5-143
SQA frames (SQA F) field
in Monitor II SRCS report 3-63
SRB field
in the Monitor III SYSTREND report 2-202
SRCS (central storage/processor/SRM) activity
Monitor II report
field descriptions 3-63
option 3-62
sample 3-62
SRM absorption (ABS) rate field
in Monitor II ARD report 3-11
in Monitor II ARDJ report 3-11
SRM service total (SESS) field
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
SRVCLASS field
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
standard deviation (STD. DEV.) field
in Monitor I TRACE report 5-136
starting and ending an RMF reporter session 2-14
static storage map field
in Monitor I VSTOR common storage summary report 5-143
status field
in Monitor I CPU report 5-48
in Monitor III ENQJ report 2-128
STC (started task control)
subsystem reporting 5-151
STC users (average) field
in Postprocessor SUMMARY report 5-204
STC users (maximum) field
in Postprocessor SUMMARY report 5-204
STD. DEV. (standard deviation) field
in Monitor I TRACE report 5-136
STG GRP (storage group) field
in Monitor II DEVV report 3-29
STOR (storage delays)
Monitor III report
command 2-155
field descriptions 2-155, 2-157
report options panel 2-157
sample 2-155
summary 2-155
Monitor III reporter session command
summary 2-18
storage delay summary
Monitor III report (STORS)
command 2-173
field descriptions 2-174
options panel 2-175
sample 2-173
summary 2-173
storage delays
Monitor III report (STOR)
command 2-155
field descriptions 2-155, 2-157
report options panel 2-157
sample 2-155
summary 2-155
storage frames
Monitor III report (STORF)
command 2-165
field descriptions 2-165, 2-167
report options panel 2-167
sample 2-165
summary 2-165
storage group
Monitor I DEVICE report 5-60
Monitor II DEV report 3-28
storage group field
in Monitor I DEVICE report 5-64
in Monitor II DEVV report (STG GRP) 3-29
storage job delays
Monitor III report (STORJ)
command 2-132
field descriptions 2-132, 2-133
summary 2-132
storage protection xxi
storage summary field
in Monitor III STORS report 2-174
STORC (common storage)
Monitor III report
command 2-158
field descriptions - job information 2-160, 2-161
field descriptions - system information 2-159, 2-160
options panel 2-162
sample 2-159
summary 2-158

STORC (common storage) *(continued)*
Monitor III reporter session command summary 2-19
STORCR (common storage remaining)
Monitor III report
command 2-163
field descriptions 2-164
sample 2-163
summary 2-163
Monitor III reporter session command summary 2-19
STORF (storage frames)
Monitor III report
command 2-165
field descriptions 2-165, 2-167
report options panel 2-167
sample 2-165
summary 2-165
Monitor III reporter session command summary 2-19
STORJ (storage job delays)
Monitor III report
command 2-132
field descriptions 2-132, 2-133
summary 2-132
Monitor III reporter session command summary 2-19
STORR (resource-oriented storage delays)
Monitor III reporter session command summary 2-19
STORR (storage resource delays)
Monitor III report
command 2-168
field descriptions 2-168, 2-172
sample 2-169
summary 2-168
STORS (storage delay summary)
Monitor III report
command 2-173
field descriptions 2-174
options panel 2-175
sample 2-173
summary 2-173
STORS (storage delays summary)
Monitor III reporter session command summary 2-19
subsystem (SUBSYS) field
in Monitor II TRX report 3-65
subsystem data field
in Monitor III SYSRTD report 2-188
subsystem menu 2-25
subsystem reports 2-25
subsystem type field
in Monitor III SYSWKM report 2-209
summary
definition 5-1
Postprocessor report 5-202
description and format 5-202, 5-204
field descriptions 5-203, 5-204
heading fields 5-202
sample 5-202
swap out field
in Monitor II SPAG report 3-61
swap rate field
in Monitor II ARD report 3-11
in Monitor II ARDJ report 3-11
swap rate field *(continued)*
in Postprocessor SUMMARY report 5-204
swap RV (recommendation value) field
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
swap sets allocated field
in Monitor I PAGESP report 5-105
swap sets used field
in Monitor I PAGESP report 5-105
switched time field
in Monitor III SYSWKM report 2-210
SYSENQ (sysplex enqueue delays)
Monitor III report
command 2-177
field descriptions 2-178
sample 2-177
Monitor III reporter session command summary 2-19
SYSINFO (system information)
Monitor III report
command 2-179
field descriptions 2-180, 2-182
sample 2-180
summary 2-179
Monitor III reporter session command summary 2-19
sysplex data server activity
Monitor II report (SDS) 3-51
sysplex environment, using MIII reports in 2-2
Sysplex Report Selection Menu 2-21
sysplex reports 5-3
sysplex summary
Monitor III report (SYSSUM)
command 2-189
field descriptions 2-194
report options panel 2-200
sample 2-190, 2-193
SYSRTD (response time distribution)
Monitor III report
command 2-184
field descriptions 2-188
option panel 2-188
sample 2-185
Monitor III reporter session command summary 2-19
SYSSUM (sysplex summary)
Monitor III report
command 2-189
field descriptions 2-194
report options panel 2-200
sample 2-190, 2-193
Monitor III reporter session command summary 2-19
system activities measured
Monitor I session 1-4
Monitor II session 1-4
Monitor III session 1-4
Postprocessor 1-4
system field
in Monitor II SENQ report 3-56
in Monitor II SENQR report 3-58
in Monitor III ENQJ report 2-128
in Monitor III SYSWKM report 2-188
system information
Monitor III report (SYSINFO)
command 2-179
field descriptions 2-180, 2-182
sample 2-180
summary 2-179
Monitor III report (SYSTREND)
command 2-201
field descriptions 2-202
sample 2-201
summary 2-201
system name (from) field
in Postprocessor XCF report 5-176
system name (to) field
in Postprocessor XCF report 5-175
system name field
in Monitor III SYSWKM report 2-211
in Postprocessor XCF report 5-177, 5-178
system queue area (SQA + ESQA) overflow field
in Monitor III STORR report 2-170
SYSTREND (system trend)
Monitor III report
command 2-201
field descriptions 2-202
sample 2-201
summary 2-201
SYSWKM (work manager delays)
Monitor III report
command 2-204
field descriptions 2-212
report options panel 2-212
sample 2-204, 2-205
Monitor III reporter session command summary 2-19

T

tabular reports, choosing 2-6
tape rate field
in Postprocessor SUMMARY report 5-203
target working set size field
in Monitor II ASD report (TAR WSS) 3-14
in Monitor II ASDJ report (TAR WSS) 3-14
in Monitor II SPAG report (TAR CWS) 3-61
tasks own (TSK OWN) field
in Monitor II SENQ report 3-55
tasks shared waiting (TSK WTS) field
in Monitor II SENQ report 3-55
tasks waiting (TSK WTE) field
in Monitor II SENQ report 3-55
TCB field
in the Monitor III SYSTREND report 2-202
TCB time field
in Monitor II ARD report 3-11
in Monitor II ARDJ report 3-11
TCB+SRB seconds field
in Monitor I WKLD report 5-163
text field
in Monitor III WFEX report definition and criteria panel 2-225

time device allocated field
in Monitor I DEVICE report 5-68
time field
in Monitor I TRACE report 5-136
in Monitor II ARDJ report 3-10
in Monitor II ASRMJ report 3-19
in Monitor II SPAG report 3-61
in Monitor II SRCS report 3-63
in Monitor III SYSSUM report 2-195
in Postprocessor EXRPTS report 5-182
in Postprocessor SUMMARY report 5-203
Time field
in the Monitor III DEVT report (Time) 2-83
in the Monitor III SYSTREND report (Time) 2-202
time overflow 5-61, 5-62
time slice group
definition 5-150
time slice group field
in Monitor I WKLD report 5-159
Time Stamp field
in the Monitor III DEVT report (Time) 2-83
in the Monitor III SYSTREND report (Time) 2-202
time weighted service (TWSR) field
in Monitor II DDMN report 3-25
top of allocated area field
in Monitor I VSTOR (private area summary) report 5-146
total event field
in Monitor I ENQ report 5-78
total field
in Monitor I CPU report 5-50
in Monitor I PAGING report 5-123
total frames field
in Monitor I PAGING report 5-118
total length of intervals field
in Postprocessor SUMMARY report 5-203
total logical processor utilization field
in Monitor I CPU report 5-49
total physical processor utilization field
in Monitor I CPU report 5-50
total slots field
in Monitor I PAGING report 5-119
total system field
in Monitor I PAGING report 5-109
total utilization (%) field
in Monitor I CHAN report 5-39
total utilization field
in Monitor II channel path activity (CHANNEL) report 2-57, 3-22
total/average field
in Monitor I CPU report 5-45
in Postprocessor SUMMARY report 5-204
TRACE (trace activity)
Monitor I report
field descriptions 5-136
sample 5-135
trace activity (TRACE)
Monitor I report
field descriptions 5-136
trace activity (TRACE) *(continued)*
Monitor I report *(continued)*
sample 5-135
TRAN TIME PAGE (page transfer time) field
in Monitor II PGSP report 3-49
TRANS CUR RES (transaction resident time) field
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
TRANS RATE (transaction rate) field
in Monitor II TRX report 3-66
TRANS/SEC (transactions per second) field
in Monitor III SYSINFO report 2-181
transaction (TX) CPU field
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
transaction (TX) MSO field
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
transaction (TX) SRB field
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
transaction active field
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
transaction activity
Monitor II report (TRX)
field descriptions 3-65, 3-66
option 3-64
sample 3-64
transaction class (TRXCLASS) field
in Monitor II TRX report 3-65
transaction count (TX CT) field
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
Transaction ended rate field
in Monitor III GROUP report 2-109
transaction I/O service consumed (TX IOC) field
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
transaction name (TRXNAME) field
in Monitor II TRX report 3-66
transaction rate (Trx Rate) field
in Monitor III SYSRTD report 2-188
transaction rate field
in Monitor II TRX report (TRANS RATE) 3-66
in Postprocessor SUMMARY report 5-204
transaction resident time (TRANS CUR RES) field
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
transaction response time field
in Monitor III SYSINFO report (RESP Time) 2-181
transaction swap count (TX SC) field
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
transactions per second (TRANS/SEC) field
in Monitor III SYSINFO report 2-181
transport class field
in Postprocessor XCF report 5-175, 5-176, 5-178
TRX (transaction activity)
Monitor II report
field descriptions 3-65, 3-66
option 3-64
sample 3-64
TRXCLASS (transaction class) field
in Monitor II TRX report 3-65
TRXNAME (transaction name) field
in Monitor II TRX report 3-66
TSK OWN (tasks own) field
in Monitor II SENQ report 3-55
TSK WTE (tasks waiting) field
in Monitor II SENQ report 3-55
TSK WTS (tasks shared waiting) field
in Monitor II SENQ report 3-55
TSO sessions (average) field
in Postprocessor SUMMARY report 5-203
TSO sessions (maximum) field
in Postprocessor SUMMARY report 5-203
TSO/E (Time Sharing Option Extensions) session
subsystem reporting 5-151
tutorial 2-15
TWSR (time weighted service) field
in Monitor II DDMN report 3-25
TX (transaction) CPU field
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
TX (transaction) MSO field
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
TX (transaction) SRB field
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
TX CT (transaction count) field
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
TX IOC (transaction I/O service consumed) field
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
TX SC (transaction swap count) field
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14
in Monitor II ASRM report 3-19
in Monitor II ASRMJ report 3-19
type (T) field
in Monitor III SYSSUM report 2-194
type field
in Monitor I CPU report 5-45
Type field
in the Monitor III DEVT report 2-83

U

UKN % (unknown percentage) field
in Monitor III DELAY report 2-64
unknown percentage (UKN %) field
in Monitor III DELAY report 2-64
user menu 2-25
user reports 2-25

user reports *(continued)*
Monitor III samples
DSD (detailed storage delays) 2-84
MSI (migration SYSINFO) 2-136
RGD (resource group data) 2-146
userid field
in Monitor II TRX report 3-65
users field
in Monitor III GROUP report 2-108
in Monitor III STORS report 2-174
in Monitor III SYSINFO report 2-181
users/active field
in Monitor III WFEX report 2-217
using
cursor-sensitive control 2-5
monitor III reports 2-1
samples
defined for Workflow of a resource 2-6
Using (%) of an address space
definition of 2-8
using percentage field
in Monitor III DELAY report (USG %) 2-63
in Monitor III DEV report (USG %) 2-72
in Monitor III DEVR report (USG %) 2-80
in Monitor III HSMJ report 2-129
in Monitor III JESJ report 2-129
in Monitor III PROC report (USG %) 2-144
using/delay (USG/DEL) field
in the Monitor III DEVN report 2-75

V

V (VIO) field
in Monitor II PGSP report 3-50
V&H I+O (hiperspace and VIO) field
in Monitor II SPAG report 3-61
V&H RT (VIO and hiperspace rate) field
in Monitor II ARD report 3-11
in Monitor II ARDJ report 3-11
vector utilization field
in Monitor III GROUP report (Vector UTIL) 2-108
velocity field
in Monitor III SYSWKM report 2-211
VF affinity percentage field
in Monitor I CPU report 5-43
VIO and hiperspace rate (V&H RT) field
in Monitor II ARD report 3-11
in Monitor II ARDJ report 3-11
VIO field
in Monitor I PAGESP report 5-106
in Monitor II PGSP report (V) 3-50
VIO slots field
in Monitor I PAGING report 5-118
virtual storage activity
changing the sample to cycles ratio 5-141
Monitor I report (VSTOR)
field descriptions 5-143, 5-147
sample (common storage detail) 5-144
virtual storage activity *(continued)*
Monitor I report (VSTOR) *(continued)*
sample (private area detail) 5-147
sample (private area summary) 5-145
VolSer field
in the Monitor III DEVT report 2-83
VOLSER field
in Monitor II DEVV report 3-29
volume field
in Monitor II SENQR report 3-59
in Monitor III DEVR report 2-78
volume serial field
in Monitor I DEVICE report 5-64
in Monitor I PAGESP report 5-105
in Monitor II PGSP report 3-49
in Monitor III STORR report 2-170
VSTOR (virtual storage activity)
Monitor I report 5-140
common storage summary section 5-142
sample (common storage summary) 5-142

W

wait completion field
in Monitor I CPU report 5-48
wait time percentage field
in Monitor I CPU report 5-44
weights field
in Monitor I CPU report 5-49
WFEX (workflow/exceptions)
Monitor III report
command 2-213
criteria names selection panel 2-227
definition and criteria panel 2-223
definition of UNAVAIL panel 2-227
field descriptions 2-216, 2-220, 2-221, 2-223, 2-227
graphic report 2-216
sample 2-215, 2-216
summary 2-213, 2-216, 2-223, 2-227
Monitor III report option
action panel 2-221
command 2-221
field descriptions 2-222
summary 2-221
Monitor III reporter session command
summary 2-19
WFEX report options definition of UNAVAIL panel 2-227
WFL % (workflow percentage) field
in Monitor III DELAY report 2-63
in Monitor III SYSINFO report 2-181
the Monitor III GROUP report 2-108
WGPER 5-157
WKLD (workload activity)
Monitor I report
options 5-149
sample 5-152
WLMGL 5-153
WLMGL (workload activity)
Postprocessor report
command 5-153
field descriptions 5-158, 5-171
sample 5-153
summary 5-153
WLMGL option list 5-153
work manager delays (SYSWKM)
Monitor III report
command 2-204
field descriptions 2-208, 2-212
report options panel 2-212
sample 2-204, 2-205
workflow (%)
of a resource 2-8
of an address space 2-7, 2-8
workflow percentage field
in Monitor III DELAY report (WFL %) 2-63
in Monitor III HSMJ report 2-129
in Monitor III JESJ report 2-129
in Monitor III SYSINFO report (WFL %) 2-181
the Monitor III GROUP report (WFL %) 2-108
workflow/exceptions (WFEX)
criteria names selection panel
sample 2-227
summary 2-227
definition of unavail panel
sample 2-227
summary 2-227
Monitor III report
command 2-213
criteria names selection panel 2-227
definition and criteria panel 2-223
definition of UNAVAIL panel 2-227
field descriptions 2-216, 2-220, 2-221, 2-223, 2-227
graphic report 2-216
sample 2-215, 2-216
summary 2-213, 2-216, 2-223, 2-227
working set field
in Monitor III STOR report 2-157
in Monitor III STORJ report 2-132
workload activity
Monitor I report
options 5-149
sample 5-152
policy summary report 5-156
Postprocessor report (WLMGL)
command 5-153
field descriptions 5-158, 5-171
sample 5-153, 5-156
summary 5-153
service class period report 5-155
service class report 5-156
service policy page 5-158
workload group and service class/period report 5-157
workload group report 5-156
workload group and service class/period report (WGPER)
Postprocessor report 5-157

workload group report
Postprocessor report 5-156
written to EXP STOR (expanded storage) field
in Monitor I PAGING report 5-114
written to REAL STOR (central storage) field
in Monitor I PAGING report 5-116
WSM recommendation value field
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14

X

XCF (cross-system coupling facility)
Monitor III report
command 2-232
field descriptions 2-233
report options panel 2-233
sample 2-232
summary 2-232
Monitor III reporter session command
summary 2-19
Postprocessor report
field descriptions 5-175, 5-176, 5-177, 5-178
sample 5-174
sample (path statistics) 5-177
sample (usage by member) 5-176
XCFJ (cross system coupling facility job delays)
Monitor III report
command 2-133
field descriptions 2-133
summary 2-133
Monitor III reporter session command
summary 2-19
XM (cross memory) field
in Monitor II ARD report 3-11
in Monitor II ARDJ report 3-11
in Monitor II ASD report 3-14
in Monitor II ASDJ report 3-14

Y

yellow and red (Yel Red) fields
in Monitor III WFEX report definition and criteria panel 2-226

Readers' Comments — We'd Like to Hear from You

z/OS
Resource Measurement Facility
Report Analysis

Publication No. SC33-7991-05

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? ☐ Yes ☐ No

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

Name

Address

Company or Organization

Phone No.

Readers' Comments — We'd Like to Hear from You
SC33-7991-05

Cut or Fold
Along Line

Fold and Tape | **Please do not staple** | Fold and Tape

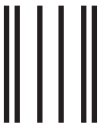

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Deutschland Entwicklung GmbH
eServer Performance Management Development
Schoenaicher Strasse 220
D-71032 Boeblingen
Federal Republic of Germany

Fold and Tape | **Please do not staple** | Fold and Tape

Cut or Fold
Along Line

Program Number: 5694–A01, 5655–G52

Printed in U.S.A.

SC33-7991-05

Spine information:

z/OS

z/OS V1R4 RMF Report Analysis

SC33-7991-05