

z/OS
3.2

*Resource Measurement Facility
Report Analysis*



Note

Before using this information and the product it supports, read the information in [“Notices” on page 527.](#)

Last updated: 2026-04-29

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

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

© **Rocket Software, Inc. 2019, 2025.**

Contents

Figures.....	xiii
Tables.....	xxv
About this document.....	xxxv
Who should use this document.....	xxxv
How this document is organized.....	xxxv
z/OS RMF library.....	xxxv
z/OS Data Gatherer library.....	xxxvi
z/OS information.....	xxxvi
How to provide feedback to IBM.....	xxxvii
Summary of changes.....	xxxix
Summary of changes for z/OS Resource Measurement Facility Report Analysis 3.2.....	xxxix
Summary of changes for z/OS Resource Measurement Facility Report Analysis 3.1.....	xl
Chapter 1. Introducing RMF data gathering and reporting.....	1
Gathering data.....	1
Short-term data collection with Monitor III.....	1
Snapshot monitoring with Monitor II.....	1
Long-term data gathering with Monitor I and Monitor III.....	1
Long-term overview reporting with the Postprocessor.....	2
Report analysis with the Spreadsheet Reporter.....	2
Monitoring on the workstation.....	2
What you can gather and report.....	3
Activity monitoring.....	3
Delay monitoring.....	4
Long-term performance analysis with RMF XP.....	4
Reporting of other SMF data.....	5
OpenMetrics with Monitor III.....	5
Chapter 2. Interactive performance analysis with Monitor III.....	7
Using Monitor III reports.....	7
System activities measured.....	7
Where to start.....	7
Using cursor-sensitive control.....	10
Common Monitor III report measurements.....	10
Monitor III MINTIME and range.....	14
Monitor III report options.....	15
Monitor III sysplex support in different time zones.....	16
The reporter session.....	16
Starting and stopping a Monitor III reporter session.....	17
The Primary Menu.....	17
Selecting a report.....	17
Monitor III report commands - Overview.....	18
Header for single-system reports.....	21
Header for sysplex reports.....	22
The Sysplex Report Selection Menu.....	22

The Overview Report Selection Menu.....	23
The Job Report Selection Menu.....	23
The Resource Report Selection Menu.....	24
The Subsystem Report Selection Menu.....	24
The User Report Selection Menu.....	25
The Data Index.....	25
How to request the Data Index.....	26
Contents of the Data Index.....	26
CACHDET - Cache Detail Report.....	32
How to request this report.....	32
Contents of the report.....	32
CACHSUM - Cache Summary Report.....	36
How to request this report.....	36
Contents of the report.....	36
CFACT - Coupling Facility Activity Report.....	39
How to request this report.....	39
Special considerations.....	39
Contents of the report.....	40
CFOVER - Coupling Facility Overview Report.....	47
How to request this report.....	47
Contents of the report.....	47
CFSYS - Coupling Facility Systems Report.....	49
How to request this report.....	49
Contents of the report.....	49
CHANNEL - Channel Path Activity Report.....	52
How to request this report.....	53
Special considerations of report output.....	53
Contents of the report.....	53
CPC - CPC Capacity Report.....	56
How to request this report.....	56
Contents of the report.....	56
CRYACC – Crypto Accelerator Activity Report.....	59
How to request this report.....	60
Contents of the report.....	60
CRYOVW – Crypto Hardware Overview Report.....	61
How to request this report.....	61
Contents of the report.....	61
CRYPKC – Crypto PKCS11 Coprocessor Activity Report.....	64
How to request this report.....	64
Contents of the report.....	65
DELAY - Delay Report.....	66
How to request this report.....	66
Contents of the report.....	66
DEV - Device Delays Report.....	73
How to request this report.....	73
Contents of the report.....	74
DEVN - Device Activity Report.....	75
How to request this report.....	76
Contents of the report.....	76
DEVR - Device Resource Delays Report.....	78
How to request this report.....	78
Contents of the report.....	78
DEVT - Device Activity Trend Report.....	82
How to request this report.....	83
Contents of the report.....	83
DSD - Detailed Storage Delays Report.....	84
How to request this report.....	84
Contents of the report.....	84

DSND - Data Set Delays Report.....	85
How to request this report.....	86
Special considerations.....	86
Contents of the report.....	86
DSNJ - Data Set Delays - Job Report.....	88
How to request this report.....	88
Contents of the report.....	88
DSNV - Data Set Delays - Volume Report.....	89
How to request this report.....	89
Contents of the report.....	90
EADM - Extended Asynchronous Data Mover (EADM) Activity Report.....	91
How to request this report.....	91
Contents of the report	91
ENCLAVE - Enclave Report.....	93
How to request this report.....	94
Contents of the report.....	94
Field descriptions.....	94
Cursor-sensitive control on the Enclave Report.....	95
Report options.....	98
ENQ - Enqueue Delays Report.....	98
How to request this report.....	99
Contents of the report.....	99
ENQR - Enqueue Resource Delays Report.....	103
How to request this report.....	103
Contents of the report.....	103
GROUP - Group Response Time Report.....	105
How to request this report.....	105
Contents of the report.....	106
HSM - Hierarchical Storage Manager Delays Report.....	112
How to request this report.....	112
Contents of the report.....	113
IOQUEUE - I/O Queuing Activity Report.....	114
How to request this report.....	115
Special considerations of report output.....	115
Contents of the report.....	115
JES - Job Entry Subsystem Delays Report.....	117
How to request this report.....	118
Contents of the report.....	118
JOB - Job Delay Report.....	120
How to request this report.....	120
Contents of the report.....	120
Displaying the Job Delays report.....	121
The top and bottom parts of the Job Delays report.....	121
Device Delay variation.....	122
Enqueue Delay variation.....	123
HSM and JES variations.....	124
Operator Message and Mount Delay variations.....	125
Processor Delay variation.....	125
Quiesce variation.....	126
Storage Delay variation.....	127
XCF variation.....	127
LOCKSP - Spin Lock Report.....	129
How to request this report.....	129
Contents of the report.....	129
LOCKSU - Suspend Lock Report.....	130
How to request this report.....	131
Contents of the report.....	131
OPD - OMVS Process Data Report.....	132

How to request this report.....	132
Contents of the report.....	133
PCIE - PCIE Activity Report.....	137
How to request this report.....	138
Contents of the report.....	138
Cursor sensitive control.....	139
PROC - Processor Delays Report.....	145
How to request this report.....	145
Contents of the report.....	146
PROCU - Processor Usage Report.....	148
How to request this report.....	148
Contents of the report.....	149
RLSLRU - VSAM LRU Overview Report.....	151
How to request this report.....	151
Contents of the report.....	151
RLSSC/RLSDS - VSAM RLS Activity Report.....	153
How to request this report.....	153
Contents of the report.....	153
SPACED - Disk Space Report.....	157
How to request this report.....	158
Contents of the report.....	158
SPACEG - Storage Space Report.....	158
How to request this report.....	159
Contents of the report.....	159
STOR - Storage Delays Report.....	159
How to request this report.....	159
Contents of the report.....	160
STORC - Common Storage Report.....	162
How to request this report.....	162
Contents of the report.....	163
STORCR - Common Storage Remaining Report.....	166
How to request this report.....	166
Contents of the report.....	167
STORF - Storage Frames Report.....	167
How to request this report.....	168
Contents of the report.....	168
STORM - Storage Memory Objects Report.....	170
How to request this report.....	170
Contents of the report.....	170
STORR - Storage Resource Delays Report.....	172
How to request this report.....	172
Contents of the report.....	173
STORS - Storage Delay Summary Report.....	176
How to request this report.....	177
Contents of the report.....	177
SYSENQ - Sysplex Enqueue Delays Report.....	180
How to request this report.....	180
Contents of the report.....	180
SYSINFO - System Information Report.....	181
How to request this report.....	181
Contents of the report.....	182
SYSRG - Resource Group Activity Report.....	186
How to request this report.....	187
Contents of the report.....	187
SYSRTD - Response Time Distribution Report.....	190
How to request this report.....	190
Contents of the report.....	190
SYSSUM - Sysplex Summary report.....	193

How to request this report.....	194
Contents of the report.....	194
SYSTREND - System Trend Report.....	203
How to request this report.....	203
Contents of the report.....	203
SYSWKM - Work Manager Delays Report.....	205
How to request this report.....	205
Contents of the report.....	206
USAGE - Monitor III Job USAGE Report.....	212
How to request this report.....	212
Contents of the report.....	212
Field descriptions.....	213
Monitor III Utility fields.....	214
Cursor-sensitive control on the USAGE Report.....	215
WFEX - Workflow/Exceptions Report.....	216
How to request this report.....	216
Contents of the report.....	217
Workflow/Exceptions graphic Report.....	231
XCF - Cross-System Coupling Facility Delays Report.....	232
How to request this report.....	232
Contents of the report.....	232
XCFGROUP - XCF Group Data.....	233
How to request the XCF Group Data report.....	233
Contents of the XCF Group Data report.....	234
XCFOVW - XCF Overview.....	234
How to request the XCF Overview report.....	234
Contents of the XCF Overview report.....	235
XCFPATH - XCF Path statistics.....	235
How to request the XCF Path statistics report.....	235
Contents of the XCF Path statistics report.....	236
XCFSYS - XCF Systems data.....	237
How to request the XCF Systems data report.....	237
Contents of the XCF Systems data report.....	237
ZFSFS - zFS File System.....	238
How to request this report.....	238
Contents of the report.....	240
ZFSKN - zFS Kernel report.....	243
How to request this report.....	243
Contents of the report.....	244
ZFSOVW - zFS Overview Report.....	244
How to request this report.....	245
Contents of the report.....	245

Chapter 3. Snapshot reporting with Monitor II.....251

Monitor II sessions.....	251
Structure of Monitor II reports.....	252
Contents of the Monitor II report header.....	252
Different formats of Monitor II reports.....	254
Different modes of Monitor II reports.....	254
ARD/ARDJ - Address Space Resource Data report.....	256
How to request the ARD report.....	256
Contents of the report.....	257
ASD/ASDJ - Address Space State Data report.....	260
How to request the ASD report.....	260
Contents of the report.....	260
ASRM/ASRMJ - Address Space SRM Data report.....	264
How to request this report.....	264

Contents of the report.....	265
Report options.....	267
CHANNEL - Channel Path Activity report.....	267
How to request this report.....	267
Special considerations of report output.....	267
Contents of the report.....	268
DEV/DEVV - Device Activity report.....	270
How to request this report.....	270
Special considerations of report output.....	271
Contents of the report.....	272
HFS - Hierarchical File System Statistics report.....	276
How to request this report.....	276
Contents of the report.....	277
ILOCK - IRLM Long Lock Detection report.....	279
How to request this report.....	279
Special considerations.....	279
Contents of the report.....	280
IOQUEUE - I/O Queuing Activity report.....	281
How to request this report.....	282
Special considerations of report output.....	282
Contents of the report.....	283
LLI - Library List report.....	285
How to request this report.....	285
Contents of the report - Link Library List.....	286
Contents of the report - LPA Library List.....	287
Contents of the report - APF Library List.....	287
OPT - OPT Settings report.....	288
How to request this report.....	288
Contents of the report.....	289
PGSP - Page Data Set Activity report.....	290
How to request this report.....	290
Special considerations of report output.....	290
Contents of the report.....	290
SDS - Sysplex Data Server report.....	292
How to request this report.....	292
Contents of the report.....	292
SENQ - System Enqueue Contention report.....	293
How to request this report.....	293
Different types of SENQ reports.....	294
Contents of the report.....	295
SENQR - System Enqueue Reserve report.....	297
How to request this report.....	298
Contents of the report.....	298
SPAG - Paging Activity report.....	299
How to request this report.....	299
Contents of the report.....	300
SRCS - Central Storage/Processor/SRM report.....	301
How to request this report.....	301
Contents of the report.....	302

Chapter 4. Real-time reporting with Monitor I.....305

Chapter 5. Long-term overview reporting with the Postprocessor.....307

Interval and duration reports.....	308
Single-system report header.....	311
Sysplex report header.....	311
Duration report header.....	312

Overview condition names.....	312
CACHE - Cache Subsystem Activity report.....	312
How to request this report.....	312
Different report levels.....	313
Cache Summary reporting.....	313
Subsystem-level reporting.....	314
Device-level reporting.....	320
Spreadsheet and Overview reference.....	321
CF - Coupling Facility Activity report.....	324
How to request this report.....	324
Contents of the report.....	324
Spreadsheet and Overview reference.....	340
CHAN - Channel Path Activity report.....	341
How to request this report.....	342
Contents of the report.....	342
Spreadsheet and Overview reference.....	346
CPU - CPU Activity report.....	347
How to request this report.....	348
Contents of the report.....	348
Using the information in the CPU Activity report.....	350
Using the information in the Partition Data Report.....	357
Using the information in the LPAR Cluster Report.....	363
Using the information in the Group Capacity Report.....	365
Using the information in the Hardware Group Report.....	366
Using the information in the Power Consumption report.....	367
Duration report.....	369
CRYPTO - Crypto Hardware Activity report.....	369
How to request this report.....	370
Contents of the report.....	370
Spreadsheet and Overview reference.....	375
DEVICE - Device Activity report.....	377
How to request this report.....	377
Contents of the report.....	378
Spreadsheet and Overview reference.....	388
DOMINO - Lotus Domino Server report.....	390
How to request this report.....	390
Contents of the report.....	390
EADM - EADM Activity Report.....	393
How to request this report.....	393
Contents of the report.....	394
ENQ - Enqueue Activity report.....	395
How to request this report.....	396
Different report levels.....	396
Contents of the report.....	397
Messages.....	399
ESS - Enterprise Disk Systems report.....	399
How to request this report.....	401
Contents of the report.....	401
FCD - FICON Director Activity report.....	407
How to request this report.....	407
Contents of the report.....	408
HFS - Hierarchical File System Statistics report.....	409
How to request this report.....	409
Contents of the report.....	410
Special considerations.....	412
HTTP - HTTP Server report.....	413
How to request this report.....	413
Contents of the report.....	413

IOQ - I/O Queuing Activity report.....	417
How to request this report.....	418
Using the information given in the report.....	418
Data gathering considerations.....	418
Missing data in report fields.....	419
Messages.....	419
Contents of the report.....	420
Spreadsheet and Overview reference.....	427
OMVS - OMVS Kernel Activity report.....	428
How to request this report.....	428
Contents of the report.....	429
PAGESP - Page Data Set Activity report.....	431
How to request this report.....	431
Using the information given in the report.....	431
Contents of the report.....	431
Overview reference.....	433
PAGING - Paging Activity report.....	433
How to request this report.....	433
Using the information given in the report.....	433
Contents of the report.....	434
PCIE - PCIE Activity Report.....	447
How to request this report.....	447
Contents of the report.....	447
Spreadsheet and Overview reference.....	452
SDELAY - Serialization Delay report.....	453
How to request this report.....	453
Contents of the report.....	454
SDEVICE - Shared Device Activity report.....	461
How to request this report.....	461
Using the information given in the report.....	461
Shared Direct Access Device Activity report.....	461
Shared Magnetic Tape Device Activity report.....	462
Spreadsheet and Overview reference.....	467
TRACE - Trace Activity report.....	467
How to request this report.....	467
Using the information given in the report.....	468
Contents of the report.....	468
VSTOR - Virtual Storage Activity report.....	472
How to request this report.....	472
Contents of the report.....	473
Data gathering considerations.....	473
Overview reference.....	475
Common Storage Detail section.....	475
Private Area Summary section.....	476
Private Area Detail section.....	479
High Virtual Memory Usage section.....	479
WLMGL - Workload Activity report.....	480
How to request this report.....	480
The WLMGL option list.....	480
Contents of the report.....	481
Service Class Period report.....	481
Service Class report.....	485
Workload Group report.....	486
Workload Group and Service Class Period report.....	486
Policy Summary report.....	487
Report Class report.....	487
Report Class Period report.....	488
Service Policy page.....	489

Field descriptions of all reports.....	490
Spreadsheet and Overview reference.....	506
XCF - Cross-System Coupling Facility Activity report.....	509
How to request this report.....	509
Contents of the report.....	509
Exception report.....	513
How to request this report.....	513
Contents of the report.....	514
Improved exception reporting.....	516
Overview report.....	517
How to request this report.....	517
Report description.....	517
Summary report.....	521
How to request this report.....	521
Generating a report.....	521
Special considerations of report output.....	522
Spreadsheet reference.....	523
Appendix A. Accessibility.....	525
Notices.....	527
Terms and conditions for product documentation.....	528
IBM Online Privacy Statement.....	529
Policy for unsupported hardware.....	529
Minimum supported hardware.....	529
Programming interface information.....	530
Trademarks.....	530
Glossary.....	531
Index.....	537

Figures

1. Suggested Sequence for Using Monitor III Reports.....	9
2. Monitor III Primary Menu.....	17
3. Header of Monitor III Single-System Reports.....	21
4. Header of Monitor III Sysplex Reports.....	22
5. Monitor III Sysplex Report Selection Menu.....	22
6. Monitor III Overview Report Selection Menu.....	23
7. Monitor III Job Report Selection Menu.....	24
8. Monitor III Resource Report Selection Menu.....	24
9. Monitor III Subsystem Report Selection Menu.....	25
10. Monitor III User-written Report Selection Menu.....	25
11. Data Index.....	27
12. Data Index with Preallocated Data Sets - Detailed View.....	28
13. Data Index - Condensed Version.....	28
14. Data Index with Preallocated Data Sets - Condensed Version.....	28
15. Data Index Options Panel.....	31
16. Data Index - Sort Order Descend.....	31
17. CACHDET Report.....	32
18. CACHDET Report - Volume Details (pop-up panel).....	33
19. RMF Cache SSID Detail (pop-up panel).....	33
20. CACHDET Report Options.....	36
21. CACHSUM Report.....	36
22. CACHSUM Report - SSID Details.....	37
23. CFACT Report.....	40

24. CFACT Report - Details for a Lock Structure.....	40
25. CFACT Report - Details for a List Structure.....	41
26. CFACT Report - Details for a Cache Structure.....	41
27. CFACT Report - Details for a List Structure using SCM storage (1).....	42
28. CFACT Report - Details for a List Structure using SCM storage (2).....	42
29. Coupling Facility Report Options Panel.....	46
30. CFOVER Report.....	47
31. CFSYS Report.....	49
32. CFSYS Report - Subchannels and Paths.....	50
33. CHANNEL Report.....	53
34. CPC Capacity report.....	57
35. Crypto Accelerator Activity Report.....	60
36. Crypto Hardware Overview Report.....	62
37. Crypto Report Options Panel.....	64
38. Crypto PKCS11 Coprocessor Activity Report.....	65
39. DELAY Report.....	66
40. DELAY Report Options Panel.....	71
41. DELAY Report Job Selection/Exclusion Panel.....	73
42. DEV Report.....	74
43. DEV Report Options Panel.....	75
44. DEVN Report.....	76
45. DEVR Report.....	78
46. DEVR Report Options Panel.....	82
47. DEVT Report.....	83
48. Modified STOR Report Showing all Storage Delays in Detail.....	85

49. DSND Report.....	86
50. DSND Report Options Panel.....	87
51. DSNJ Report.....	88
52. DSNV Report.....	90
53. DSNV Report Options Panel.....	91
54. EADM Activity Report.....	92
55. ENCLAVE Report.....	94
56. ENCLAVE Report - Enclave Details.....	96
57. ENCLAVE Report - Enclave Classification Attributes (1).....	97
58. ENCLAVE Report Options.....	98
59. ENQ Report.....	99
60. ENQR Report.....	103
61. ENQR Report Options Panel.....	104
62. GROUP Report.....	106
63. Response time Components data (pop-up panel).....	106
64. GROUP Report Options Panel.....	112
65. HSM Report.....	113
66. IOQUEUE Report.....	115
67. JES Delays report.....	118
68. Top Part of Job Delays report.....	121
69. Bottom Part of Job Delay report.....	122
70. Device Delay variation of the Job Delay report.....	122
71. Enqueue Delay variation of the Job Delay report.....	124
72. HSM Delay variation of Job Delay report.....	124
73. Operator Message Delay variation of the Job Delay report.....	125

74. Processor Delay variation of the Job Delay report.....	126
75. Quiesce Delay variation of the Job Delay report.....	126
76. Storage Delay variation of Job Delay report.....	127
77. XCF Delay variation of Job Delay report.....	128
78. Job Report Options Panel.....	129
79. Spin Lock Report.....	130
80. Lock Report Options.....	130
81. Suspend Lock Report.....	131
82. OPD Report.....	133
83. OPD Report - Details for Server Process.....	134
84. OPD Report Options Panel.....	137
85. PCIE Activity Report.....	138
86. RMF Hardware Accelerator And Compression Activity panel.....	140
87. RMF RoCE Activity panel.....	141
88. RMF NETH Activity panel.....	142
89. RMF Internal Shared Memory Activity panel.....	143
90. RMF Synchronous I/O Link Activity panel.....	143
91. PROC - Processor Delay Report.....	146
92. PROCU - Processor Usage Report.....	149
93. RLRLRU Report.....	151
94. VSAM LRU Overview - Buffer Counts by Pool.....	152
95. VSAM RLS Activity by Storage Class - Sysplex Total View.....	154
96. VSAM RLS Activity by Storage Class - System/CF Structure View.....	154
97. VSAM RLS Activity by Data Set - Sysplex Total View.....	155
98. VSAM RLS Activity by Data Set - System/CF Structure View.....	156

99. Disk Space Report.....	158
100. Storage Space Report.....	159
101. STOR Report.....	160
102. STORC Report.....	163
103. STORC Report Options Panel.....	166
104. STORCR Report.....	167
105. STORF Report.....	168
106. STORM Report.....	170
107. STORR Report.....	173
108. STORS Report.....	177
109. STORS Report Options Panel.....	179
110. SYSENQ Report.....	180
111. SYSINFO Report.....	182
112. SYSRG Resource Group Activity report.....	187
113. Pop-up panel for a resource group.....	187
114. Pop-up panel for a tenant group.....	188
115. Pop-up panel for resource group activity report options.....	189
116. SYSRTD Report - With Response Time Data.....	190
117. SYSRTD Report - Without Response Time Data.....	191
118. SYSSUM Report - GO Mode.....	194
119. SYSSUM Report - STOP Mode.....	195
120. SYSSUM Report for Workload Groups.....	197
121. SYSTREND Report.....	204
122. SYSWKM Report for Subsystem CICS.....	206
123. SYSWKM Report for Subsystem IMS.....	206

124. SYSWKM Report Options Panel.....	211
125. Job Usage report.....	213
126. WFEX Report.....	217
127. WFEX Report Options Action Panel.....	221
128. WFEX Definition and Criteria Panel.....	223
129. WFEX Definition of UNAVAIL Panel.....	226
130. WFEX Criteria Names for Class: Job.....	228
131. WFEX Criteria Names for Class: Service Class.....	229
132. WFEX Criteria Names for Class: Processor.....	229
133. WFEX Criteria Names for Class: Storage.....	230
134. WFEX Criteria Names for Class: Device.....	230
135. WFEX Graphic Report.....	231
136. XCF Report.....	232
137. XCF Group Data Report.....	234
138. XCF Overview Report.....	235
139. XCF Path statistics Report.....	236
140. XCF Systems Data Report.....	237
141. ZFSFS - Report Options.....	239
142. ZFSFS Report.....	240
143. ZFSFS Report - File system details.....	240
144. ZFSFS Report - zFS File System Details.....	241
145. zFS Kernel Report.....	244
146. zFS Overview Report.....	245
147. I/O Details by Type.....	247
148. User Cache Details.....	247

149. Vnode Cache Details.....	249
150. Metadata Cache Details.....	250
151. Header of a Monitor II TSO/E session report.....	253
152. Format of printed reports from a Monitor II display session.....	255
153. Format of printed reports from a Monitor II background session.....	256
154. ARD Report in z/Architecture.....	257
155. ARD, ASD, and ASRM Report Options Panel.....	259
156. ARDJ, ASDJ, and ASRMJ Report Options Panel.....	259
157. ASD Report.....	261
158. ASDJ Report.....	261
159. ASRM Report.....	265
160. ASRMJ Report.....	266
161. CHANNEL Report.....	268
162. DEV Report.....	272
163. DEVV Report.....	272
164. DEV Report Options Panel.....	275
165. DEVV Report Options Panel.....	276
166. HFS Report.....	277
167. HFS Report Options Panel.....	278
168. ILOCK Report.....	280
169. IOQUEUE Report.....	283
170. IOQUEUE Report Options Panel.....	285
171. LLI Report - Link Library List.....	286
172. LLI Report - LPA Library List.....	287
173. LLI Report - APF Library List.....	287

174. OPT Settings.....	289
175. PGSP Report.....	290
176. SDS Report.....	292
177. SENQ Summary Report.....	295
178. SENQ GRS Report.....	295
179. SENQ Detail Report.....	295
180. SENQ Report Options Panel.....	297
181. SENQR Report.....	298
182. SENQR Report Options Panel.....	299
183. SPAG Report.....	300
184. SRCS Report.....	302
185. Header of a Single-System Report.....	311
186. Header of a Sysplex report.....	311
187. Header of a Single-System Duration Report.....	312
188. Cache Subsystem Summary Report - Subsystem Summary.....	314
189. Cache Subsystem Activity Report - Top-20 Device Lists.....	314
190. Cache Subsystem Activity Report - Status and Overview.....	315
191. Cache Subsystem Activity Report - Device Overview.....	319
192. Cache Subsystem Activity Report - Cache Device Activity (device-level reporting).....	320
193. Structure of the Coupling Facility Activity report.....	325
194. Coupling Facility Activity Report - Usage Summary.....	326
195. Coupling Facility Activity Report - Structure Activity.....	333
196. Coupling Facility Activity Report - Subchannel Activity.....	336
197. Coupling Facility Activity Report - CF to CF Activity.....	339
198. Channel Path Activity report.....	343

199. CPU Activity Report.....	351
200. CPU Activity Report - Partition Data Report.....	359
201. CPU Activity Report - LPAR Cluster Report.....	363
202. CPU Activity Report - Group Capacity Report.....	365
203. CPU Activity Report – Hardware Group Report.....	367
204. CPU Activity Report - Power Consumption Report.....	367
205. Crypto Hardware Activity Report.....	371
206. Direct Access Device Activity Report.....	379
207. Synchronous I/O Device Activity.....	379
208. Synchronous I/O Device Activity	380
209. Communication Equipment Activity Report.....	380
210. Communication Equipment Activity Report.....	380
211. Magnetic Tape Device Activity Report.....	381
212. Lotus Domino Server Report - Summary.....	390
213. Lotus Domino Server Report - Details.....	391
214. Extended Asynchronous Data Mover (EADM) Activity Report.....	394
215. Enqueue Detail Activity Report.....	397
216. ESS Link Statistics.....	402
217. ESS Synchronous I/O Link Statistics.....	403
218. ESS Extent Pool Statistics.....	403
219. ESS Rank Statistics.....	404
220. FICON Director Activity Report.....	408
221. HFS Global Statistics Report.....	410
222. HFS (HFS statistics)Postprocessor reportsampleHFS File System Statistics Report.....	411
223. HTTP Server Report - Summary.....	413

224. HTTP Server Report - Details.....	414
225. I/O Queuing Activity Report.....	420
226. OMVS Kernel Activity Report.....	429
227. PAGESP Report.....	431
228. PAGING Report - Central Storage Paging Rates.....	434
229. PAGING Report - Central Storage Movement and Request Rates / Frame and Slot Counts.....	437
230. PAGING Report - Memory Objects and High Virtual Storage Frames.....	442
231. PAGING Report - Dedicated memory.....	445
232. Postprocessor PCIE Activity Report.....	449
233. SDELAY Report - Serialization Delay Summary.....	455
234. SDELAY Report - Serialization Delay Details - CMS Lock Details.....	457
235. SDELAY Report - Serialization Delay Details - CML and Local Lock Details.....	458
236. SDELAY Report - Serialization Delay Details - GRS Latch Details.....	459
237. SDELAY Report - Serialization Delay Details - GRS Enqueue Details.....	460
238. Shared DASD Activity Report.....	462
239. Shared Magnetic Tape Device Activity Report.....	462
240. TRACE Report.....	468
241. VSTOR report - Common Storage Summary.....	474
242. VSTOR report - Common Storage Detail.....	476
243. VSTOR report - Private Area Summary.....	477
244. VSTOR report - Private Area Detail and High Virtual Memory Usage (above 2GB).....	479
245. WLMGL - Service Class Period report - with execution velocity goal for Period 1 and response time distributions.....	482
246. WLMGL - Service Class Period report - with response time distribution (percentile).....	483
247. WLMGL - Service Class Period report - with subsystem data and response time distribution for response time goal.....	484
248. WLMGL - Service Class Period report with transaction data for enclaves.....	485

249. WLMGL - Service Class Report.....	485
250. WLMGL - Workload Group with associated service classes.....	486
251. WLMGL - Report Class Report.....	488
252. WLMGL Report - Service Policy Page.....	489
253. XCF Activity Report - Usage by System.....	510
254. XCF Activity Report - Usage by Member.....	511
255. XCF Activity Report - Path Statistics (Coupling Facility and Channel-to-Channel).....	512
256. Exception Report - Low CPU Utilization.....	514
257. Exception Report - CPU Utilization.....	515
258. Overview Report - Exception Version.....	516
259. Overview Report - Summary Version.....	519
260. Summary Report.....	522

Tables

1. Monitored activities and SMF record types.....	3
2. Report Commands.....	19
3. User-Written Report Commands.....	21
4. Monitor III Heading Information.....	22
5. Fields in the Data Index.....	29
6. Fields in the CACHDET Report.....	33
7. Fields in the CACHDET Report - Volume and SSID Details.....	34
8. Fields in the CACHSUM Report.....	37
9. Fields in the CACHSUM Report - SSID Details.....	38
10. Fields in the CFACT Report.....	42
11. Fields in the CFOVER Report.....	47
12. Fields in the CFSYS Report.....	50
13. Fields in the CFSYS Report - Subchannels and Paths.....	51
14. Fields in the CHANNEL Report.....	54
15. Additional Fields in the CHANNEL Report.....	55
16. Fields in the CPC Capacity Report.....	57
17. Fields in the Cryptographic accelerator activity Report.....	60
18. Fields in the Crypto Hardware Overview Report.....	63
19. Fields in the Crypto PKCS11 Coprocessor Activity Report.....	65
20. Fields in the DELAY Report.....	67
21. Additional Fields in the DELAY Report.....	71
22. Fields in the DEV Report.....	74
23. Fields in the DEVN Report.....	77

24. DEVN Report - Cursor-sensitive Control for Select and Sort.....	77
25. Fields in the DEVR Report.....	79
26. Additional Fields in the DEVR Report.....	81
27. Fields in the DEVT Report.....	83
28. Fields in the DSD Report.....	85
29. Fields in the DSND Report.....	86
30. Fields in the DSNJ Report.....	88
31. Fields in the DSNV Report.....	90
32. Fields in the EADM Activity Report.....	92
33. Fields in the ENCLAVE Report.....	94
34. Fields in the ENCLAVE Report - Enclave Details.....	96
35. Fields in the ENQ Report.....	99
36. Fields in the ENQR Report.....	104
37. Fields in the GROUP Report.....	106
38. Fields in the HSM Report.....	113
39. Fields in the IOQUEUE Report.....	116
40. Fields in the JES Report.....	118
41. Fields in the Job Delays report.....	121
42. Fields in the Device Delay variation of the Job Delay report.....	122
43. Fields in the Enqueue Delay variation of the Job Delay report.....	124
44. Fields in the HSM/JES delay variation of the Job Delay report.....	124
45. Fields in the Operator Message and Mount Delay variations of the Job Delay report.....	125
46. Fields in the Processor Delay variation of the Job Delay report.....	126
47. Fields in the Quiesce Delay variation of the Job Delay report.....	126
48. Fields in the Storage Delays variation of the Job Delay report.....	127

49. Fields in the XCF Delay variation of the Job Delay report.....	128
50. Additional Fields in the Job Delay report.....	128
51. Fields in the LSP Report.....	130
52. Fields in the LSU Report.....	131
53. Fields in the OPD Report.....	133
54. Fields in the OPD Details Report.....	134
55. Fields in the PCIE Activity Report.....	138
56. Monitor III Utility fields.....	139
57. Fields in the RMF Hardware Accelerator And Compression Activity pop-up panel.....	140
58. Fields in the RMF RoCE and NETH Activity pop-up panels.....	142
59. Fields in the RMF Internal Shared Memory Activity pop-up panel.....	143
60. Fields in the RMF Synchronous I/O Link Activity pop-up panel.....	143
61. Fields in the PROC Report.....	146
62. Additional Fields in the PROC Report.....	148
63. Fields in the PROCU Report.....	149
64. Additional Fields in the PROCU Report.....	150
65. Fields in the RLRLRU Report.....	152
66. Fields in the VSAM RLS Activity report.....	156
67. Fields in the Disk Space Report.....	158
68. Fields in the Storage Space Report.....	159
69. Fields in the STOR Report.....	160
70. Additional Fields in the STOR Report.....	161
71. Fields in the STORC Report - System Information Section.....	163
72. Fields in the STORC Report - Job Information Section.....	164
73. Additional Fields in the STORC Report.....	165

74. Fields in the STORCR Report.....	167
75. Fields in the STORF Report.....	168
76. Fields in the STORM Report - System Summary Section.....	170
77. Fields in the STORM Report - Address Space Section.....	171
78. Fields in the STORR Report - Central Storage Summary Section.....	174
79. Fields in the STORR Report - Page/Swap Activity Section.....	174
80. Additional Fields in the STORR Report.....	176
81. Fields in the STORS Report.....	177
82. Additional Fields in the STORS Report.....	178
83. Fields in the SYSENQ Report.....	181
84. Fields in the SYSINFO Report.....	183
85. Additional Fields in the SYSINFO Report.....	185
86. Fields in the SYSRG Report.....	188
87. Fields in the SYSRTD Report.....	192
88. Fields in the SYSSUM Report.....	197
89. Fields in the SYSSUM Report - Continuation.....	200
90. Additional Fields in the SYSSUM Report.....	202
91. Fields in the SYSTREND Report.....	204
92. SYSTREND Report - Cursor-sensitive Control for Navigation.....	204
93. Fields in the SYSWKM Report.....	208
94. Fields in the Job Usage report.....	213
95. Additional fields in the Job Usage report.....	214
96. Cursor sensitivity on Job Usage report.....	215
97. Fields in the Speed Section of the WFEX Report.....	218
98. Fields in the Exceptions Section of the WFEX Report.....	218

99. Fields in the WFEX Action Panel.....	222
100. Fields in the WFEX Definition and Criteria Panel.....	224
101. Fields in the Graphic WFEX Report.....	231
102. Fields in the XCF Report.....	233
103. Fields in the XCF Group Data Report.....	234
104. Fields in the XCF Overview Report.....	235
105. Fields in the XCF Path statistics Report.....	236
106. Fields in the XCF Systems data Report.....	237
107. zFS File System report field descriptions.....	241
108. zFS File System report field descriptions - zFS File Systems Details panel.....	242
109. zFS Kernel report field descriptions.....	244
110. zFS Overview report field descriptions.....	245
111. Fields in the zFS Overview report - I/O details.....	247
112. Fields in the zFS Overview Report - User Cache Details.....	248
113. Fields in the zFS Overview Report - Vnode Cache Details.....	249
114. Fields in the zFS Overview Report - Metadata Cache/Backing Cache Details.....	250
115. Monitor II Display Session Areas.....	253
116. Fields in the ARD and ARDJ reports.....	257
117. Fields in the ASD and ASDJ Report.....	261
118. Fields in the ASRM and ASRMJ Report.....	266
119. Fields in the CHANNEL Report.....	268
120. Fields in the DEV and DEVV Report.....	273
121. Fields in the HFS Report.....	277
122. Fields in the HFS Report Options Panel.....	278
123. Fields in the ILOCK Report.....	280

124. Fields in the IOQUEUE Report.....	283
125. Fields in the LLI Report.....	287
126. Fields in the OPT Settings report.....	289
127. Fields in the PGSP Report.....	291
128. Fields in the SDS Report.....	292
129. Fields in the SENQ Report.....	296
130. Fields in the SENQR Report.....	298
131. Fields in the SPAG Report.....	300
132. Fields in the SRCS Report.....	302
133. Available formats for Postprocessor reports.....	307
134. Interval and Duration Reports.....	309
135. Monitor II Interval Reports.....	310
136. Fields in the Cache Subsystem Summary Report.....	314
137. Fields in the Cache Subsystem Activity Report - Header.....	315
138. Fields in the Cache Subsystem Activity Report - Subsystem Status.....	316
139. Fields in the Cache Subsystem Activity Report - Subsystem Overview.....	317
140. Fields in the Cache Subsystem Activity Report - Device Overview.....	319
141. Fields in the Cache Subsystem Activity Report - RAID Rank Activity.....	320
142. Fields in the Cache Subsystem Activity Report - Cache Device Status.....	321
143. Overview names in the Cache Subsystem Activity Report.....	321
144. Fields in the Coupling Facility Activity Report - Usage Summary.....	327
145. Fields in the Coupling Facility Activity Report - Structure Activity.....	334
146. Fields in the Coupling Facility Activity Report - Subchannel Activity.....	337
147. Fields in the Coupling Facility Activity Report - Subchannel Activity - Channel Path Details.....	338
148. Fields in the CF to CF Activity Section.....	339

149. Fields in the Coupling Facility Activity Report - CF to CF Activity - Channel Path Details.....	340
150. Overview names in the Coupling Facility Activity Report.....	340
151. Fields in the Channel Path Activity report.....	343
152. Overview names in the Channel Path Activity report.....	346
153. Fields in the CPU Activity Report.....	351
154. Overview names in the CPU Activity Report.....	356
155. Fields in the Partition Data Report.....	359
156. Overview names in the Partition Data Report.....	363
157. Fields in the LPAR Cluster Report.....	364
158. Overview names in the LPAR Cluster Report.....	364
159. Fields in the Group Capacity Report.....	365
160. Overview names in the Group Capacity Report.....	366
161. Fields in the Hardware Group Report.....	367
162. Overview names in the Hardware Group Report.....	367
163. Fields in the Power Consumption Report.....	368
164. Overview names in the Power Consumption Report.....	368
165. Fields in the CRYPTO Hardware Activity Report.....	371
166. Overview names in the CRYPTO Hardware Activity Report.....	375
167. Fields in the Device Activity Reports.....	381
168. Overview names in the DASD Activity Report.....	388
169. Fields in the Domino Server Summary Report.....	391
170. Fields in the Domino Server Details Report.....	391
171. Fields in the EADM Activity Report.....	394
172. Fields in the Enqueue Activity Report.....	397
173. Overview names in the Enqueue Activity Report.....	399

174. Common Fields in the Enterprise Disk Systems Report.....	401
175. Fields in the ESS Link Statistics.....	402
176. Fields in the ESS Synchronous I/O Link Statistics Report.....	403
177. Fields in the ESS Extent Pool Statistics.....	404
178. Fields in the ESS Rank Statistics.....	404
179. Overview names in the ESS Link Statistics section.....	405
180. Overview names in the ESS Extent Pool Statistics section.....	406
181. Overview names in the ESS Rank Statistics section.....	407
182. Overview names in the ESS Synchronous I/O Link Statistics section.....	407
183. Fields in the FICON Director Activity Report.....	408
184. Overview names in the FICON Director Activity Report.....	409
185. Fields in HFS Global Statistics Report.....	410
186. Fields in the HFS File System Statistics Report.....	411
187. Fields in the HTTP Server Summary Report.....	414
188. Fields in the HTTP Server Details Report.....	415
189. Fields in the I/O Queuing Activity Report.....	420
190. Overview names in the I/O Queuing Activity Report.....	427
191. Fields in the OMVS Kernel Activity Report.....	429
192. Fields in the Page Data Set Activity report.....	432
193. Overview names in the Page Data Set Activity report.....	433
194. Fields in the Paging Activity report - Central Storage Paging Rates.....	435
195. Overview names in the Paging Activity report - Central Storage Paging Rates.....	437
196. Fields in the Paging Activity report - Central Storage Movement and Request Rates.....	438
197. Overview names in the Paging Activity report - Central Storage Movement and Request Rates.....	438
198. Fields in the Paging Activity report - Frame and Slot Counts.....	439

199. Overview names in the Paging Activity report - Frame and Slot Counts.....	441
200. Fields in the Paging Activity report - Memory Objects and High Virtual Storage Frames.....	442
201. Overview names in the Paging Activity report - Memory Objects and High Virtual Storage Frames.....	444
202. Fields in PAGING Report - Dedicated memory.....	445
203. Fields in the PCIE Activity Report.....	450
204. Overview conditions in the PCIE Activity Report.....	452
205. Fields in the Serialization Delay Summary section.....	455
206. Fields in the Serialization Delay Details section - CMS Lock Details.....	457
207. Fields in the Serialization Delay Details section - CML and Local Lock Details.....	458
208. Fields in the Serialization Delay Details section - GRS Latch Details.....	459
209. Fields in the Serialization Delay Details section - GRS Enqueue Details.....	460
210. Fields in the Shared Device Activity Reports.....	462
211. Overview names in the Shared DASD Activity Report.....	467
212. Fields in the Trace Activity Report.....	468
213. Variables in the Trace Activity Report.....	469
214. Fields in the Virtual Storage Activity Report - Common Storage Summary.....	474
215. Overview names in the Virtual Storage Activity Report.....	475
216. Fields in the Virtual Storage Activity Report - Common Storage Detail Section.....	476
217. Fields in the Virtual Storage Activity Report - Private Area Summary.....	477
218. Fields in the Virtual Storage Activity Report - Private Storage Map.....	477
219. Fields in the Virtual Storage Activity Report - Bottom Half.....	478
220. Fields in the Virtual Storage Activity report - Private Area Detail section	479
221. Fields in the Virtual Storage Activity report - High Virtual Memory Usage section.....	479
222. Fields in the Workload Activity Report.....	490
223. Fields in the WLMGL Report - POLICY.....	503

224. Exception and Overview names in the Workload Activity Report.....	506
225. Fields in the XCF Activity Report - Usage by System.....	510
226. Fields in the XCF Activity Report - Usage by Member.....	512
227. Fields in the XCF Activity Report - XCF Path Statistics.....	512
228. Fields in the Exception Report.....	515
229. Fields in the Overview Report.....	516
230. Overview Header Record - Prefix Section.....	519
231. Overview Header Record - Header Section.....	519
232. Overview Header Record - Report Column Names Section.....	520
233. Overview Data Record - Prefix Section.....	521
234. Overview Data Record - Data Section.....	521
235. Overview Data Record - Report Data Section.....	521
236. Fields in the Summary Report.....	522

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 Resource Measurement Facility User's Guide*.

About special purpose processors:

Throughout this document, zIIP refers to IBM Z® Integrated Information Processors. zAAP refers to IBM Z Application Assist Processors.

Who should use this document

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

How this document is organized

This document contains the following chapters:

Chapter 1, “Introducing RMF data gathering and reporting,” on page 1

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

Chapter 2, “Interactive performance analysis with Monitor III,” on page 7

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

Chapter 3, “Snapshot reporting with Monitor II,” on page 251

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

Chapter 4, “Real-time reporting with Monitor I,” on page 305

This chapter gives you a table of reports you can request when using a Monitor I session. Since all Monitor I reports are also Postprocessor reports, the detailed description of these reports is located in Chapter 5, “Long-term overview reporting with the Postprocessor,” on page 307

Chapter 5, “Long-term overview reporting with the Postprocessor,” on page 307

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

z/OS RMF library

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

- *z/OS Resource Measurement Facility Programmer's Guide*
- *z/OS Resource Measurement Facility User's Guide*
- *z/OS Resource Measurement Facility Report Analysis*

- *z/OS Resource Measurement Facility Messages and Codes* (includes z/OS Data Gatherer messages)

z/OS Data Gatherer library

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

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

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

z/OS information

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

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

To find the complete z/OS library, go to [IBM Documentation \(www.ibm.com/docs/en/zos\)](http://www.ibm.com/docs/en/zos).

How to provide feedback to IBM

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

Summary of changes

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

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

Summary of changes for z/OS Resource Measurement Facility Report Analysis 3.2

New

The following content is new.

April 2026

- As part of APAR OA69228 (www.ibm.com/support/pages/apar/OA69228) implementation, the Service Class Period is included in the following reports:
 - [DELAY report](#)
 - [PROC report](#)
 - [STOR report](#)
 - [STORC report](#)
- The “CPC - CPC Capacity Report” on page 56 is enhanced to display Hardware Group Cap utilization as part of [APAR OA69228](http://www.ibm.com/support/pages/apar/OA69228) (www.ibm.com/support/pages/apar/OA69228).
- Updated Fields in the Virtual Storage Activity Report - Common Storage Summary and Fields in the STORC Report - System Information Section to clarify how SQA-to-CSA conversion values are accounted for differently in Monitor III and Postprocessor reports.

December 2025

As part of [APAR OA68173](http://www.ibm.com/support/pages/apar/OA68173) (www.ibm.com/support/pages/apar/OA68173), added the following support for NETH PCIE card:

- A new device Network Express Hybrid (NETH) has been added for the PCIE Activity Report. See “[PCIE - PCIE Activity Report](#)” on page 137 and “[Contents of the report](#)” on page 138.
- A new RMF NETH activity panel and fields for the NETH Activity pop-up panels are added. See “[Cursor sensitive control](#)” on page 139.

September 2025

All features and enhancements introduced in RMF for z/OS 3.1 are also available in 3.2.

Changed

The following content is changed.

January 2026

- The WLMGL - Service Class Report sample is updated with the missing POWER USAGE subheader. See “[Service Class report](#)” on page 485.
- The WLMGL - Report Class Report sample is updated with the missing POWER USAGE subheader. See “[Report Class report](#)” on page 487.

- The PROCESSOR, I/O, and MEMORY field descriptions for WLMGL - Workload Activity report have been updated as per technical accuracy terms. See [“Field descriptions of all reports”](#) on page 490.

Summary of changes for z/OS Resource Measurement Facility Report Analysis 3.1

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

New

The following content is new.

September 2025

Support is added for the OpenMetrics exposition format for Monitor III, enabling integration with third-party monitoring tools.

June 2025

The following reports have been enhanced to include the information about Power Consumption:

- [CPU report](#)
- [WLMGL Service Class report](#)
- [WLMGL Report Class report](#)

March 2025

A new field (Report class) has been added to the following reports:

- [PROC - Processor Delays Report](#)
- [PROCU - Processor Usage Report](#)
- [DELAY Report](#)

December 2024

The Enterprise Disk Systems Statistics report has been improved to feature a COMPRESSION section, which now includes additional columns for more comprehensive data analysis. See [ESS Extent Pool Statistics](#).

June 2024

The **Paging Activity** report has been updated to include the information about Dedicated Memory. See [“Dedicated Memory”](#) on page 445.

March 2024

- WLM I/O Management/Velocity Migration
 - The Workload Activity (WLMGL) report has been enhanced to make WLM I/O Management migration aid bidirectional. Based on these enhancements, the description of the **VELOCITY MIGRATION** field in the Workload Activity Report has been updated. See [“Field descriptions of all reports”](#) on page 490.

- Implicitly CPU critical and Batch AI

- New fields in WLMGL

The latest update in the Service Class and Service Class Period reports includes a new field called **BATCH INITIATOR AI-INFUSED**. This field indicates whether the WLM-managed batch initiator for the given service class is AI-infused. It also introduces new values for the **CRITICAL** field to indicate Implicit CPU protection. See [“Service Class report”](#) on page 485 and [“Service Class Period report”](#) on page 481.

- New fields and values in the Monitor III reports

The Monitor III reports introduce new fields and values. These new fields and values indicate Implicit CPU protection and Batch AI for Monitor III reports. See [DELAY report](#), [System Information \(SYSINFO\) report](#), and [Sysplex Summary \(SYSSUM\) report](#).

September 2023

- The following new reports have been added:
 - [“XCFGROUP - XCF Group Data” on page 233](#)
 - [“XCFOVW - XCF Overview” on page 234](#)
 - [“XCFPATH - XCF Path statistics” on page 235](#)
 - [“XCFSYS - XCF Systems data” on page 237](#)
- The new fields have been added to the PROCU Report. See [“Monitor III Utility fields” on page 150](#).

Changed

The following content is changed.

April 2025

The Monitor II report does not support detailed reporting for certain Channel Measurement Groups (CMGs).

October 2024

The following reports have been updated to include the missing Monitor III Utility fields:

- [PROC - Processor Delays Report](#)
- [PROCU - Processor Usage Report](#)
- [DELAY Report](#)

September 2023

The [TRANS /SEC](#) field description has been updated to include the change in cursor-sensitive control.

Changed the references of the following words with z/VM or alternate VM due to the addition of the Wazi and zPDT virtual environments:

- z/VM
- z/VM directory entry
- z/VM LPAR
- z/VM partition

Chapter 1. Introducing RMF data gathering and reporting

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

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

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.

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 *z/OS Resource Measurement Facility User's Guide*.

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

Monitoring on the workstation

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

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

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

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

What you can gather and report

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

Activity monitoring

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

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

Table 1. Monitored activities and SMF record types

Gathering				Activity	Reporting			
Short-term Mon III		Snapshot Mon II	Long-term Mon I		Interactive Mon III	Snapshot Mon II	Real-time Mon I	Long-term Post-processor
SMF	VSAM	SMF	SMF					
	*	79.1/2/5		Address space	*	*		*
	*		74.5	Cache	*			*
	*	79.12	73	Channel path	*	*	*	*
74.4	*			Coupling facility	*			*
	*		70.2	Cryptographic hardware	*		*	*
			70.3	CPU diagnostics				
	*	79.9	74.1	Device	*	*	*	*
74.10	*			EADM Activity	*			*
	*			Enclave	*			
	*	79.7	77	Enqueue	*	*	*	*
			74.8	Enterprise Storage Server (ESS)				*
			74.7	FICON director				*
		79.15		IRLM long locks		*		
	*	79.14	78.3	I/O queuing	*	*	*	*
		79.11	75	Page data set		*	*	*
		79.4	71	Paging		*	*	*
74.9	*			PCIE Activity	*			*
	*	79.3	70.1	Processor	*	*	*	*
		79.6		Reserve		*		*
72.5				Serialization Delay				*
72.4	*	79.3		Storage	*	*		*
			76	System counters			*	*
74.3/6	*			UNIX	*	*		*
	*		78.2	Virtual storage	*		*	*

Table 1. Monitored activities and SMF record types (continued)								
Gathering				Activity	Reporting			
Short-term Mon III		Snapshot Mon II	Long-term Mon I		Interactive Mon III	Snapshot Mon II	Real-time Mon I	Long-term Post-processor
SMF	VSAM	SMF	SMF					
	*		72.3	Workload Service classes and report classes	*			*
74.2	*			XCF	*			*
	*			zFS	*			

Delay monitoring

Delayed address spaces and groups

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

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

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

Delay reasons for address spaces

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

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

Long-term performance analysis with RMF XP

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

Subtype 1-12

AIX on System p performance data

Subtype 20-31

Linux on System x performance data

Subtype 40-53

Linux on System z performance data

For information on the metric categories provided in the subtypes and how to request the collection of SMF record type 104 from the systems of all or selected supported platforms, see [Cross platform monitoring with RMF XP](#) in *z/OS Resource Measurement Facility User's Guide*.

Reporting of other SMF data

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

WebServer performance reporting

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

Lotus Domino support

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

OpenMetrics with Monitor III

Monitor III metrics can be exposed to third-party tools using the OpenMetrics exposition format.

The OpenMetrics format is supported by various solutions, including Prometheus, VictoriaMetrics, InfluxDB via Telegraf, and many others. Each Monitor III metric can be skipped or exposed by default or a custom name.

For more information on OpenMetrics and its implementation, refer to the OpenMetrics Endpoint section in *z/OS Resource Measurement Facility User's Guide*.

Chapter 2. Interactive performance analysis with Monitor III

This information unit:

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

Using Monitor III reports

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

- [“System activities measured” on page 7](#)
- [“Where to start” on page 7](#)
- [“Using cursor-sensitive control” on page 10](#)
- [“Common Monitor III report measurements” on page 10](#)
- [“Monitor III MINTIME and range” on page 14](#)
- [“Monitor III report options” on page 15](#)
- [“Monitor III sysplex support in different time zones” on page 16](#)

System activities measured

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

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

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

See chapter [“DELAY - Delay Report” on page 66](#) for more details.

Where to start

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

It is necessary to:

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

Monitor III addresses these needs by:

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

Suggested sequence of reports

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

Start the report in GO mode and let it run.

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

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

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

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

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

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

3. The scrollable section in the SYSRTD report is the link from the sysplex level to the single system. From these lines it is possible to "zoom" into any of the listed systems. Placing the cursor on the system-ID in one of the rows and pressing ENTER, leads to the SYSINFO report of that system, thus allowing further analysis based on the data shown there. Placing the cursor on a specific data column in one of the rows of the scrollable area leads to a specific report of that system that provides additional information related to the selected column.
4. Finally, when the single-system level is reached, navigation among those reports is possible as described in step "5" on page 8.
5. In the workflow/Exceptions (WFEX) report, you can identify jobs and resources with low workflow values or jobs that have met exceptional conditions. For example, you can check the Reason field to identify the user or the possible cause of delay. Once you recognize a user or a resource with a potential problem, you can analyze the situation using cursor-sensitive control. "Using cursor-sensitive control" on page 10 describes how to invoke reports using this method.

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

In case of a delay due to devices (DEV) or enqueued resources (ENQ), you can use cursor-sensitive control to further investigate a problem by looking at the resource-oriented device report (DEVVR) and

the resource-oriented enqueue report (ENQR). For storage problems involving paging or swapping delays, you can use the resource-oriented storage delay report (STORR). Use either the job entry subsystem (JES), hierarchical storage manager (HSM), or cross-system coupling facility (XCF) delays report for a delay associated with SUBS. For OPER delay, use cursor-sensitive control to see the appropriate Job Delay report.




Figure 1. Suggested Sequence for Using Monitor III Reports

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

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

[Figure 1 on page 9](#) shows a suggested sequence for using Monitor III reports to resolve potential problems.

Getting information about data to be reported

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

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

This leads to the Data Index screen.

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

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

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

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

Using cursor-sensitive control

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

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

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

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

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

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

Cursor-sensitive control is active on:

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

Common Monitor III report measurements

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

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

- Execution velocity

Using samples

PROC

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

DEV

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

Delay samples

PROC

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

DEV

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

ENQ

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

HSM

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

JES

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

OPR

The number of address spaces found waiting for operator interventions.

STR

The number of address spaces found waiting for storage operations.

XCF

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

Address space workflow (%)

The workflow of an address space represents how a job uses system resources and the speed at which the job moves through the system in relation to the maximum average speed at which the job could move through the system. The speed at which the system performs the work of one job depends on the simultaneous work requested by other jobs.

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

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

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

Single Address Space

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

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

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

Example

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

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

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

Group of Address Spaces

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

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

Resource workflow (%)

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

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

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

Resource

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

Address space using (%)

Jobs getting service from hardware resources (PROC or DEV) are *using* these resources. The use of a certain resource by an address space can vary from 0% to 100%, where 0% indicates no use of the

resource during the report interval and 100% indicates that the address space was found using the resource in every sample during that period. If you use the default range of 100 seconds, 1% of using is equal to 1 second of using to the user.

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

Single Address Space

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

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

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

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

Group of Address Spaces

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

Address space delay (%)

The delay of an address space represents a job that needs one or more resources but that must wait because it is contending for the resource(s) with other users in the system. The delay of an address space for a specific resource or for all resources can vary from 0% to 100%. A delay of 0% indicates no delay during the report interval, while a delay of 100% represents a job that was found delayed at every sample during that period. Delay is a percent of Time during the period; with the default Range of 100 seconds, 1% delay is equal to one second of delay to the user.

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

Single Address Space

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

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

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

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

Group of Address Spaces

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

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

Execution velocity

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

The calculation of the execution velocity is:

Execution Velocity

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

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

Monitor III MINTIME and range

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

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

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

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

nnnnS where nnnn represents a number from 0 to 9999

nnnM where nnn represents a number from 1 to 166

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

How the data gathered affects the data reported

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

Example

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

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

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

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

Requested Report

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

Range: 100 Sec

```

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

Presented Report

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

Range: 200 Sec

```

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

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

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

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

Shortened intervals

The following events can cause a shortened report interval:

A policy switch

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

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

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

Example

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

```

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

Monitor III report options

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

will stay valid until they will be changed explicitly. In addition, you can use the following commands in a Report Options panel:

CANCEL

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

RESET

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

Monitor III sysplex support in different time zones

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

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

Example

SYSA (local time is GMT+2)

SYSB (local time is GMT-1)

Monitor III reporter session started on SYSA

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

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

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

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

Note:

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

The reporter session

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

```

Selection ==>>      RMF Monitor III Primary Menu      z/OS V3R2 RMF

Enter selection number or command on selection line.

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

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

                    O OPTIONS      T TUTORIAL      X EXIT

                    5655-ZOS Copyright IBM Corp. 1986, 2023.
                    Licensed Materials - Property of IBM

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

```

Figure 2. Monitor III Primary Menu

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

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

Selecting a report

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

- Enter the report name or its abbreviation on any selection or command line.

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

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

Example

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

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

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

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

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

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

```
Command ===> ROPTIONS
```

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

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

Monitor III report commands - Overview

Table 2 on page 19 lists all report commands with their parameters and abbreviations. The "How to request this report" section in the *z/OS Resource Measurement Facility Report Analysis* for each report shows an example of the command and parameters.

You can enter the commands on any command line.

The **Parameters** column in Table 2 on page 19 indicates what parameters, if any, you can specify on the respective commands:

cfname

A coupling facility name

job_class

One of the following names of a job class:

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

Note:

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

dsname

A data set name

jobname

A job name

period

A service or report class period

resource

A resource name

resource_group

The name of a resource group or tenant resource group

service_class

A service class name

s/r-class

A service or report class name

ssid

A cache subsystem identifier

sstype

The name of a subsystem that schedules enclaves

storage_class

A storage class name

volser

A serial number of a volume

wlm

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

Table 2. Report Commands			
Command	Parameters	Displays	Abbreviation
CACHDET	ssid	Cache detail report	CAD
CACHSUM		Cache summary report	CAS
CFACT	cfname	coupling facility activity report	CA
CFOVER	cfname	coupling facility overview report	CO
CFSYS	cfname	coupling facility system report	CS
CHANNEL		Channel path activity report	CHAN, CH
CPC		CPC capacity report	
CRYACC		Crypto accelerator activity	CRA
CRYOVW		Crypto hardware overview	CRO
CRYPKC		Crypto PKCS11 coprocessor activity	CRP
DELAY	job_class, service_class	Delays report for all jobs or specified job groups	DEL, DLY, DL
DELAYJ	jobname	Job report variation for specified job reflecting primary delay reason	DLJ, DJ, DELJ, DLYJ, JOB, JO
DEV	job_class, service_class	Device delays report for all jobs or specified job groups	DD, DVD
DEVJ	jobname	Device delays variation of job report for specified jobname	DDJ, DVJ
DEVR	volser	Device delays report for all or specified resources	DR, DVR
DSINDEX		Data index information	DS, DI
DSND	dsname	Data set delays report for all or specified data sets	DSN
DSNJ	jobname	Data set delays - Job report for specified jobname	DSJ
DSNV	volser	Data set delays - Volume report for specified volume	DSV
EADM		EADM activity report	SCM
ENCLAVE	sstype	Enclave activity report	ENCL

Table 2. Report Commands (continued)

Command	Parameters	Displays	Abbreviation
ENQ	job_class, service_class	Enqueue delays report for all jobs or specified job groups	ED
ENQJ	jobname	Enqueue delays variation of job report for specified jobname	EJ
ENQR	resource	Enqueue delays for all or specified resources	ER
GROUP	s/r-class, period	Group response time breakdown	GP, GRP, GD, RT, GRT
HSM	job_class, service_class	HSM delays report for all jobs or specified job groups	HD
HSMJ	jobname	HSM delays variation of job report for specified jobname	HJ
IOQUEUE		I/O queuing activity report	IOQ, IQ
JES	job_class, service_class	JES delays report for all jobs or specified job groups	JD
JESJ	jobname	JES delays variation of job report for specified jobname	JJ
JOB	jobname	Job report variation for specified job reflecting primary delay reason	JO, DELAYJ, DLYJ, DELJ, DLJ, DJ
LOCKSP	HELD SPIN <u>BOTH</u>	Spin Lock Report about held spin locks and/or address spaces spinning due to a request for a spin lock	LSP
LOCKSU	LOCAL GLOBAL <u>BOTH</u>	Suspend Lock Report about local and/or global suspend locks	LSU
MNTJ	jobname	Operator delays variation for mount request of job report for specified jobname	MTJ
MSGJ	jobname	Operator delays variation for message request of job report for specified jobname	MSJ
OPD		OMVS process data	
PCIE		Statistics and performance measurements about PCI Express-based functions	PCI
PROC	job_class, service_class	Processor delays report for all jobs or specified job groups	PD
PRO CJ	jobname	Processor delays variation of job report for specified job	PJ
PROCU	job_class, service_class	Processor usage of a job per processor type (standard or special purpose processors)	PU
QSCJ	jobname	Operator delays variation for quiesce command of job report for specified jobname	QJ
RLSDS	dsname	VSAM RLS activity by data set	RLD
RLSLRU		VSAM LRU overview	RLL
RLSSC	storage_class	VSAM RLS activity by storage class	RLS
SPACED		Disk space report	SPD
SPACEG		Storage space report	SPG
STOR	job_class, service_class	Storage delays report for all jobs or specified job group	SD
STORC	job_class, service_class	Common storage report	SC
STORCR		Common storage remaining at end of job report	SCR
STORF	job_class, service_class	Detailed information on frame counts for all jobs or specified job group	SF

Table 2. Report Commands (continued)			
Command	Parameters	Displays	Abbreviation
STORM	job_class, service_class	Detailed information about the use of memory objects within the system	SM
STORJ	jobname	Storage delays variation of job report for specified job	SJ
STORR		Storage space and paging activity report for all system volumes	SR
STORS	wlm	Summarized storage information by workload group, service or report class	SS
SYSENG		Sysplex enqueue delays report	ES
SYSINFO	wlm	System information, total and by user groups	SY, SYS, SI
SYSRG	resource_group	Resource Group Activity report	SRG
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
USAGE	job_class, service_class	Job usage report	USG
WFEX		Workflow/exceptions screen	WE, WF
XCF	job_class, service_class	Cross-system coupling facility delays report	XD
XCFJ	jobname	XCF delays variation of the job report for specified jobname	XJ
ZFSFS		zFS file system	ZFF
ZFSKN		zFS file system kernel	ZFK
ZFSOVW		zFS file system overview	ZFO

Table 3 on page 21 contains commands for the examples of user-written reports that were delivered with RMF.

Table 3. User-Written Report Commands			
Command	Parameters	Displays	Abbreviation
DEVN		Device activity	DA
DEVT		Device trend	DT
DSD		Detailed storage delays	
SYSTREND		System trend	ST

Header for single-system reports

Figure 3 on page 21 shows the common header for single-system Monitor III reports.

```

RMF 3.1  TITLE                               Line 1 of 30
Command ==>                               Scroll ==> HALF
Samples: nnn    System: syst Date: mm/dd/yy Time: hh.mm.ss Range: nnn  Sec

```

Figure 3. Header of Monitor III Single-System Reports

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

Table 4. Monitor III Heading Information	
Heading	Definition
Report title	The type of measurement data
Date/Time	The starting date and time for the first set of samples included in the report
Range	The length of time (in seconds) during which samples were gathered, starting with the time specified in the Time field

Header for sysplex reports

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

1. The sysplex reports show the sysplex name, whereas the single-system reports show the SMF system identification. In Figure 4 on page 22, **SYSPLEXN** is the eight character sysplex name.
2. For some sysplex reports, for example, SYSSUM or SYSWKM, the RMF sample count is replaced by the number of WLM samples. This value is an average of the MINTIMES from the different systems, that contributed to the report.
3. The number of systems participating in the report is shown and indicates whether the complete sysplex is contributing to the report or not. To find out which specific system is not included, check the HELP-Screen of the message *Not all systems included in the report*. Furthermore, the Data Index may be useful.

The sysplex report header is shown in Figure 4 on page 22.

```

Command ==>          RMF 3.1  TITLE - SYSPLEXN                      Line 1 of 30
                                                                Scroll ==> HALF
WLM Samples:   nnn   Systems: n Date: mm/dd/yy Time: hh.mm.ss Range: nnn   Sec

```

Figure 4. Header of Monitor III Sysplex Reports

The Sysplex Report Selection Menu

```

RMF Sysplex Report Selection Menu
Selection ==>

Enter selection number or command for desired report.

Sysplex Reports
 1 SYSSUM   Sysplex performance summary                (SUM)
1A SYSRG   Resource Group activity                     (SRG)
 2 SYSRTD   Response time distribution                 (RTD)
 3 SYSWKM   Work Manager delays                       (WKM)
 4 SYSENG   Sysplex-wide Enqueue delays               (ES)
 5 CFOVER   Coupling Facility overview                 (CO)
 6 CFSYS    Coupling Facility systems                  (CS)
 7 CFACT    Coupling Facility activity                 (CA)
 8 CACHSUM  Cache summary                             (CAS)
 9 CACHDET  Cache detail                              (CAD)
10 RLSSC   VSAM RLS activity by storage class          (RLS)
11 RLSDS   VSAM RLS activity by data set               (RLD)
12 RLSLRU  VSAM LRU overview                          (RLL)
13 ZFSOVW  zFS Overview                               (ZFO)
14 ZFSFS   zFS File System                            (ZFF)
15 ZFSKN   zFS Kernel                                 (ZFK)
16 CRYOVW  Crypto hardware overview                   (CRO)
17 CRYACC  Crypto accelerator activity                 (CRA)
18 CRYPKC  Crypto PKCS11 coprocessor activity          (CRP)

Data Index
  D DSINDEX Data index                                (DI)

```

Figure 5. Monitor III Sysplex Report Selection Menu

You can navigate to the **Sysplex Report Selection Menu** by selecting **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 25.

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

The Overview Report Selection Menu

```

RMF Overview Report Selection Menu
Selection ==>

Enter selection number or command for desired report.

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

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

```

Figure 6. Monitor III Overview Report Selection Menu

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

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

The Job Report Selection Menu

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

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

```

RMF Job Report Selection Menu
Selection ==>

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

Jobname ==> _____

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

```

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

Figure 7. Monitor III Job Report Selection Menu

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

The Resource Report Selection Menu

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

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

```

RMF Resource Report Selection Menu
Selection ==>

Enter selection number or command for desired report.

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

```

Figure 8. Monitor III Resource Report Selection Menu

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

The Subsystem Report Selection Menu

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

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

```

RMF Subsystem Report Selection Menu
Selection ==>

Enter selection number or command for desired subsystem report.

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

```

Figure 9. Monitor III Subsystem Report Selection Menu

The User Report Selection Menu

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

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

```

RMF User-written Report Selection Menu
Selection ==>

Enter selection number or command for desired report.

2 DSD          Detailed Storage Delays

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

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

```

Figure 10. Monitor III User-written Report Selection Menu

The Data Index

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

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

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

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

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

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

How to request the Data Index

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

```
DI
```

Contents of the Data Index

Read the following information about the Data Index:

- [“If you are using active Monitor III gatherers in the sysplex” on page 26](#)
- [“If you are using preallocated data sets” on page 27](#)
- [“Condensed information on the Data Index” on page 28](#)
- [“Data Index — field descriptions” on page 29](#)
- [“Information and error messages on the Data Index” on page 29](#)
- [“Cursor-sensitive control” on page 30](#)
- [“Data Index options” on page 31](#)

If you are using active Monitor III gatherers in the sysplex

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

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

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

Figure 11 on page 27 shows a sample Data Index that is using data from active Monitor III data gatherers in a sysplex.

```

RMF 3.1  Data Index - RMFPLEX1                               Line 1 of 22
Command ==>                                                    Scroll ==> HALF

Samples: 118      System: MVS2  Date: 04/25/2023  Time: 10.12.00  Range: 120  Sec

----Begin/End----
System --Date-- --Time-- -DDNAME- -----Data Set Name-----
MVS1  04/25/2023 10.03.20
      10.12.00      * * *      In-storage buffer      * * *
MVS1      SYS00002 RMF.MONITOR3.DATASET1.MVS3
      * * *      Data from system MVS3      * * *
MVS1      SYS00001 RMF.MONITOR3.DATASET2.MVS3
      * * *      Data from system MVS3      * * *

MVS2  04/25/2023 09.11.00 SYS00002 RMF.MONITOR3.DATASET1.MVS2
      09.14.00
MVS2  04/25/2023 10.03.00 SYS00003 RMF.MONITOR3.DATASET2.MVS2
      10.12.00      * * *      Currently active      * * *
MVS2  04/25/2023 10.03.00
      10.12.00      * * *      In-storage buffer      * * *

MVS3  04/25/2023 09.11.00 SYS00002 RMF.MONITOR3.DATASET1.MVS3
      09.14.00
MVS3  04/25/2023 10.03.00 SYS00003 RMF.MONITOR3.DATASET2.MVS3
      10.12.00      * * *      Currently active      * * *
MVS3  04/25/2023 10.03.00
      10.12.00      * * *      In-storage buffer      * * *
TEST
      * * *      No response      * * *

```

Figure 11. Data Index

If you are using preallocated data sets

If a local session has had data sets preallocated to it before the RMF reporter is started, the Data Index will only display data from those data sets. This is independent of the active Monitor III gatherers in the sysplex. For information about how to preallocate data sets to the local session, see *z/OS Resource Measurement Facility 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 12 on page 28 shows a sample Data Index that is using data from preallocated data sets to a reporter session.

RMF 3.1 Data Index - RMFPLEX				Line 1 of 18	
Command ==>				Scroll ==> HALF	
Samples: 37	System: RMFB	Date: 04/25/2023	Time: 13.09.00	Range: 60	Sec
----Begin/End----					
System	--Date--	--Time--	-DDNAME-	-----Data Set Name-----	
RMFA	04/25/2023	04.17.00	RMFDS01	RMF.RMFA.RMFDS01	
		06.26.00			
RMFA	04/25/2023	04.20.00	RMFDS00	RMF.RMFA.RMFDS00	
		06.28.00			
RMFA	04/25/2023	11.28.00	RMFDS02	RMF.RMFA.RMFDS02	
		13.07.00			
RMFB	04/25/2023	04.16.00	RMFDS04	RMF.RMFB.RMFDS01	
		06.24.00			
RMFB	04/25/2023	04.26.00	RMFDS05	RMF.RMFB.RMFDS00	
		06.36.00			
RMFB	04/25/2023	11.58.00	RMFDS03	RMF.RMFB.RMFDS02	
		13.10.00			
RMFC	04/25/2023	04.17.00	RMFDS07	RMF.RMFC.RMFDS01	
		06.26.00			
RMFC	04/25/2023	04.27.00	RMFDS06	RMF.RMFC.RMFDS00	
		06.37.00			
RMFC			RMFDS07	RMF.RMFC.RMFDS02	
				* * *	
				Empty	
				* * *	

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

Condensed information on the Data Index

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

Figure 13 on page 28 and Figure 14 on page 28 show what the data displayed in Figure 11 on page 27 and Figure 12 on page 28, respectively, look like if the DDNAMES/DDNAMES option is used to condense the information displayed.

RMF 3.1 Data Index - RMFPLEX1				Line 1 of 4	
Command ==>				Scroll ==> HALF	
Samples: 118	System: RMFE	Date: 04/25/2023	Time: 10.12.00	Range: 120	Sec
-----Begin-----		-----End-----			
System	--Date--	--Time--	--Date--	--Time--	
MVS1	04/25/2023	10.03.20	04/25/2023	10.12.00	
MVS2	04/25/2023	09.11.00	04/25/2023	10.12.00	
MVS3	04/25/2023	09.11.00	04/25/2023	10.12.00	
TEST			* * *	No response	* * *

Figure 13. Data Index - Condensed Version

RMF 3.1 Data Index - RMFPLEX				Line 1 of 3							
Command ==>				Scroll ==> HALF							
Samples: 37		System: RMFB		Date: 04/25/2023		Time: 13.09.00		Range: 60		Sec	
-----Begin-----				-----End-----							
System		--Date-- --Time--		--Date-- --Time--							
RMFA	04/25/2023	04.17.00		04/25/2023	13.07.00						
RMFB	04/25/2023	04.16.00		04/25/2023	13.10.00						
RMFC	04/25/2023	04.17.00		04/25/2023	06.37.00						

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

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

Data Index – field descriptions

For a description of the report header area of the Data Index, refer to [“Header for single-system reports” on page 21](#), and [“Header for sysplex reports” on page 22](#), where the various header fields are described in more detail.

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

Table 5. Fields in the Data Index	
Field Heading	Meaning
System	The four-character SMF system identifier of the Monitor III gatherer that collected the data. If the identification of the system could not be determined, the field contains '????'.
Begin/End Date Time	These are the beginning and ending dates/times for the data in the usable and not-empty data sets or the in-storage buffers. If the beginning and ending dates are the same, RMF will only display the beginning date. If there is a problem with the data, the dates and times are left blank, and a message is shown.
DDNAME	For a Data Index using active gatherer's data, the DDNAME is the system generated DD name for the data set that has been dynamically allocated for the data gatherer's session. For a Data Index using preallocated data sets, this is the name that was specified in RMFDSxx on the ALLOCATE command issued before the reporter session was started. If there is a problem with the data, the field remains blank. Also, note that the field remains blank if this line is representing the gatherer's in-storage buffer.
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 29 .) Note that the first line is blank when the data represents a gatherer's in-storage buffer.
Begin Date Time	The begin date and time for which data is available on the respective system. This field is blank if there is no data to be represented for the system.
End Date Time	The end date and time for which data is available on the respective system. This field contains a comment about the status of the data if a problem was encountered.

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

Information and error messages on the Data Index

The following messages can be shown in special cases:

*** Currently active ***

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

*** In-storage buffer***

The local storage buffer entry of the Monitor III data gatherer.

*** Empty ***

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

*** No data available ***

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

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

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

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

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

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

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

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

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

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

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

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

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

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

***** Invalid CFSIZE *****

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

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

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

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

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

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

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

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

Dynamic allocation error (the data set is unusable).

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

Data set unallocated (the data set is unusable).

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

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

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

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

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

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

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

Cursor-sensitive control

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

Data Index options

```

RMF Data Index Options

Command ==>

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

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

```

Figure 15. Data Index Options Panel

The Data Index has two options:

- The *DDNAMES/DSNAMES* option allows you to determine the amount of information that will be displayed. There are two valid values for the *DDNAMES/DSNAMES* option:

YES

This value gives you a more detailed Data Index. It contains, at a data-set level, the data that is used on other RMF reports. You can see if there are any specific problems with the available data. See [Figure 11 on page 27](#) and [Figure 12 on page 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 13 on page 28](#) and [Figure 14 on page 28](#) for example screens.

- The *Sort Order* is available on the Data Index Options panel, so that it can be accessed with the *ROPTIONS* command from the Data Index screen. If more than one row with the same system identification exist, the usable data sets that are not empty are listed first, then the empty data sets, and finally the unusable data sets. The usable data sets that are not empty are sorted according to the end time of the stored data.

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

ASCEND

This value causes the entries in the Data Index to be sorted with the oldest data at the top of the individual system lists. See [Figure 11 on page 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 16 on page 31](#) for an example.

```

RMF 3.1  Data Index - RMFPLEX1                               Line 1 of 14
Command ==>                                                Scroll ==> HALF

Samples: 100      System: MVS1  Date: 04/25/2023  Time: 10.31.40  Range: 100  Sec

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

MVS1  04/25/2023 09.55.00
      10.31.40          * * *   In-storage buffer  * * *
MVS1  04/25/2023 05.53.20  SYS00003  SYS3.RMF.DS02
      10.11.40          * * *   Currently active   * * *
MVS1  04/25/2023 01.31.40  SYS00002  SYS3.RMF.DS01
      05.53.20

MVS2  04/25/2023 09.55.00
      10.31.40          * * *   In-storage buffer  * * *
MVS2  04/25/2023 01.31.40  SYS00003  SYS3.RMF.DS01
      10.11.40          * * *   Currently active   * * *

TEST                                     * * *   No response   * * *

```

Figure 16. Data Index - Sort Order Descend

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

CACHDET - Cache Detail Report

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

How to request this report

To request the CACHDET report, select **5** on the Primary Menu, and then select **9** on the Sysplex Report menu (shown in Figure 5 on page 22), or enter the following command:

```
CACHDET [subsystem_id]
```

Contents of the report

RMF 3.1 Cache Detail - SYSPLEX Line 1 of 17												
Samples: 60		Systems: 3		Date: 04/25/2023		Time: 13.35.00		Range: 60		Sec		
				CDate: 04/25/2023		CTime: 13.34.57		CRange: 59		Sec		
Volume /Num	SSID	I/O %	I/O Rate	Hit %	Cache Read	Hit DFW	Rate CFW	- DASD Total	I/O Stage	Seq Rate	Async Rate	
*ALL		100	208.6	100	208.2	0.4	0.0	0.0	0.0	0.0	0.0	
MVSTG3 041AC	4100	62.3	130.0	100	130.0	0.0	0.0	0.0	0.0	0.0	0.0	
MVSLIB 04100	4100	16.8	35.0	100	34.8	0.2	0.0	0.0	0.0	0.0	0.0	
MVSTGT 041AA	4100	5.9	12.3	100	12.3	0.0	0.0	0.0	0.0	0.0	0.0	
MVSOI4 041A4	4100	3.8	8.0	100	8.0	0.0	0.0	0.0	0.0	0.0	0.0	
MVSOI8 041A8	4100	2.5	5.1	100	5.1	0.0	0.0	0.0	0.0	0.0	0.0	
MVSOI2 041A2	4100	1.8	3.7	100	3.7	0.0	0.0	0.0	0.0	0.0	0.0	
MVSCIC 04103	4100	1.5	3.1	100	3.1	0.0	0.0	0.0	0.0	0.0	0.0	
MVSOI9 041A9	4100	1.3	2.8	100	2.8	0.0	0.0	0.0	0.0	0.0	0.0	

Figure 17. CACHDET Report

There is no graphic version of this report available.

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

RMF 3.1		Cache Detail		- SYSPLEX		Line 1 of 251	
Samples: 60	Systems: 3	Date: 04/25/2023	Time: 13.55.00	Range: 60	Sec		
		CDate: 04/25/2023	CTime: 13.54.56	CRange: 60	Sec		
Vol	-----						
	RMF Cache Volume Detail						
COB	The following details are available for Volume A0CSM2 on SSID 4C00						
MVS	Press Enter to return to the Report panel.						
COB							
AOC	DFW: Active		Pinned: None				
COB							
MVS	-----	Read	-----	-----	Write	-----	Read
COH	Rate	Hit	Hit%	Rate	Fast	Hit	Hit%
SYS	Norm	91.5	91.5	100	76.4	76.4	76.4
KEY	Seq	82.8	82.8	100	29.6	29.6	29.6
AOC	CFW	0.0	0.0	0.0	0.0	0.0	0.0
KEY	Total	174.4	174.4	100	106.0	106.0	106.0
MVS							
SYS	-----	Misc	-----	---	CKD	----	- Record Caching -
AOC	Del Op NVS	:	0.0	Write:	3.6	Read Miss	:
SAT	Del Op Cache:	0.0	Hits	:	3.6	Write Prom:	22.7
KEY	DFW Inhibit	:	0.0				
AOC							
RMF							
CPO							
IRD							
MVS	-----						
COHSMS	04B15	4B00	9.5	12.5	100	12.1	0.4
KEYSM4	04CB4	4C00	2.3	11.5	100	9.2	2.3

Figure 18. CACHDET Report - Volume Details (pop-up panel)

RMF 3.1		Cache Detail		- SYSPLEX		Line 1 of 17	
Samples: 60	Systems: 3	Date: 04/25/2023	Time: 13.35.00	Range: 60	Sec		
		CDate: 04/25/2023	CTime: 13.34.57	CRange: 59	Sec		
Vol	-----						
	RMF Cache SSID Detail						
*AL	The following details are available for SSID 4100						
MVS	Press Enter to return to the Report panel.						
MVS							
MVS		-----	Read	-----	-----	Write	-----
MVS		Rate	Hit	Hit%	Rate	Fast	Hit
MVS	Norm	197.5	197.5	100	0.4	0.4	0.4
MVS	Seq	10.7	10.7	100	0.0	0.0	0.0
MVS	CFW	0.0	0.0	0.0	0.0	0.0	0.0
MVS	Total	208.2	208.2	100	0.4	0.4	0.4
MVS							
MVS	-----	Misc	-----	---	CKD	----	- Record Caching -
MVS	Del Op NVS	:	0.0	Write:	98.2	Read Miss	:
MVS	Del Op Cache:	0.0	Hits	:	98.2	Write Prom:	0.3
MVS	DFW Inhibit	:	0.0				
MVS							
MVS							
MVS							

Figure 19. RMF Cache SSID Detail (pop-up panel)

Field descriptions

Table 6. Fields in the CACHDET Report	
Field Heading	Meaning
Device reserve activity can cause a data gatherer interface to wait until a RESERVE has been released. This in turn can cause the cache interval to be much longer than a regular RMF interval.	
Therefore, CDate, CTime, and CRange are used to show the actual point in time to which the cache interval start is related, and the actual cache interval length. All rates shown in the report are based on CRange, and not on Range.	

Table 6. Fields in the CACHDET Report (continued)	
Field Heading	Meaning
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. This line appears only if the report shows data for one specific subsystem ID.
/Num	The five-digit hexadecimal device number that identifies a physical I/O device. The first digit represents the ID of the subchannel set to which the I/O device is physically configured. If a four-digit device number is displayed, the device statistics were collected on a remote system where no subchannel set information was available.
SSID	Subsystem identifier; a number assigned during the installation of the subsystem that uniquely identifies the storage subsystem.
I/O %	Percentage of I/O requests to this device or category, compared to the total number of I/O requests sent to the subsystem.
I/O Rate	Rate of I/O requests.
Hit %	Percentage of I/Os that were processed within the cache (cache hits) based on the total number of I/Os.
Cache Hit Rate - I/O rate of all cache hits.	
Read	Rate of SEARCH/READ requests that completed without accessing the DASD.
DFW	Rate of DFW requests.
CFW	Rate of WRITE and READ-AFTER-WRITE requests that are processed in cache.
DASD I/O - I/O rate of all requests that accessed DASD.	
Total	I/O rate of all requests that accessed DASD. This is the sum of Stage rates (see below) and other request rates (inhibit cache load, DFW BYPASS, CFW BYPASS, DFW INHIBIT).
Stage	Rate of normal or sequential I/O requests that accessed DASD.
Seq Rate	Rate of tracks that have been staged due to cache misses for sequential I/O requests.
Async Rate	Rate of tracks that have been destaged asynchronously.

Table 7. Fields in the CACHDET Report - Volume and SSID Details	
Field Heading	Meaning
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. Exists Pinned data exists for the reported device. Unknown pinned status The pinned status cannot be determined.

Table 7. Fields in the CACHDET Report - Volume and SSID Details (continued)	
Field Heading	Meaning
Cache I/O Request - The channel command DEFINE EXTENT specifies the way the cache will be used. There are three categories: Norm - Seq - CFW	
Norm	Cache will be managed by least-recently-used (LRU) algorithm for making cache space available.
Seq	Tracks following the track assigned in the current CCW chain are promoted. They will be transferred from DASD to cache in anticipation of a short-term requirement.
CFW	WRITE and READ-AFTER-WRITE requests are processed in cache. The data might not be written to DASD. Because CFW does not use the NVS, the application is responsible for restoring the data after a cache or subsystem failure.
Total	This is either the sum of I/O requests, the total I/O rate, or the average hit ratio for the three categories previously described.
Read - Cache I/O requests that searched or read data from DASD. This is the number of channel operations that had at least one SEARCH or READ command but no WRITE commands. It is counted for cache devices only.	
Rate	Rate of SEARCH/READ requests.
Hit	Rate of SEARCH/READ requests that completed without accessing the DASD.
Hit%	Percentage of SEARCH/READ requests that completed without accessing the DASD.
Write - Cache I/O requests that wrote data to DASD. This is the number of channel commands that had at least one WRITE command. It is counted for cache devices only.	
Rate	Rate of WRITE requests.
Fast	Rate of DASD/CACHE FAST WRITE requests.
Hit	Rate of DASD/CACHE FAST WRITE requests that completed without accessing the DASD (fast write hit).
Hit%	Percentage of DASD/CACHE FAST WRITE requests that completed without accessing the DASD (fast write hit).
Read %	Percentage of READ requests based on the sum of all READ and WRITE requests (excluding ICL and BYPASS).
Tracks	Rate of tracks transferred from DASD to cache.
Misc - Miscellaneous cache activity rates.	
Del Op NVS	DASD Fast Write operations delayed due to non-volatile storage space constraints.
Del Op Cache	Operations delayed due to cache space constraints.
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.
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 found in the cache, and access to DASD was not required.

Report options

```

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

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

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

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

```

Figure 20. CACHDET Report Options

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

CACHSUM - Cache Summary Report

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

If you feel that further analysis is required, you may continue with the Cache Detail report (see [“CACHDET - Cache Detail Report”](#) on page 32).

How to request this report

To request the CACHSUM report, select **S** on the Primary Menu, and then select **8** on the Sysplex Report menu (shown in [Figure 5](#) on page 22), or enter the following command:

```
CACHSUM
```

Contents of the report

```

RMF 3.1  Cache Summary  - SYSDPLEX  Line 1 of 23
Samples: 60  Systems: 3  Date: 04/25/2023  Time: 13.35.00  Range: 60  Sec
                  CDate: 04/25/2023  CTime: 13.34.57  CRange: 59  Sec

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

2180  2182  2107-951  124G  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
4000  405C  2107-981  246G  23.8  100  23.8  0.0  0.0  100  0.0  0.0
4100  4102  2107-981  246G  208.6  100  208.6  0.0  0.0  99.8  0.0  0.0
4300  43AC  2107-981  246G  13.7  100  13.7  0.0  0.0  100  0.0  0.0
4400  44C1  2107-981  246G  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
4500  45C3  2107-981  246G  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
4600  46C4  2107-981  246G  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
4700  47BF  2107-981  246G  6.8  99.5  6.8  0.0  0.0  27.5  0.0  0.0
4800  48C2  2107-981  246G  47.8  100  47.8  0.0  0.0  94.3  0.0  0.0
4A00  4A27  2107-981  246G  5.5  100  5.5  0.0  0.0  56.4  0.0  0.0
4D00  4D05  2107-981  246G  5.3  100  5.3  0.0  0.0  53.7  0.2  0.5
5000  5040  2107-981  246G  0.3  100  0.3  0.0  0.0  100  0.0  0.0
5100  51A8  2107-981  246G  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
5300  53AB  2107-981  246G  5.1  100  5.1  0.0  0.0  52.5  0.0  0.0
5400  5402  2107-981  246G  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
5800  58BF  2107-981  246G  13.9  100  13.9  0.0  0.0  86.1  0.0  0.0
5A00  5AA9  2107-981  246G  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
5B00  5B70  2107-981  246G  38.4  100  38.4  0.0  0.0  49.3  0.0  0.0
5C00  5C04  2107-981  246G  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
5D00  5D5F  2107-981  246G  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
5F00  5F02  2107-981  246G  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
D600  D611  2107-951  124G  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0
D680  D6D2  2107-951  124G  0.0  0.0  0.0  0.0  0.0  0.0  0.0  0.0

```

Figure 21. CACHSUM Report

There is no graphic version of this report available.

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

RMF 3.1

Cache Summary

- SYSPLEX

Line 1 of 23

Samples: 60

Systems: 3

Date: 04/25/2023

Time: 13.35.00

Range: 60

Sec

CDate: 04/25/2023

CTime: 13.34.57

CRange: 59

Sec

SSI

RMF Cache SSID Information

218

The following details are available for SSID 4100

400

Press Enter to return to the Report panel.

410

430

CUID : 4102

Config: 246G

NVS : Active

440

Type-Mod: 2107-981

Avail : 211G

Config: 16G

450

Offl : 0

Pinned: 0

460

Pinned: 0

470

480

----- Read -----

----- Write -----

----- Read

Tracks

4A0

Rate

Hit

Hit%

Rate

Fast

Hit

Hit%

%

4D0

Norm

197.5

197.5

100

0.4

0.4

0.4

100

99.8

0.0

500

Seq

10.7

10.7

100

0.0

0.0

0.0

100

99.8

0.0

510

CFW

0.0

0.0

0.0

0.0

0.0

0.0

0.0

0.0

530

Total

208.2

208.2

100

0.4

0.4

0.4

100

99.8

540

580

5A0

5B0

5C0

5D00

5D5F

2107-981

246G

0.0

0.0

0.0

0.0

0.0

0.0

0.0

0.0

5F00

5F02

2107-981

246G

0.0

0.0

0.0

0.0

0.0

0.0

0.0

0.0

D600

D611

2107-951

124G

0.0

0.0

0.0

0.0

0.0

0.0

0.0

0.0

D680

D6D2

2107-951

124G

0.0

0.0

0.0

0.0

0.0

0.0

0.0

0.0

Figure 22. CACHSUM Report - SSID Details

Field description

Table 8. Fields in the CACHSUM Report	
Field Heading	Meaning
Device reserve activity can cause a data gatherer interface to wait until a RESERVE has been released. This in turn can cause the cache interval to be much longer than a regular RMF interval.	
Therefore, CDate, CTime, and CRange are used to show the actual point in time to which the cache interval start is related, and the actual cache interval length. All rates shown in the report are based on CRange, not on Range.	
Note: The reported storage capacities for cache and non-volatile storage (NVS) represents only the Cluster Processor Complex in the storage server, that controls the subsystem. Since a typical storage server has two clusters, you must double the reported capacities to get the actual sizes.	
CDate	Date on which the cache interval started.
CTime	Time at which the cache interval started.
CRange	Cache interval time.
SSID	Subsystem identifier: a number assigned during the installation of the subsystem that uniquely identifies the storage subsystem.
CUID	Physical control unit number of the caching subsystem. This is equal to the lowest device number, or the device that has been turned online first, respectively.
Type-Mod	Hardware type and model.
Size	Amount of physical storage that is configured in this storage subsystem (in giga- or megabytes).
I/O Rate	Rate of I/O requests.
Hit %	Percentage of I/Os that were processed within the cache (cache hits) based on the total number of I/Os.

Table 8. 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.

Table 9. Fields in the CACHSUM Report - SSID Details

Field Heading	Meaning
Note: The reported storage capacities for cache and non-volatile storage (NVS) represents only the Cluster Processor Complex in the storage server, that controls the subsystem. Since a typical storage server has two clusters, you must double the reported capacities to get the actual sizes.	
CUID	Physical control unit number of the caching subsystem. This is equal to the lowest device number, or the device that has been turned online first, respectively.
Type-Mod	Hardware type and model.
Cache Config	Amount of physical storage that is installed in this storage subsystem.
Cache Avail	Amount of storage that is available for caching. This is the total cache size minus the amount used by the subsystem for the cache directory, minus the amount pinned and offline storage.
Cache Offl	Amount of storage that is offline because of a host or subsystem error.
Cache Pinned	Amount of storage that is unavailable because a DASD failure is preventing the subsystem from destaging data. The data is pinned in cache.
NVS	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 7 on page 34](#) for a description of the other fields in the pop-up window of the CACHSUM report.

CFACT - Coupling Facility Activity Report

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

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

How to request this report

To request the Coupling Facility Activity report, select **S** from the Primary Menu and then select **7** on the Sysplex Report menu (shown in [Figure 5 on page 22](#)), or enter the following command:

```
CFACT [cfname]
```

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

Special considerations

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

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

Contents of the report

```

Command ==>
RMF 3.1  CF Activity  - TRXPlex  Line 1 of 78
Scroll ==> CSR

Samples: 30  Systems: 2  Date: 04/25/2023  Time: 08.41.30  Range: 30  Sec

CF: X7CFP87  Type  ST E System  CF  --- Sync ---  ----- Async -----
Structure Name  Util  Rate  Avg  Rate  Avg  Chg  Del
                %    %    Serv %    %    %    %

ISGLOCK         LOCK  A  - *ALL  0.0  0.5  3  0.0  0  0.0  0.0
                LOCK   R70  0.5  3  0.0  0  0.0  0.0
                LOCK   R71  0.0  0  0.0  0  0.0  0.0
ISTGENERIC      LIST  AP Y *ALL  0.0  27.1  14  0.0  0  0.0  0.0
                LIST   R70  27.1  14  0.0  0  0.0  0.0
                LIST   R71  0.0  0  0.0  0  0.0  0.0
SYSZWLM_BC772827  CACHE AP Y *ALL  0.0  0.4  8  0.1  126  0.0  0.0
                CACHE   R70  0.4  8  0.1  126  0.0  0.0
                CACHE   R71  0.0  0  0.0  0  0.0  0.0
SYSZWLM_WORKUNIT  CACHE AP N *ALL  0.0  0.0  0  0.0  0  0.0  0.0
                CACHE   R70  0.0  0  0.0  0  0.0  0.0
                CACHE   R71  0.0  0  0.0  0  0.0  0.0
SYSZWLM_7A862827  CACHE AS N *ALL  0.0  0.0  0  0.0  0  0.0  0.0
                CACHE   R70  0.0  0  0.0  0  0.0  0.0
                CACHE   R71  0.0  0  0.0  0  0.0  0.0
THRLSTSCMKP1_1   LIST  A  N *ALL  14.6  1106  6  9.0  121  0.0  0.0
                LIST   R70  1106  6  9.0  121  0.0  0.0
                LIST   R71  0  0  0.0  0  0.0  0.0

```

Figure 23. CFACT Report

There is no graphic version of this report available.

If you place the cursor on any of the lines with coupling facility structure values, a pop-up window appears (one out of Figure 24 on page 40 through Figure 26 on page 41, depending on the structure type), showing details for this structure. The pop-up window from Figure 27 on page 42 appears instead of the one shown in Figure 25 on page 41, if you select a List Structure which also uses SCM storage.

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

```

+-----+
|          RMF Coupling Facility - Structure Details          |
+-----+
| Lock Structure      : ISGLOCK                               |
| Coupling Facility   : X7CFP87                               |
| System             : R70                                     |
|                                                             |
| Structure Size      : 9M      Connection Name : ISGLOCK#R70 |
| % of CF Storage     : 0.0     Jobname         : GRS         |
| List entries       Total : 0     Status        : ACTIVE     |
|                   Current : 0     ASID         : 7           |
| Lock Entries       Total : 1049K  CF Level Req  : 0          |
|                   Current : 6890                                     |
| Contention         (%)   : 0.0                                     |
| False Contention   (%)   : 0.0                                     |
|                                                             |
| Press Enter to return to the Report panel.                 |
| If data is missing, see Help panel.                         |
+-----+

```

Figure 24. CFACT Report - Details for a Lock Structure

```

+-----+
|               RMF Coupling Facility - Structure Details               |
+-----+
| List Structure      : ISTGENERIC                                     |
| Coupling Facility   : X7CFP87                                       |
| System             : R70                                           |
|                                                             |
| Structure Size      : 11M      Connection Name : USIBMT6_T6LA      |
| % of CF Storage     : 0.0      Jobname         : VTAM390          |
| List entries        : 25438    Status          : ACTIVE           |
|   Current           : 47       ASID            : 41              |
| Data Elements       : 499      CF Level Req    : 1               |
|   Current           : 2                                               |
| Lock Entries        : 4                                               |
|   Current           : 0                                               |
| Contention          : (%)      : 0.0                                |
|                                                             |
| Press Enter to return to the Report panel.                        |
| If data is missing, see Help panel.                               |
+-----+

```

Figure 25. CFACT Report - Details for a List Structure

```

+-----+
|               RMF Coupling Facility - Structure Details               |
+-----+
| Cache Structure     : SYSZWLM_BC772827                             |
| Coupling Facility   : X7CFP87                                       |
| System             : R70                                           |
|                                                             |
| Structure Size      : 10M      Connection Name : #R70            |
| % of CF Storage     : 0.0      Jobname         : WLM              |
| Direct. Entries     : 216      Status          : ACTIVE           |
|   Current           : 5        ASID            : 12              |
| Data Elements       : 423      CF Level Req    : 9               |
|   Current           : 30                                               |
| Request Rate        : 0.5                                              |
| Read Rate           : 2.0                                              |
| Write Rate          : 0.5                                              |
| Castout Rate        : 0.0                                              |
| XI Rate             : 0.0                                              |
| Directory Reclaims   : 0                                               |
|                                                             |
| Press Enter to return to the Report panel.                        |
| If data is missing, see Help panel.                               |
+-----+

```

Figure 26. CFACT Report - Details for a Cache Structure

```
+-----+
|                                     RMF Coupling Facility - Structure Details
|
| Press Enter to return to the Report panel.
| If data is missing, see Help panel.
|
| List Structure      : THRLSTSCMKP1_1
| Coupling Facility   : X7CFP87
| System             : R70
|
|                                     More:      +
|
| Structure Size      :      4G    SCM Algorithm Type : KeyPriority1
| % of CF Storage     :      4.1    SCM Space Maximum  :      16G
| Augmented Space     Est Max : 754M    % Used         :      7.5
|                    % Used :      0.3
|
| List entries        Total   : 1428K    SCM List Entries Est Max : 9584K
|                    Current  : 1279K    Current          : 759K
| Data Elements       Total   : 8569K    SCM List Elements Est Max : 57M
|                    Current  : 7152K    Current          : 4242K
|
| Lock Entries        Total   : 1024
|                    Current  :      0
|
+-----+
```

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

Pressing PF8 displays more information as shown in [Figure 28 on page 42](#).

```
+-----+
|                                     RMF Coupling Facility - Structure Details
|
| Press Enter to return to the Report panel.
| If data is missing, see Help panel.
|
| List Structure      : THRLSTSCMKP1_1
| Coupling Facility   : X7CFP87
| System             : R70
|
|                                     More:      -
|
| Data Elements       Total   : 8569K    SCM List Elements Est Max : 57M
|                    Current  : 7152K    Current          : 4242K
|
| Lock Entries        Total   : 1024
|                    Current  :      0
| Contention          (%)    : 0.0
|
| Connection Name     : THRLSTSC0B00010D
| Jobname             : LF170C1P
| Status              : Active
| ASID                : 391
| CF Level Req        : 9
|
+-----+
```

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

Field descriptions

Table 10. Fields in the CFACT Report	
Field Heading	Meaning
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.

Table 10. Fields in the CFACT Report (continued)	
Field Heading	Meaning
ST	Status - can be one of following states in the SYSPLEX (*ALL) view data line for a structure: A Active - structure is allocated and connected to at least one system during the entire MINTIME. AP Active/primary - structure has been active as primary structure during MINTIME (rebuild-old). AS Active/secondary - structure has been active as secondary structure during MINTIME (rebuild-new). I Inactive - structure got disconnected from all systems during MINTIME. N New - structure became allocated and connected to at least one system during MINTIME. PA Active/primary - Structure is the rebuild-old (primary) structure in an asynchronous duplexing rebuild process. SA Active/secondary - Structure is the rebuild-new (secondary) structure in an asynchronous duplexing rebuild process. 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 during MINTIME. Therefore, all values for these structures are reported as blank.
E	Encryption indicator of the CF structure: Y Yes - indicates that the CF structure is encrypted. N No - indicates that the CF structure is not encrypted or cannot be encrypted. A CF structure cannot be encrypted if the structure is only used by a down-level system in the Parallel Sysplex, that does not support encryption of CF structures. - (hyphen) Not applicable - indicates that encryption of lock (LOCK) structures is not supported.
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.
CF Util %	The percentage of CF processor time used by the structure. The sum of the values in this column is less than 100%, because not all CF processor time is attributable to structures. 'N/A' is shown in this field if the CF level is lower than 15.
Sync Rate	Number of hardware operations per second that started and completed synchronously to the coupling facility on behalf of connectors to the structure.
Sync Avg Serv	Average time in microseconds required to satisfy a synchronous coupling facility request for this structure.
Async Rate	Number of hardware operations per second that started and completed asynchronously to the coupling facility on behalf of connectors to the structure.
Async Avg Serv	Average time in microseconds required to satisfy an asynchronous coupling facility request for this structure. This value also includes operations that started synchronously but completed asynchronously.
Async Chg %	Percentage of asynchronous requests for this structure that changed from synchronous to asynchronous because the requests could not be serviced as synchronous request. This field reports only those requests which were changed due to a subchannel busy condition and can be used as an indicator of a shortage of subchannel resources. Request conversions caused by heuristic sync/async algorithms used to optimize the coupling efficiency of workloads using the CF are not included.

Table 10. Fields in the CFACT Report (continued)

Field Heading	Meaning
Async Del %	Percentage of asynchronous hardware operations for this structure being delayed by either subchannel contention or dump serialization. This value can exceed 100% if there are several delays for one request during the MINTIME.
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. - Fields marked with ¹⁾ are only available for a single system, but not for the sysplex view from the *ALL summary line. You get the single systems view by selecting option Detail ==> Yes from the RMF Coupling Facility Report Options (Figure 29 on page 46). - Each rate is reported as <0.1 if the value is greater than 0 but would be rounded to 0. - Fields marked with ²⁾ only apply to List Structures exploiting SCM storage extension. - 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 in bytes that is currently allocated for this structure.
% of CF Storage (MON I)	The percentage of the total coupling facility storage allocated to this structure.
List Entries (MON I)	Number of list entries in use in a LIST or LOCK structure. 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.
SCM Algorithm Type ²⁾ (MON I)	Type of algorithm used by the coupling facility to control the movement of structure objects between coupling facility real storage and storage class memory.
Augmented Space ²⁾ (MON I)	Est Max Estimated maximum amount of augmented space in bytes that can be assigned for this structure. % Used Percentage of maximum augmented space that is in use by the structure.
SCM Space ²⁾ (MON I)	Maximum Maximum amount of storage class memory in bytes that this structure can use. % Used Percentage of maximum storage class memory that is in use by the structure.
SCM List Entries ²⁾ (MON I)	Est Max Estimated maximum number of list entries that can reside in storage class memory for the structure. Current Number of existing structure list entries that reside in storage class memory.

Table 10. Fields in the CFACT Report (continued)	
Field Heading	Meaning
SCM List Elements ²⁾ (MON I)	Est Max Estimated maximum number of list elements that can reside in storage class memory for the structure. Current Number of existing structure list elements that reside in storage class memory.
Lock Entries (MON I)	Number of lock table entries in use in a serialized LIST or a LOCK structure. Total Total number. Current Number of lock table entries in use. Note: This is an approximate number, since it is based on sampling.
Contention (%)	For serialized LIST structures and for LOCK structures only: percentage of all external requests issued by connectors delayed due to contention on a lock.
False Contention (%)	For LOCK structures only: percentage of all external requests issued by connectors that experience "hash contention". This occurs because a hashing algorithm is used to map a lock request to a lock table entry. When more than one lock request maps to a lock table entry, there is the potential for contention delay. You may need to increase the size of the lock table. 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 ¹⁾ (DET)	Name of the last connection from the selected system.
Jobname ¹⁾ (DET)	Name of the job that made the last connect from the selected system.
Status ¹⁾ (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 ¹⁾ (DET)	ASID of the job that made the last connect from the selected system.
CF Level Req ¹⁾ (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.

Table 10. Fields in the CFACT Report (continued)	
Field Heading	Meaning
Read Rate (DET)	Number of occurrences the coupling facility returned data on a read request by any connector (read hit) per second. Directory-only caches will always have a zero value reported since there is no data to be returned.
Write Rate (DET)	Number of occurrences per second data has been written to the cache structure. Directory-only caches will always have a zero value reported since there are no data writes possible.
Castout Rate (DET)	Number of CASTOUT processings per second. Castout is the process of writing changed cache data to permanent storage. This rate is of interest for store-in cache structures (for example, Db2® global buffer pool structures) in determining the volume of changed data being removed from the structure.
XI Rate (DET)	Number of times per second a data item residing in a local buffer pool was marked invalid by the coupling facility. XI values are seen for directory, store-in and store-thru caches. This rate reflects the amount of data sharing among the users of the cache and the amount of write or update activity against the data bases.
Directory Reclaims (DET)	Number of cache directory reclaims happened during the RMF MINTIME. Directory reclaims occur when the total number of used unique entities exceeds the total number of directories. Whenever this shortage of directory entries occurs, the coupling facility will reclaim in-use directory entries associated with unchanged data. All users of that data must be notified that their copy of the data is invalid. As a consequence, it may happen that this data must be re-read from DASD and registered to the coupling facility again. Directory reclaim activity can be avoided by increasing the directory entries for a particular structure.

Report options

```

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

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

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

                                     Available Coupling Facilities
CF5B      CF6B

```

Figure 29. Coupling Facility Report Options Panel

Name

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

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

Type

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

Detail

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

YES

The report contains data for the sysplex and all single systems.

NO

The report contains data for the sysplex only.

Available Coupling Facilities

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

CFOVER - Coupling Facility Overview Report

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

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

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

How to request this report

To request the Coupling Facility Overview report, select **S** from the Primary Menu and then select **5** on the Sysplex Report menu (shown in [Figure 5 on page 22](#)), or enter the following command:

```
CFOVER [cfname]
```

Contents of the report

```

RMF 3.1  CF Overview      - TRXPLEX      Line 1 of 4
Command ==>              Scroll ==> CSR

Samples: 30      Systems: 2      Date: 04/25/2023  Time: 08.41.30  Range: 30      Sec
CF Policy: IXCPOLBB      Activated at: 04/25/2023 08.31.57

--- Coupling Facility --- ----- Processor ----- Req - Storage - --- SCM ---
Name      Type Mod Lvl Dyn Util% Def Shr Wgt  Eff  Rate Size Avail Size Avail
X7CFH89   2817 E64  16 OFF  0.4  2  0      2.0  68.5  15G  13G   0M   0M
X7CFP87   2827 H66  19 OFF  52.6  2  0      2.0  241K  98G  92G  16G  15G
X7CFP89   2827 HA1  19 ON  12.4  1  0      1.0  412.4 100G  96G  512G 512G
X7CFR89   2817 M80  17 THIN 0.0  2  0  200  1.9  2.0  50G  50G   0M   0M

```

Figure 30. CFOVER Report

There is no graphic version of this report available.

Field descriptions

Table 11. Fields in the CFOVER Report	
Field Heading	Meaning
CF Policy	Name of the current coupling facility policy.
Activated at	Date and time the current coupling facility policy was activated.

Table 11. Fields in the CFOVER Report (continued)

Field Heading	Meaning
Coupling Facility	The following information about the coupling facility is provided: Name coupling facility name Type coupling facility processor type Mod coupling facility processor model Lvl coupling facility microcode level Dyn Dynamic CF dispatching status (ON, OFF, or THIN). THIN indicates that coupling thin interrupts are enabled for the coupling facility (only for CFLEVEL 19 or higher). Note: Dynamic CF dispatching is provided by PR/SM and available to all CF partitions with shared engines. It allows the installation to limit the impact of polling for CFs with low activity rates. The amount of CP resource used by the CF is reduced. There is, however, a performance trade-off when working with dynamic dispatching: though the CPU resource consumed by the CF is reduced, the responsiveness of the CF partition is also reduced. In Parallel Sysplex environments with a CFLEVEL 19 or higher, it is recommended to enable coupling thin interrupts for shared-engine coupling facilities.
Processor	The following information about the processors within the coupling facility is provided: Util% Percentage of processor utilization by 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 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. Def Number of logical processors defined for the coupling facility. Shr Number of shared processors defined for the coupling facility. Wgt Average weight of shared logical processors. This value is not displayed if no shared processors are assigned to this CF. Eff Number of effective available logical processors in a shared environment. This value is only useful in a CFCC environment. CFCC measures the time of real command execution as well as the time waiting for work. The reported value shows the ratio of the LPAR dispatch time (CFCC execute and wait time) to the RMF MINTIME length. For example, if a CFCC CEC contains 6 logical processors, and the measured CF LPAR has two logical processors and is limited at 5%, the number of effective logical processors is 0.3.
Req Rate	The sum of all requests (internal and external) that utilize the subchannel. Specifically: • External requests to send/receive data on behalf of a structure. The sum of synchronous and asynchronous requests completed against any structure within this coupling facility per second. This includes requests that changed from synchronous to asynchronous. • Internal requests that utilize the subchannels (but are not aggregated by the structure). 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 coupling facility storage in bytes, including both allocated and available space.
Storage Avail	The amount of coupling facility space in bytes that is not allocated to any structure, not allocated as dump space, and not allocated as augmented space.

Table 11. Fields in the CFOVER Report (continued)	
Field Heading	Meaning
SCM Size	The total amount of coupling facility storage class memory in bytes which may be concurrently used as structure extensions.
SCM Avail	The total amount of available storage class memory in bytes.

CFSYS - Coupling Facility Systems Report

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

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

How to request this report

To request the Coupling Facility Systems report, select **S** from the Primary Menu and then select **6** on the Sysplex Report menu (shown in [Figure 5 on page 22](#)), or enter the following command:

```
CFSYS [cfname]
```

Contents of the report

Command ==>		RMF 3.1 CF Systems		- TRXPLEX		Line 1 of 12		Scroll ==> CSR	
Samples: 60		Systems: 2		Date: 04/25/2023	Time: 15.55.00	Range: 30	Sec		
CF Name	System	Subchannel Delay %	Busy %	-- Paths -- Avail Delay %	-- Sync --- Rate Avg Serv	----- Async ----- Rate Avg Serv	Chng %	Del %	
X7CFH89	R70	0.0	0.0	4	0.0	0.8	20	53.0	94 0.0 0.0
	R71	0.0	0.0	4	0.0	0.7	17	54.2	105 0.0 0.0
X7CFP87	R70	0.0	0.1	4	0.0	1170	22	135.7	51 0.0 0.0
	R71	0.0	4.4	4	0.0	246K	5	193.4	37 0.0 0.0
X7CFP89	R70	0.0	0.0	4	0.0	451.3	5	3.5	92 0.0 0.0
	R71	0.0	0.1	4	0.0	229.5	22	224.2	55 0.0 0.0

Figure 31. CFSYS Report

There is no graphic version of this report available.

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

```

RMF 3.1  CF Systems      - TRXPLEX                      Line 1 of 12
Command ==>                               Scroll ==> CSR
Samples:  60      Systems: 2      Date: 04/25/2023   Time: 15.55.00   Range:  30   Sec

CF +-----+
    RMF Coupling Facility - Subchannels and Paths
    Press Enter to return to the Report panel.
X7C Details for System      : R70
    Coupling Facility      : X7CFP89
X7C Subchannels Generated : 28
    In Use                 : 28
X7C Max                   : 128
    Channel Path Details:
X7C ID Type Operation Mode      Deg Distance CHID  AID Port --IOP IDs--
    C4 CIB  1x IFB  HCA3-0 LR   N          <1  0708 000D  01 02
    C5 CIB  1x IFB  HCA3-0 LR   N          <1  0709 000D  01 02
    C6 CIB  1x IFB  HCA3-0 LR   N          <1  070A 001D  03 08
    C7 CIB  1x IFB  HCA3-0 LR   N          <1  070B 001D  03 08
    EE ICP
    F0 CFP  2GBit                N          1.4  01F0                10
    F1 CFP  1GBit                N          1.3  01F1                10
    F1=Help      F2=SplitScr  F3=End      F6=RMFHelp  F7=Backward
    F8=Forward    F9=SwapScr   F12=Return
+-----+

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

```

Figure 32. CFSYS Report - Subchannels and Paths

Field descriptions

Table 12. Fields in the CFSYS Report	
Field Heading	Meaning
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).
Subchannel Delay %	The percentage of all coupling facility requests z/OS had to delay because it found all coupling facility subchannels busy. If this percentage is high, you should first ensure that sufficient subchannels are defined (see MAX field below). If there are sufficient subchannels and this percentage is still high, it indicates either a coupling facility path constraint or internal coupling facility contention.
Subchannel Busy %	Percentage of the coupling facility subchannel utilization. This value is calculated from the sum of synchronous and asynchronous coupling facility request times related to the MINTIME and to the number of subchannels.
Paths Avail	Number of physical paths (coupling facility channels) available to transfer coupling facility requests between this system and the coupling facility.

Table 12. Fields in the CFSYS Report (continued)

Field Heading	Meaning
Paths Delay %	Percentage of all coupling facility requests that were rejected because all paths to the coupling facility were busy. This value can exceed 100% if requests encounter a <i>path busy</i> condition more than once. 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 z/OS is delaying the coupling facility requests and in effect gating the workload before it reaches the physical paths. Before concluding you have a capacity problem, however, be sure to check that the correct number of subchannels is defined in the I/O generation (see Subchannel Max). PR/SM environment only: If coupling facility channels are being shared among PR/SM partitions, Path Busy behaves differently. Potentially, you have many subchannels mapped to only a few coupling facility command buffers. You could have a case where the subchannels were properly configured (or even under-configured), Subchannel Busy is low, but Path Busy is high. This means the contention is due to activity from a remote partition. 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 asynchronous requests for this structure that changed from synchronous to asynchronous because the requests could not be serviced as synchronous request. This field reports only those requests which were changed due to a subchannel busy condition and can be used as an indicator of a shortage of subchannel resources. Request conversions caused by heuristic sync/async algorithms used to optimize the coupling efficiency of workloads using the CF are not included.
Async Del %	Percentage of asynchronous hardware operations on behalf of connectors from this system being delayed by either subchannel contention or dump serialization.

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

Field Heading	Meaning
Subchannels and Paths	
Subchannels	Subchannel configuration data. Generated Number of subchannels generated by the I/O configuration. This could be more than the number z/OS can optimally use for coupling facility requests. In Use Number of subchannels z/OS is currently using for coupling facility requests. Max Maximum number of coupling facility subchannels z/OS can optimally use for coupling facility requests. The limit is calculated by z/OS to be the number of physical paths to the coupling facility times the number of command buffer sets per path. It represents the number of parallel requests the coupling facility configuration can handle. If this number is less than the subchannels generated by the I/O configuration, you should reduce the number of coupling facility subchannels in the I/O to match this number. Over-specifying subchannels causes unnecessary storage usage and can cause a high number of rejected coupling facility requests due to path busy.

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

Field Heading	Meaning
Channel Path Details	
Note: If the hardware cannot provide values for a measurement, the field remains blank.	
ID	The hexadecimal identifier of a channel path (CHPID) that is connecting a system with the coupling facility. The physical path utilization for these coupling facility links is shown in the CHANNEL report.
Type	Channel path type.
Operation Mode	Channel path operation mode. It describes the data rate bandwidth, protocol, and adapter type of the channel path. A data rate of, for example, 1GBIT denotes a rate of 1.0625 gigabit per second. A bandwidth of, for example, 12X denotes a twelve-fold bandwidth. Protocols: • IFB – InFiniBand • IFB3 – InFiniBand 3 • CEE – Converged Enhanced Ethernet • GEN3 – PCIe third generation protocol Adapter types: • HCA2-O – Host Channel Adapter2-optical • HCA2-O LR – Host Channel Adapter2-optical long reach • HCA3-O – Host Channel Adapter3-optical • HCA3-O LR – Host Channel Adapter3-optical long reach • PCIE-O SR – Peripheral Component Interconnect Express short reach • ROCE LR – RDMA over Converged Ethernet long reach Unknown operation mode: • Unknown
Deg	Character Y in this column indicates that the channel path is operating at reduced capacity (degraded) or not operating at all.
Distance	Estimated distance in kilometers. The value is calculated as follows: $\frac{\text{Average round-trip path time in microseconds}}{10 \text{ microseconds / kilometer}}$ A value of zero means that the time was not measured.
CHID	Physical channel identifier.
AID	The hexadecimal coupling adapter ID associated with the CHPID.
PORT	The hexadecimal port number associated with the CHPID.
IOP IDS	The hexadecimal identifiers of I/O processors (System Assist Processors) to which the channel path is accessible.

CHANNEL - Channel Path Activity Report

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

For all channels that are managed by Dynamic Channel Path Management (DCM), additional information is available. DCM allows an installation to identify channels which they wish to be managed dynamically. These channels are not assigned permanently to a specific control unit, but belong to a pool of channels. Based on workload requirements in the system, these channels are assigned dynamically by DCM. On top of the report, there is a consolidated data section for managed channels displaying the total number of

channel paths for each type and the average activity data. The character **M** as suffix of the acronym for the channel path type is an indicator that the channel is managed by DCM.

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

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

How to request this report

To request the Channel Path Activity report, select **3** from the Primary Menu and then select **12** on the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)), or enter the following command:

```
CHANNEL
```

Special considerations of report output

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

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

Contents of the report

RMF 3.1 Channel Path Activity														Line 1 of 69			
Command ==>														Scroll ==> HALF			
Samples: 60				System: CB88		Date: 04/25/2023		Time: 08.00.00		Range: 60		Sec					
Channel ID	Path No	G	Type	S	Utilization(%)			Read(B/s)		Write(B/s)		FICON Rate	OPS Actv	zHPF Rate	OPS Actv		
					Part	Tot	Bus	Part	Tot	Part	Tot						
	4		*CNCSM		0.1	0.3											
	4		*FC_SM		0.0	0.0	0.0	0	0	0	0	0	0	0	0		
12			OSD	Y	0.0	0.0	0.0	2K	19K	0	0						
14			OSD	Y	0.0	0.0	0.0	5K	531K	511K	514K						
16			OSD	Y	0.4	1.3	0.0	511K	5M	3M	5M						
20			CTC_S	Y	0.0	0.0											
27			CNC_S	Y	0.0	0.0											
2B			CNC_S	Y	1.3	5.2											
2C			CNC_S	Y	0.2	0.5											
30	5		FC_S	Y	0.0	32.5	8.9	205	52M	205	235K	186	1	0	0		
31	5		FC_S	Y	0.0	33.3	8.5	429	50M	330	249K	185	2	0	0		
37	4		FC_S	Y	0.0	0.5	0.1	0	619K	0	42K	24	1	0	0		
38	4		FC_S	Y	0.0	0.5	0.1	0	613K	0	73K	30	2	0	0		
39	4		FC_S	Y	0.0	0.1	0.0	374	23K	0	31K	8	1	0	0		
3A	4		FC_S	Y	0.0	0.1	0.0	365	21K	0	32K	7	1	0	0		
3E	4		FC_S	Y	0.0	0.0	0.0	0	10K	0	1K	3	1	0	0		
7C			CNCSM	Y	0.3	0.8											
7D			CNCSM	Y	0.1	0.1											
81	3		FC_S	Y	1.1	14.2	3.2	801K	18M	147K	1M	738	2	132	1		
82	5		FC_S	Y	0.1	0.4	0.2	37K	870K	28K	86K	7	1	36	1		
83	5		FC_S	Y	0.0	0.4	0.2	36K	887K	27K	83K	8	1	36	1		
84	4		FC_S	Y	0.0	0.0	0.0	25	101	0	0	0	1	0	0		
85	3		FC_S	Y	0.4	6.8	0.7	62K	2M	61K	1M	420	1	157	1		
8C	3		FC_S	Y	0.6	10.8	1.4	344K	6M	61K	801K	720	2	0	0		
A6	5		FC_SM	Y	0.0	0.0	0.0	0	0	0	0	0	0	0	0		
B6	5		FC_SM	Y	0.0	0.0	0.0	0	0	0	0	0	0	0	0		
E0			IQD	Y						0	315K						
E1			IQD	Y						0	0						
E2			IQD	Y						0	0						
E3			IQD	Y						0	0						

Figure 33. CHANNEL Report

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

Field descriptions

Table 14. Fields in the CHANNEL Report	
Field Heading	Meaning
Channel Path ID	Hexadecimal channel path identifier (CHPID).
Channel Path No	For each channel type which is managed by DCM, a summary line is shown with the average values for all channels in this group. These summary lines are characterized by an '*' preceding the channel path type, and the number of channels of the group is displayed in column No.
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 4 indicates a throughput of 2Gbit/sec on a FICON Express4 card or a FICON Express2 card.
Channel Path Type	Type of channel path. You may issue the console command <code>D M=CHP(xx)</code> 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 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: 1. On a machine running in LPAR mode, but with only one LPAR defined, the <i>Part</i> columns for the <i>Read</i> , <i>Write</i> and <i>Utilization</i> fields display a zero value for channels of type FC (FICON). 2. When Channel Path Measurement Facility (CPMF) is not available, for example, on z/OS systems running as z/VM or alternate VM guests, RMF uses sampled data from SRM so that the reported channel utilization is only an approximate value. With increasing channel speed, the channel utilization value becomes more and more inaccurate. Therefore, in such cases, RMF does not provide accurate values of FICON channel utilization. Beginning with z990 processors, the channel data from SRM is no longer available. As a result, the channel utilization data on a z/OS system running as z/VM or alternate VM guest, is reported as '-----'	
Utilization (%) Part	The channel path utilization percentage for an individual partition. RMF uses the values provided by CPMF. $\text{Part Utilization (\%)} = \frac{\text{Channel Path Busy Time}}{\text{Channel Path Elapsed Time}} * 100$ For channels like FICON, OSA Express, or OSA Direct Express, which are running in extended CPMF mode, the calculation is as follows: $\text{Part Utilization (\%)} = \frac{\text{LPAR \# of Channel Work Units}}{\text{Max \# of Channel Work Units} * \text{Channel Path Elapsed Time}} * 100$ For some channels, like OSAEGbE, FICON EXPRESS/EXPRESS2, this value reflects the microprocessor utilization. For hipersockets, this value is not available.

Table 14. Fields in the CHANNEL Report (continued)	
Field Heading	Meaning
Utilization (%) Tot	The channel path utilization percentage for the CPC during an interval. For processors earlier than z990 and shared channels in LPAR mode, where CPMF is not available, the calculation is: $\text{Total Utilization (\%)} = \frac{\text{\# SRM Observations of Channel Path Busy}}{\text{\# SRM samples}} * 100$ For unshared channels in LPAR mode, the value for total utilization is the same as partition utilization. For channels like, for example, FICON, OSA Express, or OSA Direct Express, which are running in extended CPMF mode, the calculation is as follows: $\text{Total Utilization (\%)} = \frac{\text{Total \# of Channel Work Units}}{\text{Max \# of Channel Work Units} * \text{Channel Path Elapsed Time}} * 100$ For some channels like OSAEGbE, FICON EXPRESS/EXPRESS2, this value reflects the microprocessor utilization. For hipersockets, this value is not available.
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. For hipersockets, this value is not available.
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 CPC. For hipersockets, this value is not available.
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 CPC.
FICON OPS	Rate Number of native FICON operations per second. Actv The average number of native FICON operations that are concurrently active during the report interval.
zHPF OPS	Rate Number of zHPF (High Performance FICON) operations per second. Actv The average number of zHPF operations that are concurrently active during the report interval.

Monitor III Utility fields

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

Table 15. Additional Fields in the CHANNEL Report	
Field Name	Meaning
CHACDFDR	Number of deferred native FICON operations per second that could not be initiated by the channel due to a lack of available resources.
CHACXDFR	Number of deferred zHPF operations per second that could not be initiated by the channel due to lack of available resources.
CHACNET1	Physical-network identifier (PNET ID) of first channel path port.

Table 15. Additional Fields in the CHANNEL Report (continued)	
Field Name	Meaning
CHACNET2	Physical-network identifier (PNET ID) of second channel path port.
The following fields are only available for HiperSockets:	
CHACTMVC	Total message sent rate.
CHACTSVC	Total message sent size.
CHACTFVC	Total receive fail rate.
CHACPMVC	Rate of messages sent by this LPAR.
CHACPSVC	Average size of messages sent by this LPAR.
CHACPFVC	Rate of messages received by this partition that failed due to an unavailable buffer. The value could indicate that more receive buffers are required.
CHACSFVC	Rate of messages sent by this partition that failed.

CPC - CPC Capacity Report

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

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

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

How to request this report

To request the CPC Capacity report, select **1** from the Primary Menu and then **3** on the Overview Report Selection Menu (shown in [Figure 6 on page 23](#)), or enter the following command:

```
CPC
```

Contents of the report

The **CPC Capacity** report provides:

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

```

RMF 3.1 CPC Capacity
Command ==>
Line 1 of 58
Scroll ==> HALF

Samples: 120      System: TRX2  Date: 04/25/2023  Time: 13.16.00  Range: 120  Sec

Partition:  Z2      2817 Model 722      Boost: Speed
CPC Capacity: 2224  Weight % of Max: 74.2  4h Avg: 41  HWcp%: 71.32
Image Capacity: 60  WLM Capping %: 5.4  4h Max: 185  Group: CGRP0010
MT Mode IIP: 1  Prod % IIP: N/A  AbsMSUCap: Y  Limit: 100*

Partition  --- MSU ---  Cap  Proc  Logical  Util %  - Physical  Util % -
           Def  Act  Def  Num  Effect  Total  LPAR  Effect  Total

*CP
TZ1        130   122  N N N    4.3    11.2   12.5    0.4    10.1   11.2
Z1         150    89  N N N    5.2     9.3    9.6    0.1     3.5    3.6
Z2          50    58  N N Y    2.1    11.5   12.8    0.2     1.7    1.9
Z3         N/A     0  Y N N    2.4     8.8   10.6    0.3     1.5    1.8
PHYSICAL
           0.1     0.2

*ICF
CF1         N N N    1.0    99.9   99.9    0.0     7.1    7.1
CF2         N N N    1.0     0.0    0.0    0.0     0.0    0.0
PHYSICAL
           0.1     0.1

```

Figure 34. CPC Capacity report

Field descriptions

Table 16. Fields in the CPC Capacity Report	
Field Heading	Meaning
All MSU-related values are measured in MSU/h (millions of service units per hour).	
Values for the partition which requested the report	
Partition	Partition name.
Processor/Model	Processor family and model number of the measured system.
Boost	The boost type that was active at the end of MINTIME: N Boost was inactive. zIIP zIIP boost. Speed Speed boost. All zIIP and speed boost were both active.
CPC Capacity	Effective processor capacity available to the central processor complex (CPC), measured in MSU/h.
Image Capacity	Processor capacity available to the z/OS image (partition) which requested the report, measured in MSU/h. The field is calculated as minimum of the following capacities: - the capacity based on the partition's logical CP configuration (including CPs that are online or in standby state (not configured online)). - the defined capacity limit of the partition, if available (image softcap). - the capacity limit of the related WLM capacity group, if the partition belongs to a capacity group. - the absolute physical hardware capping limit. - the capacity based on the hardware group capping limit.
MT Mode IIP	The multithreading mode for processor type zIIP designates the number of active threads for each online logical zIIP core. If this value is greater than 1, multithreading becomes effective for zIIP cores. N/A is shown if the LOADxx PROCVIEW CPU parameter is in effect or no IIP is currently installed or online.

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

Field Heading	Meaning
Prod % IIP	The multithreading IIP core productivity represents the percentage of the maximum IIP core capacity that was used while the IIP cores were dispatched to physical hardware. When this value equals 100% in multithreading mode, all threads on all IIP cores that were configured ONLINE for the complete MINTIME are being used. If the LOADxx PROCVIEW CPU parameter is in effect or no IIP is currently installed or online, no core productivity is calculated and N/A is reported.
Weight % of Max	Average weighting factor in relation to the maximum defined weighting factor for this partition. With 'Initial Capping ON', which the operator can set on the Hardware Management Console, this value is not available and therefore, this field shows '*****' in this case.
WLM Capping %	Percentage of time when WLM capped the partition because the four-hours average MSU value exceeds the defined capacity limit.
4h Avg	Average value of consumed MSU/h during the last four hours.
4h Max	Maximum value of consumed MSUs during the last 4 hours (retrieved from 48 sample intervals of five minutes). This value can be greater than the defined capacity.
AbsMSUCap	Absolute MSU capping is active for the partition: Y or N.
Group	Name of the partition's capacity group. If the partition does not belong to a capacity group, N/A is displayed.
Limit	Capacity limit (in MSUs) defined for the partition's capacity group. An '*' following the limit value indicates that this partition started to be a member of this capacity group less than four hours ago. This partition will have a different view of unused group capacity and, therefore, might cap differently than existing group members.
Values for all configured partitions are grouped by general and special purpose processor types. The term logical processor refers to a logical core if the LOADxx PROCVIEW CORE parameter is in effect.	
Partition	Partition name. Notes: 1. Partitions identified by the name PHYSICAL are not configured partitions. Data reported in these lines includes the time during which a physical CPU was busy, but the time could not be attributed to a specific logical partition. 2. The summary lines (for example, *CP or *ICF) show the total percentages for the indicated processor type. 3. Starting with z9 processors, IFLs (Integrated Facility for Linux) and zAAPs are reported separately and no longer as ICFs (Internal Coupling Facility).
MSU	Millions of unweighted CPU service units per hour: Def Defined MSU capacity limit for the partition. Act Actual MSU consumption based on the logical processor effective dispatch time. These values are only provided for general purpose processors.
Cap Def	The hardware capping option of the partition. Each Cap Def value is a three position character string denoting which hardware capping mechanisms have or have not been applied in the logical partition controls of the Hardware Management Console (HMC) for the partition. The values in the first, second and third position of the string are either Y (Yes) or N (No) and have the following meaning: The first character (Y or N) indicates whether "Initial Capping ON" has been set. The second character (Y or N) indicates whether an absolute physical hardware capping limit (maximal number of CPUs) has been defined. The third character (Y or N) indicates whether an absolute hardware group capping limit (maximal number of CPUs) has been defined. An asterisk (*) to the right of a value indicates that the capping status changed during the report interval.
Proc Num	The number of logical processors which were online during the report interval.

Table 16. Fields in the CPC Capacity Report (continued)	
Field Heading	Meaning
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. $\frac{\text{Total Dispatch Time} - \text{Effective Dispatch Time}}{\# \text{ Physical Processors} * \text{Range Time}} * 100$ The calculation for the PHYSICAL partition is: $\frac{\text{Time PHYSICAL}}{\# \text{ Physical Processors} * \text{Range Time}} * 100$ Time PHYSICAL is the time that could not be attributed to a specific logical partition but was used by PR/SM to control the physical processor (LPAR management time).
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. $\frac{\text{Total Dispatch Time}}{\# \text{ Physical Processors} * \text{Range Time}} * 100$ The Total Dispatch Time for the PHYSICAL partition includes the time during which a physical CPU was busy, but the time could not be attributed to a specific logical partition. This time includes the time PR/SM was controlling the physical processor (LPAR management time), as well as any other time the processor was busy for any reason such as managing coupling facility traffic.
HWcp%	If the partition belongs to a hardware group and hardware group capping is enabled, the percentage of the hardware group cap currently used by all LPARs in that hardware group is displayed. Otherwise, N/A is displayed.

CRYACC – Crypto Accelerator Activity Report

The cryptographic accelerator (CRYACC) activity report is used to investigate performance problems that are related to the usage of cryptographic hardware configured in accelerator mode.

How to request this report

To request the Crypto accelerator activity report, select **S** from the Primary Menu, then select **17** from the Sysplex Report Selection Menu (shown in [Figure 5 on page 22](#)) or enter one of the following commands:

```
CRYACC
CRA
```

Contents of the report

For each cryptographic accelerator card in the sysplex, the cryptographic accelerator activity (CRYACC) report provides measurements about public key operations (RSA cryptography operations), both at the CPC and LPAR (cryptographic usage domain) level. The data for cryptographic accelerators is showing details for the two available algorithms, modular exponentiation (ME) and Chinese Remainder Theorem (CRT) for available key lengths (1024, 2048, and 4096 bit). Displayed are the rate at which requests are processed by an adapter card, the average execution time, and the utilization percentage.

```

RMF 3.1  Crypto Acc Activity - ENGTEST3                               Line 1 of 15
Command ===>                                                           Scroll ===> CSR

Samples: 100      Systems: 4      Date: 04/25/2023  Time: 06.00.00  Range: 100   Sec

Type  ID  --CPC---  -System-  -Key  ----- ME RSA -----  ---- CRT RSA -----
          Len   Rate ExTime Util%      Rate ExTime Util%

CEX6A  9  M88                1024  2705  0.023  6.1  64.25  0.078  0.5
CEX6A  9  M88                2048  223.7  0.040  0.9  274.6  0.334  9.2
CEX6A  9  M88                4096  661.6  0.122  8.0  1308  0.593  77.6
CEX6A  9  M88          S24    1024  2486  0.012  2.9  0.000  0.000  0.0
CEX6A  9  M88          S24    2048  0.000  0.000  0.0  0.000  0.000  0.0
CEX6A  9  M88          S24    4096  0.000  0.000  0.0  0.000  0.000  0.0
CEX6A  9  M88          S25    1024  218.5  0.147  3.2  64.25  0.078  0.5
CEX6A  9  M88          S25    2048  223.7  0.040  0.9  274.6  0.334  9.2
CEX6A  9  M88          S25    4096  661.6  0.122  8.0  1308  0.593  77.6
CEX3A  5  P88                1024  0.000  0.000  0.0  5472  0.129  70.8
CEX3A  5  P88                2048  0.000  0.000  0.0  0.000  0.000  0.0
CEX3A  5  P88                4096  0.000  0.000  0.0  0.000  0.000  0.0
CEX5A  9  S89                1024  0.000  0.000  0.0  0.000  0.000  0.0
CEX5A  9  S89                2048  0.000  0.000  0.0  147.8  0.522  7.7
CEX5A  9  S89                4096  0.000  0.000  0.0  673.0  1.371  92.2

```

Figure 35. Crypto Accelerator Activity Report

Field descriptions

Table 17. Fields in the Cryptographic accelerator activity Report	
Field Heading	Meaning
Type	Type that defines the cryptographic accelerator: Type Meaning CEX3A Crypto Express3 Accelerator. CEX4A Crypto Express4S Accelerator. CEX5A Crypto Express5S Accelerator. CEX6A Crypto Express6S Accelerator. CEX7A Crypto Express7S Accelerator.

Table 17. Fields in the Cryptographic accelerator activity Report (continued)	
Field Heading	Meaning
ID	Index that specifies the cryptographic hardware function.
CPC	Name of the CPC that used the cryptographic hardware function with the respective ID.
System	Name of the partition that used the cryptographic hardware function with the respective ID.
Key Len	RSA key length for each cryptographic accelerator and for each available RSA operation format (ME or CRT).
ME RSA	Rate, average execution time (in milliseconds), and utilization percentage of all operations in ME-format (one line for each used RSA key length).
CRT RSA	Rate, average execution time (in milliseconds), and utilization percentage of all operations in CRT-format (one line for each used RSA key length).

Cursor-sensitive control on the CRYACC Report

Cursor-sensitivity on any value in the tabular part of the report links back to the Cryptographic Hardware Overview (CRYOVW) report for total numbers, showing only the lines for cryptographic accelerator functions.

Cursor-sensitive control of the sysplex name and of the *Systems* field in the report header leads to the Data Index screen.

Report options

The Report Options panel is exactly the same as for the CRYOVW report shown in [Figure 37 on page 64](#). The only difference is, that the *Function* report option is ignored in the CRYACC report.

CRYOVW – Crypto Hardware Overview Report

The cryptographic hardware overview (CRYOVW) report is used to investigate performance problems that are related to the usage of various cryptographic hardware functions in the system. The report provides information about cryptographic hardware configured in accelerator, CCA coprocessor, or PKCS11 coprocessor mode.

How to request this report

To request the Crypto hardware overview report, select **S** from the Primary Menu and then **16** from the Sysplex Report Selection Menu (shown in [Figure 5 on page 22](#)) or enter one of the following commands:

```
CRYOVW
CRO
```

Contents of the report

For each cryptographic adapter card in the sysplex, the cryptographic hardware overview (CRYOVW) report provides measurements at both the CPC and LPAR (cryptographic usage domain) levels. The rate at which requests are processed by an adapter card, the average execution time, and the utilization percentage is displayed.

For cryptographic adapters configured in CCA coprocessor mode, the rate, execution time, and utilization percentage of RSA key-generation operations are additionally reported.

```

RMF 3.1  Crypto HW Overview  - ENGTEST3          Line 1 of 22
Command ==>                               Scroll ==> CSR

Samples: 100      Systems: 4      Date: 04/25/2023  Time: 06.00.00  Range: 100  Sec

Type  ID  --CPC--  -System-  Rate  Exec  Util%  --- Key Gen ---
              Rate ExTime Util%

CEX5C  0  M88              6223  0.160  99.7   0.730  0.191  0.0
CEX5C  0  M88      S24      66.46  0.146   1.0   0.000  0.000  0.0
CEX5C  0  M88      S25      6157  0.160  98.7   0.730  0.191  0.0
CEX5P  3  M88              949.4  1.053  100
CEX5P  3  M88      S24      280.5  1.102  30.9
CEX5P  3  M88      S25      668.9  1.032  69.0
CEX6A  9  M88              5237  0.195  100
CEX6A  9  M88      S24      2486  0.012   2.9
CEX6A  9  M88      S25      2751  0.361  99.4
CEX6C  12 M88              6111  0.163  99.7   0.030  0.127  0.0
CEX6C  12 M88      S24      317.7  0.269   8.5   0.000  0.000  0.0
CEX6C  12 M88      S25      5793  0.157  91.1   0.030  0.127  0.0
CEX4C  1  P88              468.8  0.706  33.1   0.000  0.000  0.0
CEX3C  4  P88              465.0  0.685  31.8   0.000  0.000  0.0
CEX3A  5  P88              5472  0.129  70.8
CEX4P  13 P88              28.19  12.22  34.5
CEX4C  14 P88              460.6  0.688  31.7   0.000  0.000  0.0
CEX4P  15 P88              26.62  12.65  33.7
CEX5P  4  S89              826.4  0.923  76.3
CEX5P  5  S89              827.1  0.919  76.0
CEX5C  6  S89              1219  0.820  99.9   0.000  0.000  0.0
CEX5A  9  S89              820.8  1.218  100

```

Figure 36. Crypto Hardware Overview Report

Field descriptions

Table 18. Fields in the Crypto Hardware Overview Report	
Field Heading	Meaning
Type	Type that defines the cryptographic hardware function: Type Meaning CEX3A Crypto Express3 Accelerator. CEX4A Crypto Express4S Accelerator. CEX5A Crypto Express5S Accelerator. CEX6A Crypto Express6S Accelerator. CEX7A Crypto Express7S Accelerator. CEX3C Crypto Express3 Coprocessor. CEX4C Crypto Express4S Coprocessor. CEX5C Crypto Express5S Coprocessor. CEX6C Crypto Express6S Coprocessor. CEX7C Crypto Express7S Coprocessor. CEX4P Crypto Express4S PKCS11 Coprocessor. CEX5P Crypto Express5S PKCS11 Coprocessor. CEX6P Crypto Express6S PKCS11 Coprocessor. CEX7P Crypto Express7S PKCS11 Coprocessor.
ID	Index that specifies the cryptographic hardware function.
CPC	Name of the CPC that used the cryptographic hardware function with the respective ID.
System	Name of the partition that used the cryptographic hardware function with the respective ID.
Rate	Rate of all operations on this cryptographic hardware function.
Exec Time	Average execution time (in milliseconds) of all operations on this cryptographic hardware function.
Util%	Total utilization percentage of this cryptographic hardware function.
Key Gen Rate ExTime Util%	Rate, average execution time (in milliseconds), and utilization percentage of RSA-key-generation operations, if the cryptographic hardware function is configured in CCA Coprocessor mode.

Cursor-sensitive control on the CRYOVW Report

For Crypto Express adapters that are configured in accelerator or PKCS11 coprocessor mode, cursor sensitivity on any value within a line in the tabular part links to the respective CRYACC or CRYPKC Monitor III reports that display more detailed information about these cryptographic functions.

Cursor-sensitive control of the sysplex name and of the *Systems* field in the report header leads to the Data Index screen.

Report options

The Report Options panel for the Crypto reports allows you to specify options for this report.




Figure 37. Crypto Report Options Panel

Scope

Either ALL or the name of one of the CPCs or systems that are available in the sysplex as shown in the **Available CPCs and Systems** section.

Function

Specification of the cryptographic functionality that data is reported in the CRYOVW:

ALL

The report contains data for all cryptographic functions.

ACC

The report contains data for cryptographic hardware that is configured in accelerator mode.

CCA

The report contains data for cryptographic hardware that is configured in CCA coprocessor mode.

PKC

The report contains data for cryptographic hardware that is configured in PKCS11 coprocessor mode.

Inactive

Specification about display of inactive cryptographic cards:

YES

Include all card lines in the report, even if the lines do not contain any activity data.

NO

Do not include card lines that do not contain any activity data.

Available CPCs and Systems

The list of CPCs and systems that are currently defined to the sysplex.

CRYPKC – Crypto PKCS11 Coprocessor Activity Report

The cryptographic PKCS11 coprocessor activity report is used to investigate performance problems that are related to the usage of cryptographic hardware configured in PKCS11 coprocessor mode.

How to request this report

To request the Crypto PKCS11 Coprocessor Activity report, select **S** from the Primary Menu, and then **18** from the Sysplex Report Selection Menu (shown in [Figure 5 on page 22](#)), or enter one of the following commands:

CRYPKC
CRP

Contents of the report

For each cryptographic PKCS11 coprocessor card in the sysplex, the cryptographic PKCS11 coprocessor activity (CRYPKC) report provides measurements about secure public-key operations that are executed by cryptographic symmetric-key and asymmetric-key functions (PKCS11 cryptography). The rate at which requests are processed by an adapter card and the utilization percentage are displayed both at the CPC and LPAR (cryptographic usage domain) level.

```

Command ==>          RMF 3.1  Crypto PKCS11 Act.  - ENGTEST3          Line 1 of 7
                                                                Scroll ==> CSR
Samples: 100          Systems: 4          Date: 04/25/2023  Time: 06.00.00  Range: 100  Sec
Type   ID   --CPC--- -System-  -Asym Slow  -Asym Fast  -Symm Part  -Symm Cmpl
          Rate Utl%   Rate Utl%   Rate Utl%   Rate Utl%
CEX5P   3   M88           S24       533.7 52.6   399.1 46.6   0.000 0.0   16.07 0.7
CEX5P   3   M88           S24       231.4 21.4   40.57 9.2    0.000 0.0    8.450 0.4
CEX5P   3   M88           S25       302.3 31.2   358.5 37.4   0.000 0.0    7.620 0.3
CEX4P  13   P88           S25       13.97 3.3    0.000 0.0    0.000 0.0   14.14 31.1
CEX4P  15   P88           S25       13.07 3.8    0.000 0.0    0.000 0.0   12.90 29.3
CEX5P   4   S89           S25       341.6 34.3   484.8 42.0   0.000 0.0    0.000 0.0
CEX5P   5   S89           S25       341.8 34.3   485.3 41.7   0.000 0.0    0.000 0.0

```

Figure 38. Crypto PKCS11 Coprocessor Activity Report

Field descriptions

Table 19. Fields in the Crypto PKCS11 Coprocessor Activity Report	
Field Heading	Meaning
Type	Type that defines the cryptographic PKCS11 coprocessor: Type Meaning CEX4P Crypto Express4S PKCS11 Coprocessor. CEX5P Crypto Express5S PKCS11 Coprocessor. CEX6P Crypto Express6S PKCS11 Coprocessor. CEX7P Crypto Express7S PKCS11 Coprocessor.
ID	Index that specifies the cryptographic hardware function.
CPC	Name of the CPC that used the cryptographic hardware function with the respective ID.
System	Name of the partition that used the cryptographic hardware function with the respective ID.
Rate Utl%	Rate and utilization percentage of executed PKCS11 operations, which are categorized by cryptographic function type: Type Meaning Asym Slow Slow asymmetric-key function. Asym Fast Fast asymmetric-key function. Symm Part Symmetric-key function that returns partial or incremental results. Symm Cmpl Symmetric-key function that returns a complete or final result.

Cursor-sensitive control on the CRYPKC Report

Cursor-sensitivity on any value in the tabular part of the report links back to the Crypto HW Overview (CRYOVW) report for total numbers, showing only the lines for cryptographic PKCS11 coprocessor functions.

Cursor-sensitive control of the sysplex name and of the *Systems* field in the report header leads to the Data Index screen.

Report options

The Report Options panel is exactly the same as for the CRYOVW report shown in Figure 37 on page 64. The only difference is, that the *Function* report option is ignored in the CRYPKC report.

DELAY - Delay Report

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

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

How to request this report

To request the Delay report, select **1** from the Primary Menu, and then select **4** on the Overview Report menu (shown in Figure 6 on page 23) or enter the following command using the format:

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

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

```
DELAY T, TSOPRIME
```

Contents of the report

RMF 3.1 Delay Report															Line 1 of 58	
Command ==>															Scroll ==> HALF	
Samples: 120		System: MVS1	Date: 04/25/2023	Time: 12.00.00	Range: 120 Sec											
Name	CX	Service Class	Cr	WFL %	USG %	DLY %	IDL %	UKN %	PRC	DEV	STR	SUB	OPR	ENQ	Primary Reason	
*SYSTEM				49	1	1	62	36	0	0	0	0	0	0		
*TSO				56	1	1	95	2	0	0	0	1	0	0		
*BATCH				39	2	4	0	94	1	0	0	4	0	0		
*STC				40	0	1	51	48	0	0	0	0	0	0		
*ASCH					0	0	0	0	0	0	0	0	0	0		
*OMVS					0	0	0	100	0	0	0	0	0	0		
*ENCLAVE					0	0	0	0	0	N/A	0	N/A	N/A	N/A		
JES2	S	SYSSTC		0	0	1	0	99	0	1	0	0	0	0	SCLSP4	
BMAI	T	PRDTSO	S	16	9	66	13	4	0	1	0	65	0	0	HSM	
HSM	S	STCCMD	C	30	26	62	0	23	0	1	0	0	50	11	Mount	
HIRW2	B	BATCHMED		35	6	14	0	1	2	1	0	12	0	0	HSM	
TCPNET	S0	SYSSTC	SC	60	3	2	0	97	2	0	0	0	0	0	NET	
MASTER	S	SYSTEM		67	3	2	0	95	0	2	0	0	0	0	M00202	

Figure 39. 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 20. 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 (-----).
Cr	An indication of whether WLM managed the address space has one of the following values during the report interval: C CPU critical S Storage critical SC Both storage and CPU critical I Implicitly CPU critical SI Both Storage and Implicitly CPU critical
WFL %	The workflow percentage of the job or job group. “Address space workflow (%)” on page 11 shows the formula used to calculate this value.
USG %	The using percentage for the job or job group. “Address space using (%)” on page 12 shows the formula used to calculate this value.
DLY %	The delay percentage for the job or job group. See “Address space delay (%)” on page 13 for more information.

Table 20. 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 <i>idling</i> 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: $\text{IDL \%} = \frac{\text{\# Idle Samples}}{\text{\# Samples}} * 100$ 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: $\text{IDL \%} = \frac{\sum \text{Idle Samples}}{\text{\# Samples} * \text{Avg \# Address Spaces}} * 100$ 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 20. 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 20. 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 frequently while the reported job was delayed. STR This field identifies the cause of the largest percentage of delay: COMM common storage paging (includes shared pages) LOCL local storage paging (includes shared pages) VIO virtual I/O paging SWAP swap-in delay OUTR swapped out and ready XMEM cross memory address space HIPR standard hiperspace paging delays DEV This field contains the volume serial number of the device that the reported job was most frequently delayed for. SUB This field contains either JES, HSM, or XCF depending on which subsystem is causing the most delay. OPR This field contains <i>Message</i> if most of the delay was due to a message or <i>Mount</i> if most of the delay was due to a mount request. The field can contain <i>QUIESCE</i> 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 delay variation of the Job Delay 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 <i>z/OS MVS Planning: Workload Management</i>).

Monitor III Utility fields

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

Table 21. Additional Fields in the DELAY Report

Field Heading	Meaning
JDELJES	JES delay percentage
JDELHSM	HSM delay percentage
JDELXCF	XCF delay percentage
JDELMNT	Operator mount delay percentage
JDELMES	Operator message delay percentage
JDELQUI	Operator quiesce delay percentage
JDEPRPCL	Report Class Name
JDELSCLP	Service Class Period
JDELENCT	Enclave Token

Cursor-sensitive control on the Delay report

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

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

Using cursor-sensitive control on an enclave name displays a pop-up panel that shows information you extracted from the WLM Enclave Classification Data (ECD) control block. You can use this information to identify the transactions that are processed in the enclave. See [“Enclave Classification Attributes”](#) on page 97 for an example.

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

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

Report options

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

Command ==>

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

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

GPMSERVE
OE

Available Service classes

OMVSKERN
STCDEF
TSODEF
SYSTEM
SYSSTC

Figure 40. DELAY Report Options Panel

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

Class

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

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

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

Service Class

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

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

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

Summary

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

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

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

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

Criterion

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

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

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

Jobs

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

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

Available Service classes

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

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

Job Selection/Exclusion Option panel

If you select YES for Jobs on the Delay Report Options panel, RMF displays a Job Selection/Exclusion panel shown in [Figure 41 on page 73](#).

```

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 41. 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 **s** to its left, under **SEL**; to exclude a job, type **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*-----
X    *ALL

```

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

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

DEV - Device Delays Report

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

How to request this report

To request the Device Delays report, select **3** from the Primary Menu and then select **2** on the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)), or enter the following command:

```
DEV [job_class,service_class]
```

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

```
DEV T, TSOPRIME
```

Contents of the report

```

RMF 3.1  Device Delays
Line 1 of 57
Command ==> Scroll ==> HALF

Samples: 100      System: MVS1  Date: 04/25/2023  Time: 10.03.20  Range: 100  Sec

Jobname  C  Service  DLY  USG  CON  -----  Main Delay  Volume(s) -----
          C  Class   %   %   %   % VOLSER  % VOLSER  % VOLSER  % VOLSER
MARYPATM B NRPRIME  70  51  54  70 TSOL11  1 DUMP00
MICHAEL  B NRPRIME  39  15  14  39 BPXLK1
MCPDUMP  S SYSSTC  36  18  20  36 D24PK2
CHARLES  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 TS0063  1 HSM004

```

Figure 42. DEV Report

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

Field descriptions

Table 22. 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.
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.

Table 22. Fields in the DEV Report (continued)	
Field Heading	Meaning
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. Note: 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 job delay. % 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, PROCU, STOR, STORC,
STORF, STORM, and XCF.

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

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

Available Service classes
ZOSUN      OMVSKERN  PRDTSO  STCCMD  SYSTEM  SYSSTC

```

Figure 43. 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 71](#) (the Delay Report Options panel) for a description of the fields.

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

DEVN - Device Activity Report

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

The report provides the capability to select a subset of all available devices and to sort the displayed devices. You can rearrange the displayed list of devices by any activity category you want and focus on devices with common characteristics (for example, same volser number or device number prefix, or devices having the same type or are connected to the same control unit type). This is done using cursor-sensitive control (see [“Cursor-sensitive control” on page 77](#)).

How to request this report

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

```
DEVN
```

Contents of the report

```

RMF 3.1  Device Activity
Command ==>
Line 1 of 118
Scroll ==> HALF

Samples: 100      System: AQTS  Date: 04/25/2023  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
HSM117 0703 33903 3990-3 S 52.2 .015 .000 76 22 54 0 0.2 0.6
HSM015 006E 33903 3990-3 S 11.1 .024 .001 26 3 20 3 0.0 0.3
TS0060 0056 33903 3990-3 S 8.9 .034 .001 30 9 18 3 DB 2 0.1 0.2
D22SHR 0B70 3380 3880-3 S 13.5 .014 .000 18 2 16 0 0.0 0.1
HSM110 0043 33902 3990-3 S 68.4 .008 .000 53 43 3 7 0.4 0.1
CLR014 0149 33903 3990-3 S 37.2 .005 .000 19 11 7 1 0.1 0.1
TS0024 0842 33902 3990-3 S 5.8 .015 .000 9 1 8 0 0.0 0.1
HLPV0L 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
HSM104 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 44. 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 377).

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.

To get the subchannel set ID of the displayed devices, either request the DEVR or DEVT report where the first digit of the displayed five-digit device number identifies the subchannel set ID to which the device is physically configured.

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 25 on page 79](#)) except for the following:

Table 23. Fields in the DEVN Report	
Field Heading	Meaning
Devices reported:	The criteria selected for the devices being reported. The devices being reported can be selected using cursor-sensitive control from the fields listed under Device Identification.
Report is sorted by:	The sort criteria for the devices being reported. The sort criteria can be selected using cursor-sensitive control from any of the fields listed under the columns between Activity and Jobs.
Activity IosQ	The average number of seconds an I/O request must wait on an IOS queue before a SSCH instruction can be issued. A delay occurs when a previous request to the same sub-channel is in progress. The value is calculated as: $\text{IosQ} = \frac{\text{IOS Queue Count} / \# \text{ Samples}}{\text{Device Activity Rate}}$ This field is not shown on the DEVR report but is available in the ISPF table of the DEVR report.
Pending Reasons	Only the highest delay reason and percentage is listed.
Jobs	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 are 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 24. 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.

Table 24. DEVN Report - Cursor-sensitive Control for Select and Sort (continued)		
Cursor- sensitive Column	Cursor Position within Column	Result
Type	any	Device Activity report for all online devices with the same device type.
CU	any	Device Activity report for all online devices with the same CU type.
S	--	No cursor-sensitive control.
Rsn %	--	No cursor-sensitive control.
All other	any	Device Activity report sorted in descending order by the selected column.

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

DEVR - Device Resource Delays Report

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

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

DB

Device busy delay

CMR

Initial command response time

How to request this report

To request the DEVR report, select **3** from the Primary menu, and then select **3** on the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)) or enter the following command:

```
DEVR [volser]
```

Contents of the report

```

RMF 3.1  Device Resource Delays                               Line 1 of 374
Command ==>                                                    Scroll ==> HALF

Samples: 100      System: MVS1  Date: 04/25/2023  Time: 10.03.20  Range: 100  Sec

Volume S/  Act  Resp  ACT  CON  DSC  PND %,  DEV/CU  Service  USG  DLY
 /Num  PAV  Rate  Time  %    %    %  Reasons Type  Jobname  C Class  %    %
160444 S   17.5  1.2   2    2    0  PND   0 33909  GPMSRVPM S GPMSERVE  0    1
08803   1.2H   0.3  0.4   0    0    0  PND   0 33903  RMF      S SYSSTC   2    0
SYSUSR S   2*    0.8  0.9   0    0    0  PND   0 33909  *MASTER* S SYSTEM   1    0
0D440   2*    0.8  0.9   0    0    0  PND   0 33903  CATALOG  S SYSTEM   1    0
MVSTGT S   1.4  0.7   1    1    0  PND   0 33903  GPMSRVPM S GPMSERVE  1    0
07790   2105   0.8  0.9   0    0    0  PND   0 33909  GPMSRVPM S GPMSERVE  1    0
SYSXCP S   1.4  0.7   1    1    0  PND   0 33903  XCFAS    S SYSTEM   1    0
0D32A   2105
```

Figure 45. 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 25. 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. The first digit of the device number represents the ID of the subchannel to which the I/O device is physically configured.
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. If the number has been changed during the report interval, it is followed by an '*'. If the device is a HyperPAV base device, the number is followed by H, for example, 1.2H. The value is the average number of HyperPAV volumes (base and alias) for that range. $\text{Average \# of HPAV devices} = \frac{\text{Accumulated \# of HPAV devices}}{\text{Number of Samples}}$
Act Rate	The rate per second that I/O instructions (SSCH, RSCH, and HSCH) to a device completed successfully. The calculation is: $\text{Act Rate} = \frac{\text{\# I/O Instructions}}{\text{Range Time}}$
Resp Time	The average response time (in milliseconds) that the device required to complete an I/O request. The calculation is: $\text{Resp Time} = \frac{\text{Active Time}}{\text{\# I/O Instructions}} + \text{IOS Queue Time}$
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: $\text{ACT \%} = \text{PEND \%} + \text{CON \%} + \text{DSC \%}$ 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.

Table 25. Fields in the DEVR Report (continued)

Field Heading	Meaning
CON %	The percent connect time. See the description under % ACT. RMF calculates the value as follows: $\text{CON \%} = \frac{\text{Accumulated Connect Time}}{\text{Range Time}} * 100$
DSC %	The percent disconnect time. See the description under %ACT. RMF calculates the value as follows: $\text{DSC \%} = \frac{\text{Accumulated Disconnect Time}}{\text{Range Time}} * 100$
Note: 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: $\text{PND \%} = \frac{\text{Accumulated Pending Time}}{\text{Range Time}} * 100$ 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. $\text{DLY DB\%} = \frac{\text{Accumulated DB Delay Time}}{\text{Range Time}} * 100$ 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. $\text{DLY CMR\%} = \frac{\text{Accumulated CMR Delay Time}}{\text{Range Time}} * 100$ 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 25. 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.
USG %	The percentage of time when the job has had a request accepted by the channel for the specified Volume, but the request is not yet complete.
DLY %	Delay the waiting job (address space) is experiencing because of contention for a specific volume during the report interval, expressed as a percentage.

Monitor III Utility fields

Table 26 on page 81 shows additional fields for the Device Resource Delay report.

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

Table 26. 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  D608F0  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 46. DEVR Report Options Panel

VOLSER

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

Enter

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

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

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

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

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

Available Volumes

The list of the online volumes in the system.

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

DEVT - Device Activity Trend Report

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

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

How to request this report

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

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

Contents of the report

RMF 3.1 CLR010 Activity Trend										Line 1 of 20				
Command ==>										Scroll ==> HALF				
Samples: 100		System: AQT5		Date: 04/25/2023		Time: 11.58.20		Range: 100		Sec				
VolSer: CLR010		Number: 00051		Type and CU-Type: 33903 3990-3										
Latest: 04/25/2023		at 11.58.20		Range/Line: 100 Sec										
Earliest: 04/25/2023		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 47. DEVT Report

The DEVT report has two parts.

- The top part provides information about the selected device, its volser, device number, device and control unit type, and information about the reported range.
- The bottom part is based on the DEVN report.

Each row is preceded by a time stamp to identify the start time of the reporting range. The device activity columns are exactly the same as those shown on the DEVN report.

At the right, a column showing the device workflow percentage is added. The workflow column is calculated from the average number of users using or being delayed for the device from the DEVN ISPF table. Please keep in mind that the value is not as precise as workflow values shown on the Workflow/Exception report.

Field descriptions

Table 27. Fields in the DEVT Report	
Field Heading	Meaning
VolSer:	The name of an online volume.
Number:	The five-digit device number where the volume is mounted. The first digit represents the ID of the subchannel to which the I/O device is physically configured.
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.

Table 27. Fields in the DEVT Report (continued)	
Field Heading	Meaning
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 25 on page 79](#) or in [Table 23 on page 77](#).

Cursor-sensitive control

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

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

DSD - Detailed Storage Delays Report

Figure 48 on page 85 shows a modified version of the Storage Delays report that replaces **Working Set Central** and **Expanded** with three columns: **VIO**, **XMEM** and **HIPR**. On the Storage Delays report, this information is combined and shown in the **OTHR** column.

How to request this report

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

```
DSD
```

Contents of the report

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

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

The COMM and LOCL fields include shared storage paging.

```

RMF 3.1  Storage Delays
Line 1 of 206
Command ==> Scroll ==> HALF
Samples: 100 System: MVS1 Date: 04/25/2023 Time: 10:31:40 Range: 100 Sec

Jobname  C  Service  DLY  ----- % Delayed for -----
          C  Class   %    COMM  LOCL  SWAP  OTR  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 48. 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 69 on page 160](#)) except for the XMEM and HIPR fields.

Table 28. Fields in the DSD Report	
Field Heading	Meaning
Delayed for XMEM	This column contains the paging delays from cross memory address spaces.
Delayed for HIPR	This column contains the paging delays from standard hiperspaces (including waits during scroll wait), but not ESO hiperspaces.

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

DSND - Data Set Delays Report

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

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

RMF sorts the data sets by descending overall delay percentages.

You can use this report as base for further analysis:

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

How to request this report

To request the Data Set Delays report, select **3** from the Primary Menu, and then select **3A** on the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)). or enter the following command:

```
DSND [dsname]
```

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

Special considerations

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

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

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

Contents of the report

```

RMF 3.1  Data Set Delays                               Line 1 of 6
Command ==>                                           Scroll ==> HALF
Samples: 100      System: MVS1  Date: 04/25/2023  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  B0ECL2  003C   20    15
                                           BGGEETEO 0201   15    10
BDA.CTT.MSPCT.SP41XCTT.GUIDE              EDSS99  B0ECL2  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 49. DSND Report

There is no graphic version of this report available.

Field descriptions

Table 29. Fields in the DSND Report	
Field Heading	Meaning
Input Data Set Name	Name of the data set or group of data sets to be reported on. This is the name which has been specified as command parameter or has been selected via report options.
Data Set Name	Name of a data set which was utilized during the report interval and selected for reporting. RMF lists the data set names by descending overall delay percentages.
Volume	Name of the volume on which the data set resides.

Table 29. Fields in the DSND Report (continued)	
Field Heading	Meaning
Jobname	Name of a job using or being delayed by the data set. RMF lists all jobs for each data set by descending delay percentages.
ASID	Hexadecimal address space identifier (ASID) of the job using the data set or waiting for its availability.
DUSG%	Percentage of time when the job has had an I/O request accepted by the channel for the volume on which the data set resides, but the request is not yet complete. Note: See "Special considerations" on page 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 report interval. You can select a data set name by placing **S** in front of the displayed data set name. The selected name is re-displayed in the header field "Selected Data Set Name". This field is an input field which can be used also to enter a data set name or a group of data set names using a "wild card", for example: BDA.CTT*, directly. The wild card example **BDA.CTT*** lets RMF select all data set names which start with the character sequence **BDA.CTT**. If someone specifies only "*", RMF reports on all data sets which are being utilized in the report interval.

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

```

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

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

Selected Data Set Name: BDA.CTT*

Sel  Data Set Name
-    BDA.CTT.MSPCT.SP41XCTT.CTTGUIDE.BOOK
-    BDA.CTT.MSPCT.SP41XCTT.GUIDE
-    BDA.CTT0.INFORM.SEQ
-    BDA.CTTX.TEST
-    BHEW.DATA.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 50. DSND Report Options Panel

Selected Data Set Name

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

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

Sel

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

Data Set Name

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

Note:

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

DSNJ - Data Set Delays - Job Report

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

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

You can use this report as base for further analysis:

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

How to request this report

To request the Data Set Delays - Job report, select **2** from the Primary Menu, and then select **1A** on the Job Report menu (shown in [Figure 7](#) on page 24). or enter the following command:

```
DSNJ [jobname]
```

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

Contents of the report

```

RMF 3.1  Data Set Delays - Job                               Line 1 of 5
Command ==>                                                Scroll ==> HALF

Samples: 100      System: MVS1  Date: 04/25/2023  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 00312    20    15
      BSHR.FIX.LINKLIB                DATA68 00257     2    10
      BDA.CTT.MSPCT.SP41XCTT.GUIDE      EDSS99 00312    13    10
      RMF.R430.NLS.OLDENG.ERBCOPS3.SEQ EDSS09 00312     3     8
      -- N/A --                      ----- 1     7
      BHEW.DATA.LST90514                DATA38 00122     1     6

```

Figure 51. DSNJ Report

There is no graphic version of this report available.

Field descriptions

Table 30. 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.

Table 30. Fields in the DSNJ Report (continued)	
Field Heading	Meaning
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, for example: • I/Os to system data sets (such as 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 report interval.
Num	The five-digit device number where the volume is mounted. The first digit represents the ID of the subchannel to which the I/O device is physically configured.
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 86.
DDL%	Percentage of time when the job was waiting to use the data set because of contention for the volume where the data set resides.

DSNV - Data Set Delays - Volume Report

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

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

RMF sorts the data sets by descending overall delay percentages.

You can use this report as base for further analysis:

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

How to request this report

To request the Data Set Delays - Volume report, select **3** from the Primary Menu, and then select **3B** on the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)). or enter the following command:

```
DSNV [volser]
```

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

Contents of the report

```

Command ==>
RMF 3.1 Data Set Delays - Volume
Line 1 of 10
Scroll ==> HALF

Samples: 100 System: MVS1 Date: 04/25/2023 Time: 10.03.20 Range: 100 Sec

----- Volume EDSS99 Device Data -----
Number: 00B4A Active: 84% Pending: 22% Average Users
Device: 3380A Connect: 10% Delay DB: 22% Delayed
Shared: Yes Disconnect: 52% Delay CM: 0% 0.4
PAV: 1.6H

----- Data Set Name ----- Jobname ASID DUSG% DDLY%
BDA.CTT.MSPCT.SP41XCTT.CTTGUIDE.BOOK BOECL2 003C 20 15
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 52. DSNV Report

There is no graphic version of this report available.

Field descriptions

Table 31. Fields in the DSNV Report	
Field Heading	Meaning
Device Data Section	This sections contains identical information as provided in the Device Delays variation of the Job Delay report (see “Device Delay variation” on page 122).
Data Set Name	Name of a data set which was utilized during the report interval and resides on the selected volume. RMF lists the data set names by descending overall delay percentages.
Jobname	Name of a job using or being delayed by the data set. RMF lists all jobs for each data set by descending delay percentages. 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, for example: • I/Os to system data sets (such as 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 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

```

Command ==>
RMF DSNV Report Options
Line 1 of 2
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 53. DSNV Report Options Panel

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

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

VOLSER

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

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

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

Available DASD Volumes

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

Note:

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

EADM - Extended Asynchronous Data Mover (EADM) Activity Report

The Extended Asynchronous Data Mover (EADM) activity report can be used to investigate performance problems that are related to the extended asynchronous data mover facility.

How to request this report

To request the EADM Activity Report, select **3** from the Primary Menu, then select **15** from the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)) or enter the following command:

```
EADM
```

Note: This report was formerly called SCM Activity Report. The **SCM** command is still accepted and has the same meaning as **EADM**.

Contents of the report

The extended asynchronous data mover (EADM) activity report provides these types of information:

EADM level information

The EADM (extended asynchronous data mover) summary section at the top of the report provides the rate of start subchannel (SSCH) instructions for all EADM devices together with response time statistics consisting of pending, IOP queue and initial command response time.

The values cover Storage Class Memory (SCM) activity as well as EADM compression and decompression activity.

Furthermore, the section provides request rates, throughput, and ratios of compression and decompression.

The values related to compression and decompression cover asynchronous compression/decompression activity by EADM. Synchronous compression/decompression activity is not reported.

Flash Express card level information

For each Flash Express card, the report provides measurements at both the LPAR and CPC level. The rate at which internal requests are processed by the adapter card, the rate at which data units were read and written, the average response and IOP queue time is displayed.

Note: If the hardware supports Virtual Flash Memory, Flash Express cards are simulated by cache and SCM activity is reported in one report line.

Figure 54 on page 92 shows an example of the EADM Activity Report.

Command ==>		RMF 3.1 EADM Activity		Line 1 of 1	
				Scroll ==> CSR	
Samples: 60	System: SYSF	Date: 04/25/2023	Time: 13.54.00	Range: 60	Sec
----- EADM Summary -----					
SSCH Total	SSCH Rate	PEND Time	IOPQ Time	ICMR Time	
0	0.00	0.000	0.000	0.000	
Compress:	Rate	Throughput	Ratio	Decompress: Rate	Throughput
	12.30	65321	45.78	23.17	43216
					0.67
Card ID	Util(%)	Read(B/s)	Write(B/s)	Req Rate	Resp Time
	Part Total	Part Total	Part Total	Part Total	Part Total
VFM	0.00 0.00	0.00 0.00	0.00 0.00	0.00 0.00	0.000 0.000
					0.000

Figure 54. EADM Activity Report

Table 32. Fields in the EADM Activity Report	
Field heading	Meaning
EADM summary This section provides summary information about the extended asynchronous data mover (EADM) devices or subchannels. EADM subchannels are similar to I/O subchannels in a way that I/O instructions can be issued. However, they do not have channel paths or device numbers assigned, and they are not defined in the I/O configuration. They are created automatically during IPL.	
SSCH Total	The total number of SSCH instructions to all EADM devices in the report interval.
SSCH Rate	The number of SSCH instructions to all EADM devices per second.
PEND Time	The average function pending time across all EADM devices in milliseconds. This is similar to function pending time for traditional I/O devices, which is the amount of time between when the SSCH is issued and the first command in the channel program is accepted. $\text{PEND} = \frac{\text{Sum(Function Pending Time)}}{\text{Measurement Event Count}}$
IOPQ Time	The average IOP queue time across all EADM devices in milliseconds. This is unique to EADM devices. It represents the amount of time the request is not accepted by the adapter because it would exceed its maximum capacity. For a particular I/O request, this may occur multiple times. $\text{IOPQ} = \frac{\text{Sum(IOP Queue Time)}}{\text{Measurement Event Count}}$

Table 32. Fields in the EADM Activity Report (continued)	
Field heading	Meaning
ICMR Time	The average initial command response time across all EADM devices in milliseconds. This is the time from when the first command does not immediately proceed to execute until the successful start of execution at the SCM resource part. $\text{ICMR} = \frac{\text{Sum(Initial Command Response Time)}}{\text{Measurement Event Count}}$
Compress: Rate	The number of compression requests per second.
Compress: Throughput	The number of bytes compressed per second.
Compress: Ratio	The ratio between input bytes compressed and output bytes compressed within this interval.
Decompress: Rate	The number of decompression requests per second.
Decompress: Throughput	The number of bytes decompressed per second.
Decompress: Ratio	The ratio between input bytes decompressed and output bytes decompressed within this interval.
Flash adapter measurements	
Card ID	The identifier of the flash adapter card. <i>VFM</i> is reported if the hardware has configured Virtual Flash Memory.
Following fields are displayed at a system-wide level (Total) and for the current LPAR (Part) whereby IOPQ Time is only available at the total level.	
Util(%)	The average utilization of the flash card during the interval as reported by the SCM measurement facility. The average utilization of Virtual Flash Memory is reported as the percentage of the time spent on System Assist Processors (SAP) for SCM processing compared to the total available SAP time in this reporting interval.
Read(B/s)	Bytes read per second.
Write(B/s)	Bytes written per second.
Req Rate	The requests processed per second.
Resp Time	The average response time per request in milliseconds. The response time represents the CHSC execution time and does not include pending, IOP queue and initial command response time.
IOPQ Time	The average IOP queue time per request in milliseconds.

ENCLAVE - Enclave Report

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

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

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

How to request this report

To request the ENCLAVE report, select **1** on the Primary Menu, and then **6** on the Overview Report Selection Menu (shown in [Figure 6 on page 23](#)), or enter the following command:

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

Contents of the report

```

RMF 3.1  Enclave Report -
Command ==>                                     Line 1 of 13
                                           Scroll ==> HALF
Samples: 120      System: SYSE  Date: 11/30/22  Time: 11.44.00  Range: 120  Sec
Current options:  Subsystem Type: ALL          -- CPU Util --
                  Enclave Owner:              Appl%   EAppl%
                  Class/Group:                19.5    19.9

Enclave  Attribute  CLS/GRP  P Goal  % D X  EAppl%  TCPU  USG  DLY  IDL
*SUMMARY
ENC00003  A          DISCRETN 1          Y      0.732
          BRMFSEZ4  0.665  40.31  3.3  51  0.0
          SRVUSER
ENC00004  -----  SYSSTC  1 N/A    Y      0.042  5.105  0.0  0.0  0.0
          RMFGAT
          STCUSER
ENC00001  A          DISCRETN 1          Y      0.014  7.308  0.8  60  0.0
          BRMFSEZ8  0.012  7.180  0.0  66  0.0
          SRVUSER
ENC00002  A          DISCRETN 1          Y      0.012  7.180  0.0  66  0.0
          BRMFSEZ6  0.012  7.180  0.0  66  0.0
          SRVUSER

```

Figure 55. ENCLAVE Report

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

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

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

Field descriptions

Table 33. Fields in the ENCLAVE Report	
Field Heading	Meaning
Subheader Section - You can define the setting of each field in the Enclave Report Options panel.	
Subsystem Type	Reporting only on enclaves that belong to this subsystem type, for example, Db2 or DDF.
Enclave Owner	Reporting only on enclaves that are owned by the address space with this jobname.
Class/Group	Reporting only on enclaves that run in this service class.
Appl%	Percentage of the maximum general purpose processor capacity used by all address spaces during the report interval. This value is divided by the number of logical processors or cores that have been active during this interval.

Table 33. Fields in the ENCLAVE Report (continued)	
Field Heading	Meaning
EAppl%	Percentage of the maximum general purpose processor capacity used by all address spaces and enclaves during the report interval. This value is divided by the number of logical processors or cores that have been active during this interval.
Enclave Identification.	
Enclave	Generated name to allow association of an enclave with instances shown on other Monitor III reports. *SUMMARY is shown in the summary line that totals up the CPU time for the reported enclaves.
Attribute	Dynamic list of attributes. The reporter lists the attributes (maximal eight characters) in the order at which they are specified in the Enclave Report Options panel.
CLS/GRP	Service class the enclave is associated with.
P	Service class period the enclave is currently running in.
Goal	Response time goal (in seconds, minutes (M) or hours (H)).
%	Response time percentile or velocity.
D	Dependent enclave indication. 'Y' if the enclave is an extension of an address space transaction, otherwise blank.
X	Multi-system Indicator This column gives an indication about the origin of the enclave: 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 the maximum general purpose processor capacity consumed by the individual enclave or by all reported enclaves (in the Monitor III range).
TCPU	Total CPU time (in seconds) consumed by the enclave (see 'Detailed Performance Statistics').
USG	Percentage of total USING samples (use samples for CPU and I/O), based on #STS (total number of state samples in the enclave).
DLY	Percentage of total DELAY samples (delay samples for CPU, I/O, capping, storage, queuing) based on #STS.
IDL	Percentage of idle samples based on #STS.

Cursor-sensitive control on the Enclave Report

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

- [“Enclave Details” on page 95](#)
- [“Enclave Classification Attributes” on page 97](#)

Enclave Details

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

RMF Enclave Details

Details for enclave ENC00003 with token 00000128 0000000A
Press Enter to return to the Report panel.

```

- CPU Time --      - zCBP Time -      - zIIP Time -
Total    40.31      Total    40.31      Total    0.000
Delta    5.589      Delta    5.589      Delta    0.000

State  ---- Using ---- ----- Delay ----- IDL  UNK
Samples CPU CBP IIP I/O  CPU CBP IIP I/O STO CAP QUE
    123   0.0 3.3 0.0 0.0   0.0 3.3 0.0 0.0 0.0 48 0.0  0.0  46

```

Figure 56. ENCLAVE Report - Enclave Details

Table 34. Fields in the ENCLAVE Report - Enclave Details	
Field Heading	Meaning
CPU Time	Total Total CPU time (in seconds) consumed by the enclave on general purpose processors and special purpose processors. Delta CPU time (in seconds) consumed by the enclave on general purpose processors and special purpose processors in the reported Monitor III range.
zCBP Time	Total Total CPU time (in seconds) consumed by the enclave on zCBPs. Delta CPU time (in seconds) consumed by the enclave on zCBPs in the reported Monitor III range.
zIIP Time	Total Total CPU time (in seconds) consumed by the enclave on zIIPs. Delta CPU time (in seconds) consumed by the enclave on zIIPs in the reported Monitor III range.
State Samples	Total number of state samples in the enclave.
Using% and Execution Delays% - In contrast to other Monitor III fields, these states shown in the pop-up panel are multistate. This means, they reflect the real amount of work executing in the enclave. All percentages are based on the number of state samples.	
Using%	Percentage of: CPU CPU using samples CBP zCBP using samples IIP zIIP using samples I/O I/O using samples

Table 34. Fields in the ENCLAVE Report - Enclave Details (continued)	
Field Heading	Meaning
Execution Delays%	Percentage of: CPU CPU delay samples CBP zCBP delay samples IIP zIIP delay samples I/O I/O delay samples STO 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 CAP CPU capping samples QUE Queue delay samples
IDL	Percentage of idle samples.
UNK	Percentage of unknown samples.

Enclave Classification Attributes

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

```

RMF Enclave Classification Attributes

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

Subsystem Type: DDF      Owner: ENC00003      System: RMF3      More:      +
Accounting Information . . . :
    Q123ERF7

Collection Name . . . . . : COLLECTION
Connection Type . . . . . :
Correlation Information . . : CTT
LU Name . . . . . :
Net ID . . . . . :
Plan Name . . . . . : TEST
Priority . . . . . :
Process Name . . . . . :
Transaction/Job Class . . . : JES3
...

```

Figure 57. ENCLAVE Report - Enclave Classification Attributes (1)

Report options

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

- An enclave filter by one of the following criteria:
 - Subsystem type, for example DDF, IWEB, or SOM
 - Enclave owner job name, for example DB2MSTR
 - Service class
 - Performance group
- A list of classification attributes.

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

```

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

Select one of the following options:

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

Select (S) one or more classification attributes:

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

Figure 58. ENCLAVE Report Options

Subsystem Type

Report on enclaves that belong to this subsystem type, for example, Db2 or DDF.

ALL selects all active subsystems.

Enclave Owner

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

Service Class / Performance Group

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

Classification Attributes

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

ENQ - Enqueue Delays Report

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

How to request this report

To request the Enqueue Delays report, select **3** from the Primary Menu, and then select **4** on the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)), or enter the following command:

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

Contents of the report

```

RMF 3.1  ENQ Delays                               Line 1 of 14
Command ==>                                         Scroll ==> HALF
Samples: 100      System: MVS1  Date: 04/25/2023  Time: 10.03.20  Range: 100  Sec
Jobname  DLY  ----- Resource Waiting -----  --- Holding ---
      %    %  STAT Major/Minor Names (Scope)      %  Name/SYS  STAT
SPEWAK2 100  100  SW  SYSDSN          (SYS)          100 AMOLLOY   SO
                        SYS1.NUCLEUS                100 SCHMATE   SO
                                                100 DRAGON    SO
                                                100 SCHUMAC   SO
                                                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  SYSZVVDV        (SYS)           1 STEVEB    EO
+CATALOG                        SYS1.MVS3.MCAT
*MASTER*  1    1  EW  SYSIKJBC        (SYSS)          1 *MASTER*  EO
                        PETEG

```

Figure 59. ENQ Report

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

Field descriptions

Table 35. 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: $\text{DLY \%} = \frac{\text{Delay samples}}{\text{\# Samples}} * 100$ 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.

Table 35. Fields in the ENQ Report (continued)	
Field Heading	Meaning
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: $\text{Waiting \%} = \frac{\text{Delay samples}}{\text{\# Samples}} * 100$ 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.

Table 35. Fields in the ENQ Report (continued)

Field Heading	Meaning
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. 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 CHPNT 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 SYSMF01 SMF SYS1.MANx data set SYSVSAM VSAM data sets SYSVTOC VTOC (Minor name is volser) SYSAVM AVM queue or data areas SYSZBDT z/OS bulk data transfer (Minor name is node name) SYSZCAXW Catalog auxiliary work area SYSZCMD5 Master trace command or Message loss detection SYSZCOMM Global Resource Serialization ring processing table SYSZCSD CSD control block field

Table 35. Fields in the ENQ Report (continued)

Field Heading	Meaning
Major/Minor Names continued	Major Name Resources SYSZEC16 Purge data set SYSZIGGI TSB (Minor name is ASID) SYSZISTOC 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)
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: $\text{Holding \%} = \frac{\text{Holding samples}}{\text{\# Samples}} * 100$ 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.

Table 35. Fields in the ENQ Report (continued)	
Field Heading	Meaning
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 43 on page 75 for an example. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed. See Figure 41 on page 73 for an example.

ENQR - Enqueue Resource Delays Report

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

How to request this report

To request the ENQR report, select **3**, and then select **5** on the Resource Report Selection Menu (shown in Figure 8 on page 24), or enter the following command:

```
ENQR [resourcename]
```

Contents of the report

Command ==>		RMF 3.1 ENQ Resource Delays		Line 1 of 14	
				Scroll ==> HALF	
Samples: 100	System: MVS1	Date: 04/25/2023	Time: 10.03.20	Range: 100	Sec
Major/Minor	Resource Name (Scope)	Delayed %	Name STAT	Holding %	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					
SYSZVDS	(SYS)	1	LUCKYSM EW	1	STEVEB EO
SYS1.MVS3.MCAT			+CATALOG		+CATALOG

Figure 60. ENQR Report

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

Field descriptions

Table 36. 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: $\text{Delayed \%} = \frac{\text{\# Delay Samples}}{\text{\# Samples}} * 100$ 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: $\text{Holding \%} = \frac{\text{\# Holding Samples}}{\text{\# Samples}} * 100$ 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 61. ENQR Report Options Panel

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

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

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

Major

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

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

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

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

Available ENQ Major Names

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

GROUP - Group Response Time Report

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

How to request this report

To request the Group Response Time report, select **1** on the Primary Menu, and then select **5** on the Overview Report menu (shown in [Figure 6 on page 23](#)), or enter one of the following commands:

```
GROUP service_class,period
```

```
GROUP report_class,period
```

Parameter report_class designates either a report class or tenant report class.

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

```
GROUP HOTBATCH,1
```

Contents of the report




Figure 62. GROUP Report

There is no graphic version of this report available.

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




Figure 63. Response time Components data (pop-up panel)

Field descriptions

Table 37. Fields in the GROUP Report	
Field Heading	Meaning
Class	The name of the service or report class.
Period	The period number.
Description	The description of the specified class, it is derived from the service policy.

Table 37. 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 216 for more details.
Frames %ACT	The percentage of central storage frames used by active users in the specified class during the report interval. For a definition of active users, see Users/Active under “WFEX - Workflow/Exceptions Report” on page 216. RMF accumulates the number of central storage frames for all active users during the report interval, then calculates the percentage as follows: $\text{Frames \%ACT} = \frac{\text{ACSF}}{\text{OCSF}} * 100$ ACSF Accumulated central storage frames OCSF Online central 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 37. 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 (in milliseconds) 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 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 37. 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 delay variation of the Job Delay 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 69 on page 160 for more information about the storage delays.

Table 37. 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: TO Terminal output wait TI Terminal input wait LW Long wait XS Auxiliary storage shortage RS Real storage shortage DW Detected long wait MP Memory Pool shortage 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 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 37. 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 delay variation of the Job Delay 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 69 on page 160 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 for Average Users OUTR Swap Reason for a list of possible 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.

Report options




Figure 64. GROUP Report Options Panel

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

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

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

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

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

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

HSM - Hierarchical Storage Manager Delays Report

The Hierarchical Storage Manager (HSM) Delays report allows you to investigate situations where jobs are delayed when requesting service from HSM.

RMF lists all jobs delayed during the refresh period in order by descending delay percentage.

How to request this report

To request the HSM report, select **4** on the Primary menu, and then select **1** on the Subsystem Report menu (shown in Figure 9 on page 25), or enter the following command:

```
HSM [job_class,service_class]
```

Contents of the report

```

Command ==>
RMF 3.1  HSM Delays
Line 1 of 3
Scroll ==> HALF

Samples: 100      System: MVS1  Date: 04/25/2023  Time: 10.03.20  Range: 100  Sec

Jobname    DLY  ----- Main Delay Reason -----
           %    % F-Code Explanation
AUDTRPTZ   94   94   3  Dataset recall from auxiliary storage.
APETER     82   82   3  Dataset recall from auxiliary storage.
TJSMITH    77   77   3  Dataset recall from auxiliary storage.

```

Figure 65. HSM Report

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

Field descriptions

Table 38. 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: $\text{DLY \%} = \frac{\text{\# Delay Samples}}{\text{\# Samples}} * 100$ 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.

Table 38. Fields in the HSM Report (continued)

Field Heading	Meaning																
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: $\% = \frac{\text{\# Delay Samples}}{\text{\# Samples}} * 100$ 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 DLY % value of the job if there is no overlap in delay states and there are no more than two function codes responsible for the delay. The HSM F-codes (in decimal) and their explanations are as follows: <table> <tr> <th>F-Code</th><th>Explanation</th></tr> <tr> <td>03</td><td>A data set is being recalled from auxiliary storage.</td></tr> <tr> <td>05</td><td>A data set is being recovered.</td></tr> <tr> <td>06</td><td>A data set is being migrated.</td></tr> <tr> <td>07</td><td>A data set is being backed up.</td></tr> <tr> <td>08</td><td>A control data set record is being read.</td></tr> <tr> <td>08</td><td>A JES3 C/I locate is being done.</td></tr> <tr> <td>12</td><td>A data set is being deleted.</td></tr> </table>	F-Code	Explanation	03	A data set is being recalled from auxiliary storage.	05	A data set is being recovered.	06	A data set is being migrated.	07	A data set is being backed up.	08	A control data set record is being read.	08	A JES3 C/I locate is being done.	12	A data set is being deleted.
F-Code	Explanation																
03	A data set is being recalled from auxiliary storage.																
05	A data set is being recovered.																
06	A data set is being migrated.																
07	A data set is being backed up.																
08	A control data set record is being read.																
08	A JES3 C/I locate is being done.																
12	A data set is being deleted.																

Report options

The HSM Report Options panel is similar to the Device Report Options panel. See [Figure 43 on page 75](#) for an example. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed. See [Figure 41 on page 73](#) for an example.

IOQUEUE - I/O Queuing Activity Report

The I/O Queuing Activity report (IOQUEUE) provides information, grouped by LCU (logical control unit), on the I/O configuration. The information includes contention rate, queue lengths, and percentages of time when one or more I/O components were busy. Information about the LCU is useful because the LCU is the focus of I/O configuration and path management measurements for a related group of I/O devices.

For all channels that are managed by **Dynamic Channel Path Management (DCM)**, additional information is available. DCM allows an installation to identify channels which they wish to be managed dynamically. These channels are not assigned permanently to a specific control unit, but belong to a pool of channels. Based on workload requirements in the system, these channels are assigned dynamically by DCM. For each LCU with DCM managed channels, a summary line displays the minimum and maximum number of connected DCM managed channels, the number of defined DCM managed channels and accumulated activity data.

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

This report can tell you about the cause of performance problems associated with channel paths and devices. You could, for example, find the reason for an unusually long pending time reported on the device report. Check the relationship between the percentage of requests deferred for device busy and control unit busy for the LCU on the I/O Queuing Activity report.

How to request this report

To request the I/O Queuing Activity report, select **3** from the Primary Menu, and then select **13** on the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)), or enter the following command:

```
IOQUEUE
```

Special considerations of report output

In a VM guest system environment, the report for a z/OS system that is authorized with the VM RMCHINFO directory option, shows static configuration data, only. Measurement data is not available.

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

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

Contents of the report

RMF 3.1 I/O Queuing Activity										Line 1 of 54						
Command ==>										Scroll ==> HALF						
Samples: 30		System: S5C		Date: 04/25/2023		Time: 03.23.30		Range:		30 Sec						
Path	DCM	CTL	Units	DCM Group		LCU	Cont Rate	Del Lngth	Q CSS	CHPID Taken	%DP Busy	%CU Busy	AVG CUB	AVG CMR		
D7		5F00				0048				1.13	0.0	0.0	0.0	0.2		
D6		5F00				0048				0.97	0.0	0.0	0.0	0.2		
						0048	0.0	0.00	0.3	2.10	0.0	0.0	0.0	0.2		
B0 PF		8000				0069				82.17	0.0	0.0	0.0	0.2		
B1 PF		8000				0069				83.83	0.0	0.0	0.0	0.2		
B2 NP		8000				0069				0.00	0.0	0.0	---	---		
B3 NP		8000				0069				0.00	0.0	0.0	---	---		
95 PF		8000				0069				83.17	0.0	0.0	0.0	0.2		
						0069	0.0	0.00	0.4	249.17	0.0	0.0	0.0	0.2		
B0 NP		8100				006A				0.00	0.0	0.0	---	---		
B1 NP		8100				006A				0.00	0.0	0.0	---	---		
B2 PF		8100				006A				124.53	0.0	0.0	0.0	0.2		
B3 PF		8100				006A				124.87	0.0	0.0	0.0	0.2		

Figure 66. IOQUEUE Report

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

Field descriptions

Table 39. 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. If applicable, the path attribute is indicated with the CHPID: PF preferred path NP non-preferred path NS path attribute not specified For devices residing in control units that do not support path attributes, only the CHPID is displayed.
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.
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. $\text{Cont Rate} = \frac{\# \text{ Enqueued Requests}}{\text{Range Time}}$
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. $\text{Del Q Lngth} = \frac{\text{Accumulated Queue Length} - \# \text{ Enqueued Requests}}{\# \text{ Enqueued Requests}}$
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. $\text{AVG CSS} = \frac{\text{Channel Subsystem Time}}{\# \text{ I/O Operations Accepted}}$

Table 39. Fields in the IOQUEUE Report (continued)	
Field Heading	Meaning
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. $\text{CHPID Taken} = \frac{\# \text{ I/O Operations Accepted on that Path}}{\text{Range Time}}$
% 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. $\% \text{ DP Busy} = \frac{\text{DPB}}{\text{DPB} + \text{CUB} + \text{SUC}} * 100$ 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
% 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. $\% \text{ CU Busy} = \frac{\text{CUB}}{\text{DPB} + \text{CUB} + \text{SUC}} * 100$ 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. $\text{AVG CUB} = \frac{\text{Control Unit Busy Time}}{\# \text{ I/O Operations Accepted on that Path}}$
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). $\text{AVG CMR} = \frac{\text{Initial Command Response Time}}{\# \text{ I/O Operations Accepted on that Path}}$

JES - Job Entry Subsystem Delays Report

The JES Delays report allows you to investigate situations where executing jobs are delayed when requesting service from JES. RMF lists all jobs delayed during the report interval in descending delay percentages.

How to request this report

To request the JES report, select **4** on the Primary menu, and then select **2** on the Subsystem Report menu (shown in [Figure 9 on page 25](#)), or enter the following command:

```
JES [job_class,service_class]
```

Contents of the report




Figure 67. JES Delays report

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

Field descriptions

Table 40. 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 the description of the Main Delay Reason field for valid JES function codes. Note: This DLY% value is also found in the JES field on the Job Delay report.

Table 40. Fields in the JES Report (continued)

Field Heading	Meaning
Main Delay Reason	The subsystem function code that indicates the main reason for the delay. RMF reports the function code with the highest count as the main delay reason. The % column indicates how much of the JES delay of the job is caused by the reported subsystem function. This value is calculated as follows: $\% = \frac{\text{\# Delay Samples}}{\text{\# Samples}} * 100$ 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 DLY % field in this report description. Note: The Main Delay Reason % values add up to the DLY % value of the job if there is no overlap in the delay states and there are no more than two function codes responsible for the delay. The JES F-codes (function codes in decimal) and their explanations are: F-Code Explanation 1 Processing TSO OUTPUT command request 2 Waiting for JES to cancel a job 3 Waiting for job status information 9 Waiting for WTO/WTOR request 12 Waiting for job termination 13 Waiting for JES to restart a job 23 Waiting for dynamic alloc via SSOBDYCD 23 Waiting for SETUP request 26 Waiting for change DD name 27 Waiting for change ENQ use attribute 75 Processing notify user 132 Waiting for JDS access 138 Validating SYSOUT destination 138 Waiting for JES to cancel a job 138 Waiting for job status information 141 Waiting for ENDREQ 144 Processing TSO OUTPUT command request 149 Dynamically allocating data set to JES3 151 Changing JES3 DD via dynamic allocation

Table 40. Fields in the JES Report (continued)

Field Heading	Meaning
Main Delay Reason continued	F-Code Explanation 152 Waiting for FSS request 153 Waiting for CI driver 157 Waiting for SYSOUT API to process request 158 Processing SSI Extended Status 162 Waiting for TCPIP NJE global services 163 Waiting for job class information 164 Waiting for initiator information 165 Waiting for NJE node information 166 Waiting for spool partition information 167 Waiting for JESPLEX information 168 Waiting for SJF services 169 Waiting for JES device information 255 Waiting for SPOOL space

Report options

The JES Report Options panel is similar to the Device Report Options panel. See Figure 43 on page 75 for an example. Select YES for Jobs on the Report Options panel to display the Job Selection/Exclusion panel. See Figure 41 on page 73 for an example.

JOB - Job Delay Report

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

How to request this report

To request the Job Delay report, select **2** from the Primary Menu and then select **5** from the Job Report Selection Menu (shown in Figure 7 on page 24), or enter the following command:

```
JOB jobname
```

Contents of the report

The Job Delay report is available in the following delay variations:

- Device delay
- Enqueue delay
- HSM address space delay

- JES address space delay
- Operator message delay
- Operator tape mount delay
- Processor delay
- Storage delay
- Quiesce delay
- XCF address space delay

The following sections explain each Job Delay report variation. The Job Delay report is divided into three sections. Each variation has the same fields in the top and bottom sections. The middle section varies depending on the type of delay being reported. Only the middle section will be described for each variation.

Delay can be either primary delay or requested delay. Primary delay shows information about the type of delay contributing most to the overall delay of a job. Requested delay shows information about a type of delay that you select. Use the Job Report Selection Menu, or cursor-sensitive control to select the type of delay to be reported. See [“The Job Report Selection Menu” on page 23](#) for more information about requesting delay types.

Displaying the Job Delays report

You can display the Job Delays report in several ways:

- Place the cursor on a jobname in any RMF report and press ENTER.
- Place the cursor on a % Delayed for field in a delay report or a Job Delay report and press ENTER.
- Issue the JOB <jobname> command on any command line.
- Use the Job Report Selection Menu to specify a particular job and the type of delay to be reported.

The highlighted line in the bottom section of the Job Delays report identifies the job which is analyzed in the top and middle sections.

Since more than one job with the same jobname may be in the system during the report interval, the address space identifier (ASID) is displayed instead of the jobname. To display a different job with the same name, place the cursor on the appropriate ASID field in the bottom section and press ENTER.

The top and bottom parts of the Job Delays report

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

```

RMF 3.1  Job Delays                               Line 1 of 1
Command ==>                                       Scroll ==> HALF

Samples: 100      System: MVS1  Date: 04/25/2023  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 68. Top Part of Job Delays report

Field descriptions

Table 41. Fields in the Job Delays report	
Field Heading	Meaning
Job:	Name of delayed job.

Table 41. Fields in the Job Delays report (continued)	
Field Heading	Meaning
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 69. Bottom Part of Job Delay report

The bottom part of the Job Delay report has similar fields as the Delay report. Similar fields also appear on the Delay report. Table 20 on page 67 describes these fields. In the Job Delay report, the address space identifier (ASID) is displayed instead of the jobname since more than one job with the same jobname may have been in the system during the report interval. The Using% field shows the percentage of time the job was using a processor (PRC) and the percentage of time the job was using a device (DEV).

Device Delay variation

This variation of the Job Delay report shows a job delayed by a device.

How to request this variation

To request the Device Delay variation, select **2** from the Primary Menu, and then select **1** on the Job Report Selection Menu (shown in Figure 7 on page 24), or enter the following command:

```
DEVJ jobname
```

```

----- Volume BPXLK1 Device Data -----
Number:    00409      Active:    62%      Pending:    3%      Average Users
Device:    3390A      Connect:   17%      Delay DB:    1%      Delayed
Shared:    Yes       Disconnect: 42%      Delay CM:    0%      0.7
PAV:       1.6H

```

Figure 70. Device Delay variation of the Job Delay report

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

Field descriptions

Table 42. Fields in the Device Delay variation of the Job Delay report	
Field Heading	Meaning
Number:	The five-digit device number where the volume is mounted. The first digit represents the ID of the subchannel to which the I/O device is physically configured.
Device:	The device type.
Shared:	Indicates if the device is shared.

Table 42. Fields in the Device Delay variation of the Job Delay report (continued)	
Field Heading	Meaning
PAV	The number of parallel access volumes (base and alias) which were available at the end of the report interval. If the number has changed during the report interval, it is followed by an '*'. If the device is a HyperPAV base device, the number is followed by an 'H'. The value is the average number of HyperPAV volumes (base and alias) for that range. $\text{Average \# of HPAV devices} = \frac{\text{Accumulated \# of HPAV devices}}{\text{Number of Samples}}$ This field appears only for parallel access volumes.
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 25 on page 79 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 25 on page 79 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 25 on page 79 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
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 25 on page 79 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 Job Delay report variation shows a job delayed by a serially reusable resource.

How to request this variation

To request the Enqueue Delay variation, select **2** from the Primary Menu, and then select **2** from the Job Report Selection Menu (shown in [Figure 7 on page 24](#)), or enter the following command:

```
ENQJ jobname
```

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

Figure 71. Enqueue Delay variation of the Job Delay report

The fields in the middle section of this report contain information about the three main holders of the resource for which this job is delayed.

Field descriptions

Table 43. Fields in the Enqueue Delay variation of the Job Delay report	
Field Heading	Meaning
Job:	Name of a job holding the resource. Up to three jobs can be displayed in this section.
Holding:	Indicates how much a specific job is contributing to the holding of the resource. See Table 36 on page 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 Delay report have the same format. They show a job delayed by either the HSM or JES address space.

How to request these variations

Select **2** from the Primary Menu, and then select **3** for HSM or **4** for JES from the Job Report Selection Menu (shown in [Figure 7 on page 24](#)) or enter one of the following commands:

```
HSMJ jobname
JESJ jobname
```

```
----- HSM Performance Summary -----
Job: DFHSM      Workflow:   50%      Primary delay category:   OPER
                  Using:      89%      Primary delay reason:      Mount
                  Delay:      72%      Primary delay percent:     66%
```

Figure 72. HSM Delay variation of Job Delay report

Field descriptions

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

Table 44. Fields in the HSM/JES delay variation of the Job Delay report	
Field Heading	Meaning
Job:	Name of the JES/HSM address space.

Table 44. Fields in the HSM/JES delay variation of the Job Delay report (continued)	
Field Heading	Meaning
Workflow:	The workflow percentage of the JES/HSM address space. See “Address space workflow (%)” on page 11 for the calculation.
Using:	The using percentage of the JES/HSM address space. See “Address space using (%)” on page 12 for the calculation.
Delay:	The delay percentage of the JES/HSM address space. See Table 20 on page 67 for the calculation.
Primary delay category:	The category of delay contributing most to the overall delay. See the % Delayed for field in Table 20 on page 67 for details.
Primary delay reason:	The contents of this field depend on the primary delay category. See the Primary Reason field in Table 20 on page 67 for an explanation.
Primary delay percent:	The percentage of delay for the primary delay category.

Operator Message and Mount Delay variations

The message and mount delay variations of the Job Delay report have the same format. They show a job delayed by either an operator message request or mount request.

How to request these variations

Select **2** from the Primary Menu, and then select **6** for the delay caused by volume mount or select **7** for the delay caused by operator reply from the Job Report Selection Menu (shown in Figure 7 on page 24) or enter one of the following commands:

```
MSGJ jobname
MNTJ jobname
```

Job: MISTYDFS Primary delay: Awaiting reply to operator request 53.

```
----- Job Performance Summary -----
      Service      WFL -Using%- DLY IDL UKN ---- % Delayed for ---- Primary
CX ASID Class    P Cr % PRC DEV % % % PRC DEV STR SUB OPR ENQ Reason
B 0167 NRPRIME  1   0  0  0 100  0  0  0  0  0  0 100  0 Message
```

Figure 73. Operator Message Delay variation of the Job Delay report

Field descriptions

Table 45. Fields in the Operator Message and Mount Delay variations of the Job Delay report	
Field Heading	Meaning
Job:	Name of delayed job.
Primary Delay: or Requested Delay:	A more detailed description of the reason stated in the Primary Reason field of the highlighted line in the bottom section of this report. Requested Delay appears instead of Primary Delay if you select operator delays as the type of delay to be reported.

Processor Delay variation

This Job Delay report variation shows a job delayed by a processor.

How to request this variation

To request the Processor Delay variation, select **2** from the Primary Menu, and then select **8** from the Job Report Selection Menu (shown in Figure 7 on page 24) or enter the following command using the format:

```
PROGJ 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 74. Processor Delay variation of the Job Delay report

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

Field descriptions

Table 46. Fields in the Processor Delay variation of the Job Delay report	
Field Heading	Meaning
Job:	Name of job holding the processor. Up to three jobs can be displayed in this section.
Holding:	The percentage of time the holding job used the processor while the delayed job was waiting for the processor.
PROC Using:	The percentage of time the holding job spent using the processor. See the USG % field in Table 61 on page 146 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 22 on page 74 for the calculation of this value.

If the sum of the PROC DLY% (not shown on this report) and PROC Using fields of the holding job is 100, the PROC and DEV Using fields in the middle section will be highlighted. This indicates the main job holding the processor may be looping. In this case, “Job may be looping” is listed as a probable cause.

Quiesce variation

How to request this variation

To request the Quiesce Delay variation, select **2** from the Primary Menu, and then select **9** from the Job Report Selection Menu (shown in [Figure 7 on page 24](#)) or enter the following command using the format:

```
QSCJ jobname
```

This report shows a job delayed because it was quiesced by the operator.

```

----- Job Performance Summary -----
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 75. Quiesce Delay variation of the Job Delay report

Field descriptions

Table 47. Fields in the Quiesce Delay variation of the Job Delay 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 Delay report shows a job delayed by contention of storage.

How to request this variation

To request the Storage Delay variation, select **2** from the Primary Menu, and then select **10** from the Job Report Selection Menu (shown in [Figure 7 on page 24](#)) or enter the following command using the format:

```
STORJ jobname
```

```
----- Job Storage Usage Data -----
Average Frames:  294      Working set:  341      Fixed Frames:   38
Active Frames:   249      Aux Slots:   2928     DIV Frames:    0
Idle Frames:     45      Page In Rate:  3.3     ES Move Rate:  12.4
```

Figure 76. Storage Delay variation of Job Delay report

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

Field descriptions

Table 48. Fields in the Storage Delays variation of the Job Delay report	
Field Heading	Meaning
Average Frames	The sum of active and idle frames.
Active Frames	The average number of frames held by the job while it was active. See Table 69 on page 160 for the calculation of this value.
Idle Frames	The average number of frames held by the job when it was idle. See Table 69 on page 160 for the calculation of this value.
Working Set	The average amount of storage a user occupied while in storage. See Table 69 on page 160 for the calculation of this value.
Aux Slots	The average number of auxiliary slots for each address space.
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 75 on page 168 for the calculation of this value.
DIV Frames	The number of central storage frames used by DIV. See Table 75 on page 168 for the calculation of this value.
ES Move Rate	The rate of pages moved from expanded storage. This includes both single and blocked pages; but does not include hiperspace or VIO pages.

XCF variation

The cross-system coupling facility (XCF) variation of the Job Delay report shows a job delayed by XCF.

How to request this variation

To request the XCF variation, select **2** from the Primary Menu, and then select **11** from the Job Report Selection Menu (shown in [Figure 7 on page 24](#)) or enter the following command:

```
XCFJ jobname
```

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

Figure 77. XCF Delay variation of Job Delay report

Field descriptions

Table 49. Fields in the XCF Delay variation of the Job Delay report	
Field Heading	Meaning
Delay:	The delay percentage of the XCF address space. See Table 20 on page 67 for the calculation.
Dev Number:	The path number corresponding to the delay percentage.

Monitor III Utility fields

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

Table 50. Additional Fields in the Job Delay report	
Field Heading	Meaning
Using percentage	The percentage of time a job was using the processor or a device.
JES delay percentage	The percentage of time a job was delayed when requesting service from JES.
HSM delay percentage	The percentage of time a job was delayed when requesting service from HSM.
XCF delay percentage	The percentage of time a job was delayed when requesting service from XCF.
Operator mount delay percentage	The percentage of time a job was delayed by an operator mount request.
Operator message delay percentage	The percentage of time a job was delayed by an operator message request.
Operator quiesce delay percentage	The percentage of time a job was delayed because the operator quiesced the address space.
WLM resource capping delay percentage	The percentage of time a job was delayed because • 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 <i>z/OS MVS Planning: Workload Management</i>).

Report options

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

```

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 78. Job Report Options Panel

RMF saves the Jobname you enter across sessions.

Jobname

The name of the job for which you want data reported in your job delay report.

Available Jobs

The list of jobs that were active during the report interval.

If the job you want is not listed, it had no activity during the current report interval. If you specify the job, it will appear on your Job Delay reports when it is available.

LOCKSP - Spin Lock Report

Through locking, the system serializes the use of system resources by authorized routines and, in a Parallel Sysplex, by processors. Lock holders can impede other work units that need the same lock and must wait until the lock holder releases the lock.

RMF reports about the various types of system resource locks in the **Spin Lock Report** described in this section and in the **Suspend Lock Report** described in [“LOCKSU - Suspend Lock Report” on page 130](#).

If a spin lock is unavailable, the requesting processor continues testing the lock until the other processor releases it (spinning). As soon as the lock is released, the requesting processor can obtain the lock and thus can obtain control of the protected resource.

No symptoms for delays due to locks are visible except excessive spinning. Therefore, RMF periodically checks all types of system resource locks.

How to request this report

To request the Spin Lock Report, select **1** from the Primary Menu and then select **12** on the Overview Report Selection Menu (shown in [Figure 6 on page 23](#)), or enter one of the following commands:

```
LOCKSP [HELD | SPIN | BOTH]
```

```
LSP [HELD | SPIN | BOTH]
```

Contents of the report

The **Spin Lock Report** provides information about how often a spin lock is held and about jobs that are spinning because of a lock request. It consists of two sections:

- The upper part (Held section) displays information about spin locks which have been observed as held, either exclusively (EXCL) or shared (SHR).
- The lower part (Spin section) displays spin locks and address spaces which are spinning due to a request for this lock.

```

RMF 3.1 Spin Lock Report
Command ==>
Line 1 of 10
Scroll ==> PAGE

Samples: 120 System: CB88 Date: 04/25/2023 Time: 09.21.00 Range: 120 Sec

Resource    Type/   CPUID/   Held   Spin
Jobname     ASID    Address  %      %
SRM         EXCL    07       0.83
DISP        EXCL    0E       0.41
SRM         EXCL    09       0.41
DISP        EXCL    11       0.41
DISP        EXCL    07       0.41
DISP        EXCL    04       0.41
VFIX        EXCL    0C       0.41
VFIX        EXCL    06       0.41
DISP        EXCL    01       0.21
XCFQ        SHR     06       0.21
IXLSHR      SHR     01       0.21
IXLSHR      SHR     0C       0.21
VFIX        RMFGAT  0196    0147A128 0.05
IXSH S/E    RMFGAT  0196    00FF14D4 0.05

```

Figure 79. Spin Lock Report

Table 51. Fields in the LSP Report	
Field Heading	Meaning
Resource	The resource name or the address of the spin lock.
Type/Jobname	Held section: The type of the lock (exclusively or shared). Spin section: The jobname (address space), which is spinning due to the lock request.
CPUID/ASID	Held section: The ID of the logical CPU holding the lock. Spin section: The decimal address space identifier of the spinning job.
Address	The address of the instruction which obtained the lock.
Held %	The percentage of samples where the lock has been held.
Spin %	The percentage of samples where the requesting address space (ASID) has been found spinning due to the unavailable lock.

Lock report options

On the **Lock Report Options** panel, you can specify options for the **Spin Lock Report** (LOCKSP) and the **Suspend Lock Report** (LOCKSU).

```

RMF Lock Report Options
Command ==>
Scroll ==> CSR

Change or verify parameters. To exit press END.
Changes will apply to the LOCKSP and the LOCKSU report.

Spin Lock ==> BOTH Information (HELD, SPIN or BOTH) in LOCKSP report
Lock Type ==> BOTH Lock type (GLOBAL, LOCAL or BOTH) in LOCKSU report

```

Figure 80. Lock Report Options

LOCKSU - Suspend Lock Report

Through locking, the system serializes the use of system resources by authorized routines and, in a Parallel Sysplex, by processors. Lock holders can impede other work units that need the same lock and must wait until the lock holder releases the lock.

RMF reports about the various types of system resource locks in the **Suspend Lock Report** described in this section and in the **Spin Lock Report** described in [“LOCKSP - Spin Lock Report”](#) on page 129.

If a suspend lock is unavailable, the unit of work requesting the lock is delayed until the lock is available. Other work is dispatched on the requesting processor. All local locks are suspend locks.

No symptoms for delays due to suspend locks are visible. Therefore, this report provides information about the jobs that are holding a suspend lock, because the overall workflow can be impacted by contention situations for the same lock. Especially, if a work unit that is holding a lock is suspended for a longer period of time, other work units can be significantly delayed.

How to request this report

To request the **Suspend Lock Report**, select **1** from the Primary Menu and then select **13** on the Overview Report Selection Menu (shown in [Figure 6 on page 23](#)), or enter one of the following commands:

```
LOCKSU [GLOBAL | LOCAL | BOTH]
```

```
LSU [GLOBAL | LOCAL | BOTH]
```

Contents of the report

The report contains one segment for local suspend locks in the upper part and one for global suspend locks in the lower part of the report. Within the segments, the report lines are sorted by descending Held%. A separate work unit within the same address space can be identified by the value in column **Address**, which is the address of the instruction that obtained the lock.

RMF 3.1 Suspend Lock Report								
Command ==>			Line 1 of 6 Scroll ==> PAGE					
Samples: 120	System: CB88	Date: 04/25/2023	Time: 09.21.00	Range: 120	Sec			
Resource	Type	Jobname	ASID	Address	Held %	Intr %	Disp %	Susp %
OS390R1	L	OS390R1	0045	07072C52	0.83	0.00	0.83	0.00
ZFS	L	ZFS	0234	012B2A66	0.83	0.00	0.83	0.00
GRS	L	GRS	0007	015EC438	0.83	0.00	0.83	0.00
ZFS	L	ZFS	0236	012B2A66	0.83	0.00	0.00	0.00
CMSSMFLK	G	ZFS	0062	00CE4652	0.83	0.00	0.00	0.00
CMSDLK	G	OS390R1	0045	2A64085A	0.83	0.00	0.83	0.00
CMSLOCK	G	HZSPROC	0059	00D3EC70	0.83	0.00	0.00	0.00

Figure 81. Suspend Lock Report

Table 52. Fields in the LSU Report	
Field Heading	Meaning
Resource	The resource name of the suspend lock. - for local locks: the address space name where the local lock resides - for cross memory local (CML) locks: the primary address space name (which is different from the holder's job name) - for all types of cross memory services (CMS) locks: the lock word name.
Type	The type of the suspend lock: L Local Suspend Lock LX Cross Memory Local (CML) Suspend Lock G Global CMS Suspend Lock
Jobname	The name of the job/address space holding the lock.
ASID	The decimal address space identifier of the job holding the lock.

Table 52. Fields in the LSU Report (continued)	
Field Heading	Meaning
Address	The address of the instruction that obtained the lock. For local locks, the address of the instruction can be in the lock address space or the requestor address space.
Held %	The percentage of samples where the address space held the lock during the report interval.
Intr %	The percentage of samples where the address space was interrupted while holding the lock.
Disp %	The percentage of samples where the address space was dispatchable while holding the lock.
Susp %	The percentage of samples where the address space has been found suspended while another dispatchable unit was holding the lock.

Note: You can specify options for this report on the **Lock Report Options** panel described in [“Lock report options”](#) on page 130.

OPD - OMVS Process Data Report

z/OS Unix address spaces can consist of several processes, which in turn might run one or more threads. Each process is typically associated with a UNIX command, consumes a certain amount of CPU, and also provides state information. UNIX System Services is the brand for UNIX on z/OS. In this context, it is referred to as open MVS or OMVS.

In addition to other reports that show OMVS address spaces with their jobname and using or delay information, the OPD report can be used for problem determination. It assists the performance analyst to find answers to the following questions:

- What are the delayed processes?
- What command is associated with them?
- What is the status of each of the processes?
- Which processes are high CPU consumers?

Address spaces under OMVS control are indicated by an additional letter **O** in the class column of the DELAY, the PROC, and the JOB report. The performance analyst can then use cursor sensitivity to navigate to this report or, alternatively, invoke it directly depending on the task to be accomplished.

The report provides basic performance metrics on the first screen, while additional information specifically related to server processes can be shown by activating a pop-up panel.

How to request this report

To request the OMVS Process Data report, select **1** from the Primary Menu and then select **7** on the Overview Report menu (shown in [Figure 6 on page 23](#)), or enter the following command:

```
OPD
```

Contents of the report

RMF 3.1 OMVS Process Data							Line 1 of 24		
Command ==>							Scroll ==> HALF		
Samples: 18		System: SYS4		Date: 04/25/2023		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
BPX0INIT	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 82. OPD Report

There is no graphic version of this report available.

Table 53. Fields in the OPD Report	
Field Heading	Meaning
Kernel Procedure	Name of the procedure used to start the OMVS kernel address space.
Kernel ASID	Decimal ID of the kernel address space.
BPXPRM	List of suffixes indicating the BPXPRM Parmlib member concatenation.
Option	Displays the current report option as specified on the Report Options panel.
Jobname	Jobname associated with the process.
User	User name associated with the process.
ASID	Decimal ID of the address space the process is associated with.
PID	Process ID.
PPID	Parent process ID.
LW	If the reported process is waiting for the process latch of another process, 'Y' is shown, otherwise blank.
State	Cumulated state information of the address space and process. You can place the cursor on any field (except Jobname and PPID) in a process line and press Enter - this will show you a pop-up panel with an explanation of the process state.
Appl%	Percentage of TCB and local/global SRB time consumed by the address space during the reporting range. Note: The calculated value is based on uniprocessor capacity and can exceed 100% on systems with more than one processor. To get the system utilization, this value has to be divided by the number of logical processors or cores.
Total	Total computing time in seconds, consumed by the address space the process is running within. When only one process is running in the address space, this time represents the accumulated CPU time for that process. In case of multiple processes running in an address space, it is the sum of the CPU time used by all of the work running in that address space.

Table 53. 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 pop-up panel shows an example of process details for a server process. For a non-server process, RMF displays 'N/A' in the fields below **Server Information**.

```

RMF OMVS Process Data - Details

Press Enter to return to the Report panel.

Start Time/Date : 12.08.57 04/25/2023
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 83. OPD Report - Details for Server Process

Table 54. Fields in the OPD Details Report	
Field Heading	Meaning
Start Time/Date	Start time and date when the process has been started.
Command	The command that created the processes truncated to 40 characters.
Process-ID	Process ID.
Parent Process-ID	Parent process ID.
Jobname	Jobname associated with the process.
User Name	User name associated with the process.
ASID	Decimal ID of the address space the process is associated with.
Hexadecimal ASID	Hexadecimal identifier of the address space.
Appl%	Percentage of TCB and local/global SRB time consumed by the address space during the reporting range. Note: APPL% shows CPU utilization based on uniprocessor capacity. On systems with more than one processor this value has to be divided by the number of processors to get the system utilization.
Total CT	Total computing time in seconds, consumed by the address space the process is running in. When only one process is running in the address space, this time represents the accumulated CPU time for that process. In case of multiple processes running in an address space, it is the sum of the CPU time used by all of the work running in that address space.
LW-PID	Process ID of the process on whose latch the reported process is waiting for.

Table 54. Fields in the OPD Details Report (continued)	
Field Heading	Meaning
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.

Table 54. Fields in the OPD Details Report (continued)

Field Heading	Meaning																																												
Process State	Cumulated state information of the address space and process. For each possible state a separate line is shown below the field. The following translation table is used: <table> <tr> <th>State</th><th>Meaning</th></tr> <tr> <td>1</td><td>Single thread</td></tr> <tr> <td>A</td><td>Message queue receive wait</td></tr> <tr> <td>B</td><td>Message queue send wait</td></tr> <tr> <td>C</td><td>Communication system kernel wait</td></tr> <tr> <td>D</td><td>Semaphore operation wait</td></tr> <tr> <td>E</td><td>Quiesce frozen</td></tr> <tr> <td>F</td><td>File system kernel wait</td></tr> <tr> <td>G</td><td>MVS pause wait</td></tr> <tr> <td>H</td><td>Multiple threads, pthread_create used</td></tr> <tr> <td>I</td><td>Swapped out</td></tr> <tr> <td>K</td><td>Other kernel wait</td></tr> <tr> <td>L</td><td>Cancelled, parent waits</td></tr> <tr> <td>M</td><td>Multiple threads, no pthread_create used</td></tr> <tr> <td>P</td><td>Ptrace kernel wait</td></tr> <tr> <td>Q</td><td>Quiesce termination wait</td></tr> <tr> <td>R</td><td>Running</td></tr> <tr> <td>S</td><td>Sleeping</td></tr> <tr> <td>T</td><td>Stopped</td></tr> <tr> <td>W</td><td>Waiting for child</td></tr> <tr> <td>X</td><td>Creating new process</td></tr> <tr> <td>Z</td><td>Zombie. Cancelled, parent does not wait</td></tr> </table>	State	Meaning	1	Single thread	A	Message queue receive wait	B	Message queue send wait	C	Communication system kernel wait	D	Semaphore operation wait	E	Quiesce frozen	F	File system kernel wait	G	MVS pause wait	H	Multiple threads, pthread_create used	I	Swapped out	K	Other kernel wait	L	Cancelled, parent waits	M	Multiple threads, no pthread_create used	P	Ptrace kernel wait	Q	Quiesce termination wait	R	Running	S	Sleeping	T	Stopped	W	Waiting for child	X	Creating new process	Z	Zombie. Cancelled, parent does not wait
State	Meaning																																												
1	Single thread																																												
A	Message queue receive wait																																												
B	Message queue send wait																																												
C	Communication system kernel wait																																												
D	Semaphore operation wait																																												
E	Quiesce frozen																																												
F	File system kernel wait																																												
G	MVS pause wait																																												
H	Multiple threads, pthread_create used																																												
I	Swapped out																																												
K	Other kernel wait																																												
L	Cancelled, parent waits																																												
M	Multiple threads, no pthread_create used																																												
P	Ptrace kernel wait																																												
Q	Quiesce termination wait																																												
R	Running																																												
S	Sleeping																																												
T	Stopped																																												
W	Waiting for child																																												
X	Creating new process																																												
Z	Zombie. Cancelled, parent does not wait																																												

Report options

```

RMF OMVS Process Data Report Options
Command ==>
Scroll ==> HALF

Change or verify parameters. To exit press END.
Select one of the following options:

1 1. Process ID ==> 1      ALL or a process ID
   2. ASID      ==> 0100   ID of an address space in decimal or
                           hexadecimal (with preceding X) format
   3. Jobname    ==>       Jobname associated with a process
   4. User       ==>       User name associated with a process

```

Figure 84. OPD Report Options Panel

You can specify a process ID, an address space ID (in decimal or hexadecimal format), a jobname, or a user name to tailor the OPD report.

PCIE - PCIE Activity Report

The PCIE Activity Report can be used to investigate performance problems that are related to PCI Express based functions (PCIE functions). This report provides these types of information:

- General PCIE activity metrics that are partially dependent of the type of the exploited hardware feature and reflect the activity of the z/OS system on which RMF data collection took place. These metrics include data rates for the communication between z/OS programs and the PCIE functions (like PCI LOAD, PCI STORE, PCI STORE BLOCK, REFRESH PCI TRANSLATIONS, and Read/Write Transfer data rates).
- Additional metrics that can be displayed on PCIE function type specific pop-up panels. These metrics provide more detailed configuration and performance information for
 - Hardware Accelerators:

Single system scope metrics including device driver buffer statistics, common accelerator metrics (for example, total request execution time or the amount of transferred data).
 - Hardware Accelerators with compression activity:

Single system scope metrics including device driver buffer statistics, common accelerator metrics (for example, total request execution time or the amount of transferred data) as well as compression specific metrics (for example, the amount of compressed data and the number and throughput of compression requests).
 - RoCE / NETH devices:

Adapter information and transfer statistics for data reads and writes on the external Ethernet interface.
 - Internal Shared Memory devices:

Adapter information and the data write rate on the SMC-D device.
 - Synchronous I/O links:

The port ID the synchronous I/O link uses, the type-model and serial number of the storage controller the synchronous I/O link connects to, and in addition data transfer and request information with a virtual function scope as well as with a CPC-wide scope.

The PCIE Activity Report provides statistics and performance measurements on PCIE functions allocated by at least one z/OS address space. A PCIE function is captured by the report if one of the following feature activities has been detected:

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

- IBM zHyperLink (zHypL)
- Network Express Hybrid (NETH)

How to request this report

To request the PCIE Activity Report, select a **3** from the Primary Menu, then select **14** from the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)) or enter the following command:

```
PCIE
```

Contents of the report

```

RMF 3.1  PCIE Activity                               Line 1 of 5
Command ==>                                         Scroll ==> CSR

Samples: 120      System: SYS1  Date: 04/25/2023  Time: 09.20.00  Range: 120  Sec

----- Function ----- Alloc - PCI Operations Rate - -Xfer Rate-
ID  CHID Type  Jobname  ASID Status  Time%  Load Store Block  Refr  Read Write
0041 013C zEDC  FPGHWAM  0012 Alloc    100    0 0.017    0 0.533
006C 0184 RoCE  VTAM      002B Alloc    100 0.120  3185    0 0.090  37.9  42.5
0430 0278 zHypL  IOSAS     001A Alloc    100
2624 0204 RoCE2  VTAM      002B Alloc    100 0.060  7094    0 0.130    0    0
2625 0204 RoCE2  VTAM      002B Alloc    100 0.060  7031    0 0.130    0    0

```

Figure 85. PCIE Activity Report

Table 55. Fields in the PCIE Activity Report

Field Heading	Meaning
Function ID	Hexadecimal identifier of the PCIE Function for which performance data is reported.
Function CHID	Physical or virtual channel identifier for the PCIE function.
Function Type	Device type for the PCIE function which can be one of the following: HWA Hardware Accelerator ISM SMC-Direct over Internal Shared Memory Oth Unknown function type RoCE Remote Direct Memory Management (RoCE and RoCE Express) RoCE2 Remote Direct Memory Management Express 2 zEDC zEnterprise Data Compression zHypL zHyperLink NETH Network Express Hybrid
Function Jobname	Name of the job who allocated the PCIE function.
Function ASID	Address space ID of the job who allocated the PCIE function.

Table 55. Fields in the PCIE Activity Report (continued)

Field Heading	Meaning
Function Status	The PCIE function status at the end of this reporting interval: Alloc The function is allocated and in use DPend The function is in the process of de-allocation Error The function is in permanent error DeAlloc The function is de-allocated Unknown The function status is unknown
Alloc Time %	Percentage of the reporting interval for which the PCIE function was allocated or in the process of de-allocation.
PCI Operations Rate Load	Rate of PCI Load operations that were executed within this reporting interval.
PCI Operations Rate Store	Rate of PCI Store operations that were executed within this reporting interval.
PCI Operations Rate Block	Rate of PCI Store Block operations that were executed within this reporting interval.
PCI Operations Rate Refresh	Rate of Refresh PCI Translations operations that were executed within this reporting interval.
Xfer Rate Read	The number of megabytes per second that a RoCE / RoCE-2 / NETH device received on the external Ethernet interface, respectively read on a synchronous I/O link. On zEC12 or zBC12, this field designates the number of megabytes per second that were transferred by DMA reads from all defined DMA address spaces to the PCIE function.
Xfer Rate Write	The number of megabytes per second transmitted on a RoCE / RoCE-2 / SMC-D / NETH device, or on a synchronous I/O link. On zEC12 or zBC12, this field designates the number of megabytes per second that were transferred by DMA writes from the PCIE function to all defined DMA address spaces.

Monitor III Utility fields

You can use the Monitor III Utility to customize the PCIE Activity report to show the fields listed in [Table 56 on page 139](#).

Table 56. Monitor III Utility fields

Field Heading	Meaning
PCIEDMAN	Number of DMA address spaces.
PCIEFTR	Hardware Accelerator transfer rate.

Cursor sensitive control

If the cursor is placed on one of the values in the Function ID, Function CHID, or Function Type columns, additional metrics can be displayed on a pop-up panel.

Depending on the function type one of the following pop-up panels is displayed:

Function Type	Pop-up panel
HWA, zEDC	RMF Hardware Accelerator And Compression Activity
RoCE, RoCE2	RMF RoCE Activity
ISM	RMF Internal Shared Memory Activity
zHypL	RMF Synchronous I/O Link Activity
NETH	RMF Network Express Hybrid

Fields available on all pop-up panels:

Field Heading	Meaning
Function ID	The hexadecimal identifier of the PCIE function for which performance data is reported.
Alloc Time %	The percentage of the reporting interval for which the PCIE function was allocated or in the process of de-allocation.
Allocated	Date and time when the PCIE function was allocated.


RMF Hardware Accelerator And Compression Activity Press Enter to return to the Report panel. Function ID : 0041 Alloc Time % : 100 More: + Allocated : 08.07.11 on 11/30/22 Hardware Accelerator Time Busy % : 0.286 Adapter Utilization : 66.67 Work Units Processed : 7.40 Transfer Read Rate : Write Rate : Request Execution Time : 28.00 Std. Deviation: 8.07 Queue Time : 65.70 Std. Deviation: 140.00 Size : 47.60 Buffer Pool Memory Size : 16 Utilization % : 0 Request Rate : 102 Compression Decompression Throughput : 2.91 0.437 Ratio : 2.79 0.009 0.652	
<i>Figure 86. RMF Hardware Accelerator And Compression Activity panel</i>	

Table 57. Fields in the RMF Hardware Accelerator And Compression Activity pop-up panel.

Field Heading	Meaning
Hardware Accelerator Time Busy %	The percentage of time that this partition kept the Hardware Accelerator busy.
Hardware Accelerator Adapter Utilization	Utilization of the Hardware Accelerator. This value is not reported on zEC12 and zBC12 hardware.
Hardware Accelerator Work Units Processed	The number of work units per second that were processed by the Hardware Accelerator. This value is not reported on zEC12 and zBC12 hardware.
Transfer Read Rate	On zEC12 or zBC12, this field designates the number of megabytes per second that were transferred by DMA reads from all defined DMA address spaces to the PCIE function. Otherwise, this field is not set.

Table 57. Fields in the RMF Hardware Accelerator And Compression Activity pop-up panel. (continued)

Field Heading	Meaning
Transfer Write Rate	On zEC12 or zBC12, this field designates the number of megabytes per second that were transferred by DMA writes from the PCIe function to all defined DMA address spaces. Otherwise, this field is not set.
Request Execution Time	The average time in microseconds the Hardware Accelerator took to process a request.
Request Execution Time Std. Deviation	The standard deviation of the request execution time.
Request Queue Time	The average queue time in microseconds that was spent for a request.
Request Queue Time Std. Deviation	The standard deviation of the request queue time.
Request Size	The average number of kilobytes that were transferred per request.
Buffer Pool Memory Size	The total size of memory in megabytes that is allocated to the buffer pool.
Buffer Pool Utilization %	The average utilization percentage of the buffer pool that z/OS kept for in-use buffers.
Request Rate	The number of compression or decompression requests per second.
Throughput	The number of megabytes that were compressed or decompressed per second.
Ratio	The ratio between input and output bytes that were compressed or decompressed within this reporting interval.






Table 58. Fields in the RMF RoCE and NETH Activity pop-up panels.

Field Heading	Meaning
Adapter Physical Network ID Port 1 and Port 2	Physical-network identifier (PNET ID) that identifies the first or second port of the RoCE / RoCE-2 / NETH device.
Transfer Read Rate	Megabytes The number of megabytes per second that a RoCE / RoCE-2 / NETH device received on the external Ethernet interface. On zEC12 or zBC12, this field designates the number of megabytes per second that were transferred by DMA reads from all defined DMA address spaces to the PCIE function. Packets The number of packets per second that were received on the external Ethernet interface of the RoCE / RoCE-2 / NETH device. This value is not reported on zEC12 and zBC12 hardware.
Transfer Write Rate	Megabytes The number of megabytes per second transmitted from a RoCE / RoCE-2 / NETH device on the external Ethernet interface. On zEC12 or zBC12, this field designates the number of megabytes per second that were transferred by DMA writes from the PCIE function to all defined DMA address spaces. Packets Number of packets per second transmitted from the RoCE / RoCE-2 / NETH device on the external Ethernet interface. This value is not reported on zEC12 and zBC12 hardware.




Table 59. Fields in the RMF Internal Shared Memory Activity pop-up panel.

Field Heading	Meaning
Adapter Physical Network ID	Physical-network identifier (PNET ID) that identifies the port of the Internal Shared Memory (ISM) virtual PCIe function.
Transfer Write Rate	The number of megabytes per second transmitted on an ISM function.




Table 60. Fields in the RMF Synchronous I/O Link Activity pop-up panel.

Field Heading	Meaning
Data in “Link(CPC)” fields is only reported if Global Performance Reporting is enabled in the LPAR image profile of the Hardware Management Console (HMC).	
Synchronous I/O Link Port ID	ID of the port on the zHyperlink card where the synchronous I/O link is attached to.

Table 60. Fields in the RMF Synchronous I/O Link Activity pop-up panel. (continued)	
Field Heading	Meaning
Synchronous I/O Link Type-Model	The type and model of the storage controller the synchronous I/O link connects to.
Synchronous I/O Link Serial Number	The serial number of the storage controller the synchronous I/O link connects to.
Synchronous I/O Link ID	The identifier of the synchronous I/O link that is configured in the storage controller.
Adapter Time Busy %	This Function The percentage of time spent on synchronous I/O processing by this synchronous I/O function. Link (CPC) The percentage of time spent on synchronous I/O processing on the synchronous I/O link this function is defined on.
Request Rate	This Function The total number of synchronous I/O requests per second for this function. Link (CPC) The total number of synchronous I/O requests per second for the synchronous I/O link this function is defined on.
Request Success %	This Function The percentage of synchronous I/O requests that completed successfully for this function. Link (CPC) The percentage of requests that completed successfully for the synchronous I/O link this function is defined on.
Transfer Read Rate	This Function The number of megabytes per second that were read from the storage controller by this synchronous I/O function. Link (CPC) The number of megabytes per second that were read from the storage controller on the synchronous I/O link this function is defined on.
Transfer Read Ratio	This Function The number of megabytes read per request processed by this synchronous I/O function. Link (CPC) The number of megabytes read per request processed on the synchronous I/O link this function is defined on.

Table 60. Fields in the RMF Synchronous I/O Link Activity pop-up panel. (continued)

Field Heading	Meaning
Transfer Write Rate	This Function The number of megabytes per second that were written to the storage controller by this synchronous I/O function. Link (CPC) The number of megabytes per second that were written to the storage controller on the synchronous I/O link this function is defined on.
Transfer Write Ratio	This Function The number of megabytes written per request processed by this synchronous I/O function. Link (CPC) The number of megabytes written per request processed on the synchronous I/O link this function is defined on.

PROC - Processor Delays Report

The Processor Delays report (PROC) displays all jobs that were waiting for or using the processor during the report interval.

RMF reports the jobs by descending overall delay percentages. Because use of the processor by many jobs might contribute to the delay of another job, RMF reports up to three jobs in the Holding Job(s) field. The jobs in this field are those that were most often found using the processor while the job was delayed.

How to request this report

To request the Processor Delay report, select 3 from the Primary Menu, then **1** from the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)) or enter the following command:

```
PROC [job_class,service_class]
```

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

```
PROC T, TSOPRIME
```

Contents of the report

```

RMF 3.1 Processor Delays
Line 1 of 138
Command ==>
Scroll ==> HALF

Samples: 60      System: MVS1  Date: 04/25/2023  Time: 09.10.00  Range: 60 Sec

Jobname      CX  Service  CPU  DLY  USG  EApp1  -----  Holding Job(s) -----
              Class  Type  %    %    %    %  Name      %  Name      %  Name
WSWS7        0  OMVS     CP   11  46  59.4   9 *ENCLAVE  7 DBS3DIST  7 WSP1S2F
WSP1S2FS     SO  WASCR     CP    4   4  42.5   2 DBS3DIST  2 WWS7      2 VTAM44
              AAP    6   0  98.4   6 *ENCLAVE
WSP1S6FS     SO  WASCR     CP    0   0   5.3
              AAP    6   0   7.7   6 *ENCLAVE
DBS3DBM1     S  DB2HIGH  CP    2   6   0.8   2 XCFAS     2 DBS3DIST  2 WSP1S2F
WSP1S6F      SO  WASCR     CP    0   2   1.9
              AAP    2   2   0.7   2 *ENCLAVE
U078069      0  OMVS     CP    2   4   1.2   2 WWS7      2 DBS3DIST  2 U078069
WSP1S4F      SO  WASCR     CP    0   0   0.1
              AAP    2   0   0.4   2 WSP1S6F
U078068      0  OMVS     CP    2   0   0.2   2 XCFAS     2 WWS7      2 *ENCLAVE
DBS3DIST     SO  DB2HIGH  CP    0  78 111.0
              IIP    0   2  21.3
XCFAS        S  SYSTEM   CP    0  28  24.1
TCPIP        SO  SYSSTC   CP    0  22  16.1
VTAM44       S  SYSSTC   CP    0  19  14.5
WSP1S2F      SO  WASCR     CP    0  15  14.0

```

Figure 91. PROC - Processor Delay Report

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

Field descriptions

Table 61. 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.
CPU Type	The processor type: CP general purpose processor AAP Application Assist Processor (zAAP) CBP Container Based Processor (zCBP) IIP Integrated Information Processors (zIIP)

Table 61. Fields in the PROC Report (continued)	
Field Heading	Meaning
DLY %	Delay percentage that the waiting job (address space) is experiencing because of contention for the processor of the type indicated in column CPU Type during the report interval. $\text{DLY \%} = \frac{\text{\# Delay Samples}}{\text{\# Samples}} * 100$ 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 of the type indicated in column CPU Type. $\text{USG \%} = \frac{\text{\# Using Samples}}{\text{\# Samples}} * 100$ 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.
EAppl%	Percentage of the processor time used by transactions that executed on the type of processor indicated in column CPU Type. This calculation is based on uniprocessor capacity which means that this value can exceed 100% in systems with more than one processor. To get the system utilization, this value has to be divided by the number of logical processors or cores. Note: The processor times that are used to calculate this value is the sum of TCB time, global and local SRB time, preemptable or client SRB time, and enclave CPU time consumed within this address space.
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. 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. $\text{\%} = \frac{\text{\# Holding Samples}}{\text{\# Samples}} * 100$ 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.

Monitor III Utility fields

Table 62 on page 148 shows the additional fields you can select for this report.

Table 62. Additional Fields in the PROC Report

Field Heading	Meaning
PRCPCAP	This column contains the actual delay caused by WLM • due to a resource group maximum • or due to discretionary goal management. This means that the work in question may be overachieving its goal. It may be capped to divert its resources to run discretionary work (see also section 'Using Discretionary Goals' in <i>z/OS MVS Planning: Workload Management</i>). It shows the Dispatchable Unit (TCB or SRB) capped delay, which should be distinguished from the address space capping state shown in the CAPP field on the Work Manager Delays report.
PRCPODEL	Overall delay percentage for this address space.
PRCPOUSE	Overall using percentage for this address space.
PRCPTST	Percentage of the processor time used by non-enclave work that executed on behalf of this address space.
PRCPETST	Percentage of the processor time used by enclave and non-enclave work that executed within this address space.
PRCPAPPL	Percentage of the processor time used by non-enclave work that executed on behalf of this address space and processor type.
PRCPTWFL	Overall workflow percentage of this address space and processor type.
PRCPTUSE	Overall using percentage for this address space and processor type.
PRCPUCP	Overall using percentage on general purpose processors for zAAPs and zIIPs.
PRCPASI	Address space ID of the job.
PRCTCPUT	Total CPU time (milliseconds).
PRCPRPCL	Report Class Name.
PRCPSCLP	Service Class Period.

Report options

The PROC Report Options panel is similar to the DEV Report Options panel. See [Figure 43 on page 75](#) for an example. Selecting YES for Jobs on the Report Options panel displays the Job Selection/Exclusion panel (see [Figure 41 on page 73](#)).

PROCU - Processor Usage Report

The Processor Usage report (PROCU) displays all jobs that were using a general purpose or special purpose processor during the report interval. RMF reports the jobs by descending CP EAppl % time. The report gives you information about the percentage of CPU time on general purpose processors consumed on behalf of the job. In addition, the percentage of CPU time used by work that is eligible for being offloaded to an Application Assist (zAAP) or Integrated Information (zIIP) processor is shown. You can use this information to understand the benefit of adding a zAAP or zIIP into the configuration.

The EAppl fields also display the percentage of task time, SRB and enclave CPU time consumed within the address space on general purpose processor or special purpose (zAAP and zIIP) processors.

How to request this report

To request the Processor Usage report, select **3** from the Primary Menu, then **1A** from the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)) or enter the following command:

```
PROCU [job_class,service_class]
```

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

```
PROCU T, TSOPRIME
```

Contents of the report

RMF 3.1									
Command ==>					Processor Usage			Line 1 of 153	
					Scroll ==> CSR				
Samples: 60		System: CB88		Date: 04/25/2023	Time: 08.00.00		Range: 60	Sec	
Jobname	CX	Service Class	--- Total	Time on CP	% --- AAP IIP	----- CP	EAppl % AAP	----- IIP	Rcvy Prc Boost
JAVA564M	BO	JESMED4	125.3	121.5	0.000	125.3	54.00	0.000	NO
GRS	S	SYSTEM	74.55	0.000	0.000	74.55	0.000	0.000	NO
XCFAS	S	SYSTEM	25.70	0.000	0.000	25.70	0.000	0.000	NO
JAVA664M	BO	JESMED4	20.11	19.86	0.000	20.11	9.495	0.000	NO
JAVA664A	BO	JESMED4	18.18	17.83	0.000	18.18	10.61	0.000	NO
TPNSADSW	S	STCHI	13.35	0.000	0.000	13.35	0.000	0.000	NO
CICS2A23	SO	STCHI	13.15	0.000	0.000	13.15	0.000	0.000	NO
CICS2A21	SO	STCHI	11.97	0.000	0.000	11.97	0.000	0.000	NO
CICS2A22	SO	STCHI	10.54	0.000	0.000	10.54	0.000	0.000	NO
CICS2T21	SO	STCHI	5.512	0.000	0.000	5.512	0.000	0.000	NO
CICS2T23	SO	STCHI	5.510	0.000	0.000	5.510	0.000	0.000	NO
CICS2T22	SO	STCHI	5.497	0.000	0.000	5.497	0.000	0.000	NO
CATALOG	S	SYSTEM	4.085	0.000	0.000	4.085	0.000	0.000	NO
SMSVSAM	S	SYSTEM	3.133	0.000	0.000	3.133	0.000	0.000	NO
IEFSCHAS	S	SYSTEM	3.070	0.000	0.000	3.070	0.000	0.000	NO
ZFS	S	SYSSTC	2.278	0.000	0.000	2.278	0.000	0.000	NO

Figure 92. PROCU - Processor Usage Report

Field descriptions

Table 63. Fields in the PROCU Report	
Field Heading	Meaning
Jobname	The name of a job using processor time. The processor usage report does not summarize data by job groups; all jobs within a job group are reported individually.
CX	Abbreviation for the job class as follows: 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.
Time on CP %	Total Percentage of CPU time spent on general purpose processors as sum of TCB time, global and local SRB time, and preemptable or client SRB time consumed on behalf of this address space. IIP Percentage of CPU time on general purpose processors by this address space which was used by zIIP eligible work. This is a subset of the Total percentage. CBP Percentage of CPU time on general purpose processors by this address space which was used by zCBP eligible work. This is a subset of the Total percentage. Note: The calculated values are based on uniprocessor capacity, which means that they can exceed 100% on systems with more than one processor. To get the system utilization, this value has to be divided by the number of logical processors or cores.

Table 63. Fields in the PROCU Report (continued)

Field Heading	Meaning
EAppl %	CP Percentage of CPU time on general purpose processors (CPs) as sum of TCB time, global and local SRB time, preemptable or client SRB time, and enclave CPU time consumed within this address space. CBP Percentage of CPU time consumed on zCBPs within this address space. IIP Percentage of CPU time consumed on zIIPs within this address space. Note: EAppl% shows CPU utilization based on uniprocessor capacity. This means that the value can exceed 100% in systems with more than one processor. To get the system utilization, this value has to be divided by the number of logical processors or cores.
Rcvy Prc Boost	YES or NO . Indicates whether the address space had Recovery Process Boost active during the interval.

Monitor III Utility fields

Table 64 on page 150 shows the additional fields you can select for this report.

Table 64. Additional Fields in the PROCU Report

Field Heading	Meaning
PRUPCLA	Class (A, B, E, O, S, or T).
PRUPTOTC	Percentage of total accumulated CPU time as sum of TCB time, global and local SRB time and preemptable or client SRB time, consumed on behalf of this address space.
PRUPTOTE	Percentage of total accumulated CPU time as sum of TCB time, global and local SRB time, preemptable or client SRB time, and enclave time consumed within this address space.
PRUPTCB	Percentage of TCB time consumed in this address space.
PRUPSRB	Percentage of SRB time consumed in this address space by local or global SRBs.
PRUPPCS	Percentage of preemptable or client SRB time consumed on behalf of this address space.
PRUPEPS	Percentage of preemptable or client SRB and enclave CPU time consumed within this address space.
PRUPASI	Address space ID of the job.
PRUTCPT	Total time on CP in milliseconds
PRUTCBCT	CBP time on CP in milliseconds
PRUTAACT	AAP time on CP in milliseconds
PRUTIICT	IIP time on CP in milliseconds
PRUTCPE	CP EAppl in milliseconds
PRUTCBPE	CBP EAppl in milliseconds
PRUTAAPE	AAP EAppl in milliseconds
PRUTIIPE	IIP EAppl in milliseconds
PRUTTOTC	Total accumulated CPU time as sum of TCB time, global and local SRB time, and preemptable or client SRB time consumed on behalf of this address space in milliseconds.
PRUTTOTE	Total accumulated CPU time as sum of TCB time, global and local SRB time, preemptable or client SRB time, and enclave time consumed within this address space in milliseconds.
PRUTTCB	TCB time consumed in this address space in milliseconds.
PRUTSRB	SRB time consumed in this address space by local or global SRBs in milliseconds.
PRUTPCS	Preemptable or client SRB time consumed on behalf of this address space in milliseconds.
PRUTEPS	Preemptable or client SRB and enclave CPU time consumed within this address space in milliseconds.
PRUPCLP	Service class period.

Table 64. Additional Fields in the PROCU Report (continued)	
Field Heading	Meaning
PRUTCPUT	Total CPU time (milliseconds).
PRUPRPL	Report class name.

Report options

The PROCU Report Options panel is similar to the Device Report Options panel. See [Figure 43 on page 75](#) for an example. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed (see [Figure 41 on page 73](#)).

RLSLRU - VSAM LRU Overview Report

This report provides Local Buffer Manager LRU statistics for each system. The data in this report can help you in adjusting the goal and the limit for the local cache size.

In Parmlib member IGDSMSxx, there is a goal response limit for the local cache size, it defaults to 100 MB. You can specify a limit up to 1.5 GB (if a bigger value will be given, the report will display MAX as buffer size goal). Each LRU cycle, it is determined whether the system is over the goal and the buffer aging algorithms are accelerated. If the system is 5 times over the goal or reaches the 1.5 GB limit, the system starts clearing the buffers. If systems appear where BMF is over the goal (status *Accelerated* or *Reclaimed*), you could adapt the goal in Parmlib member IGDSMSxx by changing the RLS_MAX_POOL_SIZE value.

How to request this report

To request the VSAM LRU Overview report, select **S** from the Primary Menu and then select **12** on the Sysplex Report menu (shown in [Figure 5 on page 22](#)), or enter the following command:

```
RLSLRU
```

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

Contents of the report

```

RMF 3.1 VSAM LRU Overview  - SYSPLEX                      Line 1 of 2
Command ==>                      Scroll == => HALF

Samples: 120      Systems: 2      Date: 04/25/2023  Time: 13.25.00  Range: 120  Sec

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

SYS4
  Below 2GB  0.023      MAX    1M     0.0    0.0     0.0    0.0    0.0
  Above 2GB  3.543      MAX    1M     0.0    0.0    97.5    0.0    2.5
SYS5
  Below 2GB  4.457      MAX    1M     0.0    0.0     0.0    0.0    0.0

```

Figure 93. RLSLRU Report

Cursor-sensitive control on a system line displays a pop-up panel with buffer counts by pool for the selected system. Sixteen storage pools (2K, 4K, ... 32K) are available.

RMF VSAM LRU Overview - Buffer Counts by Pool

The following details are available for MVS System: SYSF
Press Enter to return to the Report panel.

Fixed Pages Low : 305 Fixed Storage : 500
 High : 305 Real Storage % : 46
 Average : 305

Size	----- Below 2 GB -----			----- Above 2 GB -----			More: +
	Low	High	Avg	Low	High	Avg	
2K	163	315	226	0	0	0	
4K	713	1537	1299	0	0	0	
6K	0	0	0	0	0	0	
8K	0	0	0	0	0	0	
10K	0	0	0	0	0	0	
12K	0	0	0	347	458	412	
14K	0	0	0	0	0	0	
16K	460	678	656	0	0	0	
18K	0	0	0	0	0	0	
20K	0	0	0	0	0	0	
22K	0	0	0	0	0	0	
24K	0	0	0	0	0	0	
26K	0	0	0	0	0	0	
28K	0	0	0	0	0	0	
30K	0	0	0	0	0	0	
32K	0	0	0	0	0	0	

Figure 94. VSAM LRU Overview - Buffer Counts by Pool

Field descriptions

Note: If applicable, the measurements in the RLSLRU Report are presented for storage addresses below and above the 2GB bar. Two lines of data are then displayed for each system and a label **Above 2GB** or **Below 2GB** precedes the corresponding line as shown in Figure 93 on page 151. Also, all measurements in the Buffer Counts by Pool panel are presented for storage addresses below and above the 2GB bar. This is indicated by the corresponding column headings as shown in Figure 94 on page 152.

Table 65. Fields in the RLSLRU Report	
Field Heading	Meaning
MVS System	System name.
Avg CPU Time	Average CPU time spent by BMF LRU processing during each report interval (milliseconds).
Buffer Size: Goal	Buffer size goal (MB). If no valid goal has been defined, MAX will be shown.
Buffer Size: High	Buffer size actual high value (MB).
Accel%	Percentage of Buffer Manager LRU intervals when BMF was over the goal and buffer aging algorithms were accelerated.
Reclaim%	Percentage of Buffer Manager LRU intervals when BMF was over the goal and buffer aging algorithms were bypassed to reclaim buffers.
Read	BMF% Percentage of READ requests that could be satisfied from local buffers being managed by SMSVSAM. CF% Percentage of CF cache structure READ requests. DASD Read % Percentage of READ requests to DASD.

Fields in the Buffer Counts by Pool panel: There are 16 buffer pools with different buffer sizes between 2K and 32K incremented by 2K. For each pool, this panel presents a line with the high, low and average numbers of BMF buffers during this interval. These values are provided for storage pools allocated to addresses below and above the 2GB-bar. If values above 2GB are not available, '-' is displayed.

Field Heading	Meaning
Fixed Storage	The amount of buffer pool storage that is specified to be fixed. If no value is available, a '-' is shown.
Real Storage %	The value specified by the RLSFIXEDPOOLSIZE parameter divided by the amount of real storage in the system (in percent). If no value is available, a '-' is shown.
Fixed Pages	Low, high and average actual number of fixed 4KB-pages. If no value is available, a '-' is shown.

RLSSC/RLSDS - VSAM RLS Activity Report

With VSAM RLS, GETs and PUTs are executed by SMSVSAM on behalf of the application. When the application's data request can be satisfied from SMSVSAM's local buffers, no I/O is necessary. If the data in the buffers is invalid, SMSVSAM accesses the VSAM RLS cache structures in the coupling facility to fulfill the request. If data in the cache structures is invalid, a DASD I/O is performed.

The report is providing VSAM RLS activity data regarding READ and WRITE requests accessing the local buffers, the CF cache structures and DASD. This data might help you in answering important questions like

- Are there problems with LRU (Least Recently Used algorithms) or buffer pool sizes?
- Are the CF cache structures too small?

How to request this report

The VSAM RLS Activity report has two different versions:

- VSAM RLS Activity by Storage Class
- VSAM RLS Activity by Data Set

Both versions of the report have a similar structure, but a different scope of data. From the initial Sysplex Total View for each version, you can navigate to a System/CF Structure View.

To request the VSAM RLS Activity report, select **S** from the Primary Menu and then select **10** (for storage class) or **11** (for data set) on the Sysplex Report menu (shown in [Figure 5 on page 22](#)), or enter one of the following commands:

```
RLSSC [storage_class]
```

```
RLSDS [data_set_name]
```

In addition, you can navigate between the two versions of the report with cursor-sensitive control.

Contents of the report

VSAM RLS Activity by Storage Class

This report provides a VSAM RLS activity view by storage class. For each storage class, sysplex wide totals will be displayed for direct and sequential access.

```

RMF 3.1  VSAM RLS Activity  - SYSPLEX          Line 1 of 12
Command ==>                               Scroll == => HALF

Samples: 59      Systems: 2      Date: 04/25/2023  Time: 13.16.00  Range: 60      sec

          < 2GB / > 2GB
LRU Status  : Good / Accel
Contention % : 0.0 / 0.0
False Cont % : 0.0 / 0.0

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

RLS
  Below 2GB  DIR     0.004  665.6  88.2   0.5   11.3   100     0.01   0.00
              SEQ     0.000   0.00   0.0    0.0    0.0    0.0     0.00   0.00
  Above 2GB  DIR     0.004  665.6  88.2   0.5   11.3   100     0.01   0.00
              SEQ     0.000   0.00   0.0    0.0    0.0    0.0     0.00   0.00
RLS1
  Below 2GB  DIR     0.005  200.0  90.5   0.0    9.5   100     0.00   0.00
              SEQ     0.000   0.00   0.0    0.0    0.0    0.0     0.00   0.00
RLS2
  Below 2GB  DIR     0.003  213.3  90.5   0.0    9.5   100     0.00   0.00
              SEQ     0.000   0.00   0.0    0.0    0.0    0.0     0.00   0.00
RL3
  Above 2GB  DIR     0.004  665.6  88.2   0.5   11.3   100     0.01   0.00
              SEQ     0.000   0.00   0.0    0.0    0.0    0.0     0.00   0.00

```

Figure 95. VSAM RLS Activity by Storage Class - Sysplex Total View

Cursor-sensitive control on the LRU STATUS field displays the VSAM LRU Overview report (see [Figure 93](#) on page 151) with the Local Buffer Manager LRU statistics for each system.

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

```

RMF 3.1  VSAM RLS Activity  - SYSPLEX          Line 1 of 23
Command ==>                               Scroll == => HALF

Samples: 120     Systems: 2      Date: 04/25/2023  Time: 13.25.00  Range: 120     Sec

          < 2GB / > 2GB
LRU Status  : Good / Accel
Contention % : 0.0 / 0.0
False Cont % : 0.0 / 0.0
Stor Class  : RLS
Cache Set   : PUBLIC1
Lock Set    : RLSLOCKSET
Lock Struct : RLSLOCKSTR

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

*ALL
  Below 2GB  DIR     0.000  14.98  83.0   0.0   17.0   100     0.00   0.00
              SEQ     0.000   0.00   0.0    0.0    0.0    0.0     0.00   0.00
  Above 2GB  DIR     0.000  14.98  83.0   0.0   17.0   100     0.00   0.00
              SEQ     0.000   0.00   0.0    0.0    0.0    0.0     0.00   0.00
SYS4
  CACHE01
    Below 2GB  DIR     0.000   7.49  83.0   0.0   17.0   100     0.00   0.00
              SEQ     0.000   0.00   0.0    0.0    0.0    0.0     0.00   0.00
    Above 2GB  DIR     0.000   7.49  83.0   0.0   17.0   100     0.00   0.00
              SEQ     0.000   0.00   0.0    0.0    0.0    0.0     0.00   0.00
  CACHE02
    Below 2GB  DIR     0.000   0.00   0.0    0.0    0.0    0.0     0.00   0.00
              SEQ     0.000   0.00   0.0    0.0    0.0    0.0     0.00   0.00
  CACHE03
    Above 2GB  DIR     0.000   0.00   0.0    0.0    0.0    0.0     0.00   0.00
              SEQ     0.000   0.00   0.0    0.0    0.0    0.0     0.00   0.00

```

Figure 96. VSAM RLS Activity by Storage Class - System/CF Structure View

At the top of the report, the sysplex wide totals for the storage class will be displayed (indicated by *ALL) followed by report lines per system and CF cache structure.

Cursor-sensitive control on a CF structure name displays the CF Activity report (see [Figure 23 on page 40](#)) for this structure.

VSAM RLS Activity by Data Set

This report provides a VSAM RLS activity view by VSAM data sets. The information is grouped by VSAM spheres. A sphere consists of components, that is, data sets such as BASE.DATA, BASE.INDEX, ALT.DATA and ALT.INDEX. For each data set, sysplex wide totals are displayed for direct and sequential access. Only those VSAM spheres for which data collection has been requested are presented .

```

RMF 3.1 VSAM RLS Activity  -  SYSPLEX                      Line 1 of 20
Command ==>                                           Scroll == => HALF

Samples: 120      Systems: 2      Date: 04/25/2023  Time: 13.25.00  Range: 120  Sec

LRU Status      : < 2GB / > 2GB
                  Good / Accel
Contention %    : 0.0 / 0.0
False Cont %    : 0.0 / 0.0

Sphere/DS      Access  Resp  ----- Read  ----- BMF ----- Write
                Time    Rate  Rate  CF%  DASD%  Valid%  False Inv%  Rate
BMAI.VSAMIN.MEGA
BMAI.VSAMIN.MEGA.AIX.DATA
  Below 2GB  DIR    0.003  0.01  0.0  0.0  100  0.0  0.00  0.00
              SEQ    0.000  0.00  0.0  0.0  0.0  0.0  0.00  0.00
  Above 2GB  DIR    0.003  0.01  0.0  0.0  100  0.0  0.00  0.00
              SEQ    0.000  0.00  0.0  0.0  0.0  0.0  0.00  0.00
BMAI.VSAMIN.MEGA.AIX.INDEX
  Below 2GB  DIR    0.003  0.03  50.0  0.0  50.0  100  0.00  0.00
              SEQ    0.000  0.00  0.0  0.0  0.0  0.0  0.00  0.00
  Above 2GB  DIR    0.003  0.03  50.0  0.0  50.0  100  0.00  0.00
              SEQ    0.000  0.00  0.0  0.0  0.0  0.0  0.00  0.00
BMAI.VSAMIN.MEGA.DATA
  Below 2GB  DIR    0.000  7.45  83.2  0.0  16.8  100  0.00  0.00
              SEQ    0.000  0.00  0.0  0.0  0.0  0.0  0.00  0.00
MAI1.VSAMIN.MEGA
MAI1.VSAMIN.MEGA.AIX.DATA
  Above 2GB  DIR    0.003  0.01  0.0  0.0  100  0.0  0.00  0.00
              SEQ    0.000  0.00  0.0  0.0  0.0  0.0  0.00  0.00

```

Figure 97. VSAM RLS Activity by Data Set - Sysplex Total View

Cursor-sensitive control on the LRU STATUS field displays the VSAM LRU Overview report (see [Figure 93 on page 151](#)) with the Local Buffer Manager LRU statistics for each system.

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

```

RMF 3.1 VSAM RLS Activity - SYSPLEX                      Line 1 of 9
Command ==>                                              Scroll == => HALF

Samples: 120      Systems: 2      Date: 04/25/2023  Time: 13.25.00  Range: 120  Sec

                < 2GB / > 2GB                      Stor Class : RLS01
LRU Status      : Good / Accel                      Cache Set   : RLSCSET
Contention %    : 0.0 / 0.0                        Data Set    : RLSADSW.VF01D.INVENTOR.INDEX
False Cont %    : 0.0 / 0.0                        Lock Struct: RLSLOCKSTR

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

*ALL
Below 2GB  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
Above 2GB  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
Below 2GB  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
Above 2GB  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 98. VSAM RLS Activity by Data Set - System/CF Structure View

Cursor-sensitive control on a CF structure name displays the CF Activity report (see [Figure 23 on page 40](#)) for this structure.

Field descriptions

Note: If applicable, all measurements in the VSAM RLS Activity reports are presented for storage addresses below and above the 2GB bar. This is either indicated by the heading **< 2GB / > 2GB** or by the labels **Above 2GB** or **Below 2GB**.

Table 66. Fields in the VSAM RLS Activity report	
Field Heading	Meaning
The following information is displayed in the report header depending on the report view:	
LRU Status	LRU status of local buffers under control of BMF (Buffer Management Facility). 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. If the value above 2 GB is not available, '-' is displayed.
False Cont %	Percentage of false LOCK contentions: All external requests issued by connectors that experience "hash contention". This occurs because a hashing algorithm is used to map a lock request to a lock table entry. When more than one lock request maps to a lock table entry, there is the potential for contention delay. You may need to increase the size of the lock table. If the value above 2 GB is not available, '-' is displayed.
Stor Class	Storage class name.
Cache Set	DFSMS cache set name.
Data Set	VSAM data set name.

Table 66. Fields in the VSAM RLS Activity report (continued)	
Field Heading	Meaning
Lock Set	Lock set name. Multiple Lock Structure support allows defining an additional Coupling Facility DFSMS lock structure to be associated with a single SMS storage class. Each lock set can contain a single lock structure name. If multiple lock structures are not supported, the line remains blank.
Lock Struct	Name of the lock structure associated with the lock set. If no lock structure is associated, IGWLOCK00 is used and reported. If multiple lock structures are not supported, IGWLOCK00 is also used, however, the line remains blank.
The body of the report contains the following columns. The first column depends on the report flavour and view.	
Stor Class	Storage class name (first column in the Sysplex Total View in the Storage Class flavour).
System/CF	System name (*ALL indicates a sysplex-wide view) and the CF cache structure name (first column in the System/CF Structure View in both flavours).
Sphere/DS	VSAM sphere name and the VSAM data set name (first column in the Sysplex Total View in the Data Set flavour).
Access	Indicates whether the values in this row are shown for direct access (DIR) or sequential access (SEQ). There may be up to two sets of data depending on whether buffers above or below 2 GB were accessed.
Resp Time	Average response time of all requests (seconds).
Read Rate	Total number of BMF READ requests per second. BMF READ requests is sum of BMF valid READ hits, CF READ hits and DASD READs. The value is reported as '<0.01' if the rate is greater than 0 but below 0.01.
BMF Read%	Percentage of BMF valid READ hits that is the percentage of READ requests that were satisfied from local buffers being managed by SMSVSAM.
CF Read%	Percentage of READ requests that were satisfied by the CF cache structure.
DASD Read%	Percentage of READ requests to DASD.
BMF Valid%	Percentage of BMF READ hits that were valid. If a buffer is found in the local cache and is determined to be valid according to the information in local control blocks, this counts as a BMF valid READ hit. 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: • $\text{BMF Invalid Read Hits} = \text{True} + \text{False Invalid Read Hits}$ • $\text{BMF Read Hit\%} = \text{BMF Read\%} / \text{BMF Valid\%} * 100$ • $\text{BMF Invalid Read Hit\%} = \text{BMF Read Hit\%} - \text{BMF Read\%}$
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.

SPACED - Disk Space Report

The Disk Space Report displays capacity and disk space information for volumes. This report displays only those volumes that belong to storage groups specified with the Monitor III SGSPACE gatherer option. You can use this information to decide whether a certain volume provides sufficient free disk space for new allocation requests.

Together with the Storage Space Report, this report can help to make decisions for long-term disk space capacity planning.

Although the Disk Space Report is a single system report, the report combines the data collected from all systems within the sysplex. This allows you to gather the data for an SMS Storage Group only on one system of the sysplex so that the collection of redundant data can be avoided.

How to request this report

To request the Disk Space Report, select **1** from the Primary Menu, then select **11** from the Overview Report Selection Menu (shown in [Figure 6 on page 23](#)) or enter one of the following commands:

```
SPACED
SPD
```

In addition, you can invoke the Disk Space Report for all volumes of a selected storage group by using cursor-sensitive control from the Storage Space Report (SPACEG).

Contents of the report

RMF 3.1 Disk Space Report						Line 1 of 90	
Samples: 60	System: TRX2	Date: 04/25/2023	Time: 08.42.00	Range: 60	Sec		
Volume	Total (MB)	Free (MB)	Free (%)	Largest Ext (MB)	Storage Group		
SYSSD1	8120	2922	36.0	2922	DB2		
SYSSD3	8120	2291	28.2	2291	DB2		
SYSSD2	8120	2074	25.5	2074	DB2		
SYS0PE	8120	6326	77.9	6326	OMVSSYS		
SYSSM5	8120	1164	14.3	40	SMS		
SYSSM3	8120	1034	12.7	233	SMS		
SYSSM6	8120	1017	12.5	294	SMS		
SYSSM2	8120	1004	12.4	198	SMS		
SYSSM5	8120	982	12.1	62	SMS		
SYSSM4	8120	947	11.7	34	SMS		
SYSSM7	8120	728	9.0	139	SMS		

Figure 99. Disk Space Report

Field descriptions

Table 67. Fields in the Disk Space Report	
Field Heading	Meaning
Volume	Name of the volume belonging to a monitored storage group.
Total (MB)	Total amount of disk space (in megabytes) on the volume.
Free (MB)	Total amount of free disk space (in megabytes) on the volume.
Free (%)	Percentage of free disk space on the volume.
Largest Ext (MB)	Largest block (extent) in megabytes of unallocated disk space available on the volume.
Storage Group	Name of the storage group to which the volume belongs.

SPACEG - Storage Space Report

A storage group is a collection of storage volumes and attributes, defined by the storage administrator and treated as a single object storage hierarchy. The Storage Space Report allows you to keep track of disk space consumption on a storage group level. This report displays only those volumes that belong to storage groups specified with the Monitor III SGSPACE gatherer option.

From this report, you can see whether the system can provide sufficient disk space for new allocation requests. This report may also be useful for making decisions in long-term disk space capacity planning.

Although the Storage Space Report is a single system report, the report combines the data collected from all systems within the sysplex. This allows you to gather the data for an SMS Storage Group only on one system of the sysplex so that the collection of redundant data can be avoided.

How to request this report

To request the Storage Space report, select **1** from the Primary Menu, then select **10** from the Overview Report Selection Menu (shown in [Figure 6 on page 23](#)) or enter one of the following commands:

```
SPACEG
SPG
```

Contents of the report

RMF 3.1 Storage Space Report				
Samples: 30	System: TRX2	Date: 04/25/2023	Time: 08.42.00	Range: 30 Sec
SGroup	Total (MB)	Free (MB)	Free (%)	Volumes
*ALL	322116	165322	51.3	100 *
SGSMB	70382	35610	50.6	28 *
SGZFS	78503	39376	50.2	32 *
SMSCAT1	2707	398	14.7	1
SMS3390B	16240	8683	53.5	2
S1P01	18947	18802	99.2	3
S1P03	24360	23946	98.3	3
TSODA1	21654	1473	6.8	4
TSODA2	8120	411	5.1	1
TSODA3	5414	139	2.6	2
TSODA4	5414	86	1.6	2
USSFS	70375	36398	51.7	12

Figure 100. Storage Space Report

Cursor-sensitive control on a storage group name displays the Disk Space Report with all volumes of this storage group.

Field descriptions

Table 68. Fields in the Storage Space Report	
Field Heading	Meaning
SGroup	Name of the storage group connected to the system. The line showing *ALL in this column presents the accumulated values or average percentage values for all storage groups.
Total (MB)	Total amount of disk space (in megabytes) on all online volumes in the storage group.
Free (MB)	Total amount of free disk space (in megabytes) on all online volumes in the storage group.
Free (%)	Percentage of free disk space in the storage group.
Volumes	Number of volumes in the storage group. If at least one volume did not return any space information, the number is followed by an * .

STOR - Storage Delays Report

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

How to request this report

To request the Storage Delays report, select **3**, from the Primary Menu, then select **6** from the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)) or enter the following command:

```
STOR [job_class,service_class]
```

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

```
STOR T, TSOPRIME
```

Contents of the report

```

RMF 3.1  Storage Delays
Command ==>
Line 1 of 103
Scroll ==> CSR

Samples: 119      System: SYSF  Date: 04/25/2023  Time: 10.52.00  Range: 120  Sec

Jobname  C  Service  DLY  ----- % Delayed for -----  -- Working Set --
          C  Class   %    COMM  LOCL  SWAP  OUTR  OTHR  Central  Expanded

*MASTER* S  SYSTEM    0      0      0      0      0      0      3514
PCAUTH   S  SYSSTC    0      0      0      0      0      0      119
RASP     S  SYSTEM    0      0      0      0      0      0      284
TRACE    S  SYSSTC    0      0      0      0      0      0      271
DUMPSRV  S  SYSTEM    0      0      0      0      0      0      423
XCFAS    S  SYSTEM    0      0      0      0      0      0      8817
GRS       S  SYSTEM    0      0      0      0      0      0      8048
SMSPDSE  S  SYSTEM    0      0      0      0      0      0      1972
SMSVSAM  S  SYSTEM    0      0      0      0      0      0      3350
CONSOLE  S  SYSTEM    0      0      0      0      0      0      3303
WLM       S  SYSTEM    0      0      0      0      0      0      2603
ANTMAIN  S  SYSTEM    0      0      0      0      0      0      1162
ANTAS000 S  STCDEF    0      0      0      0      0      0      1194
DEVMAN   S  SYSTEM    0      0      0      0      0      0      178

```

Figure 101. 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 69. 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: A ASCH B Batch O OMVS S Started task T TSO ? Data is missing or invalid.
Service Class	The name of the service class that a specified job has been running in.

Table 69. Fields in the STOR Report (continued)	
Field Heading	Meaning
DLY %	Delay the waiting job (address space) is experiencing because of contention for storage during the range period, expressed as a percentage. $\text{DLY \%} = \frac{\text{\# Delay Samples}}{\text{\# Samples}} * 100$ 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.
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 dataspace 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 previously described, you can use the Utility to have the delay percentages in [Table 70 on page 161](#) shown in the Storage Delays report.

Table 70. 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.

Table 70. Additional Fields in the STOR Report (continued)

Field Heading	Meaning
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.
Service Class Period	The WLM service class period that has been assigned to this address space.

Report options

The STOR Report Options panel is similar to the Device Report Options panel. See [Figure 43 on page 75](#) for an example. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed. See [Figure 41 on page 73](#) for an example.

STORC - Common Storage Report

This report provides information about the use of common storage (CSA, ECSA, SQA, and ESQA) within a system.

The top section of the report provides overall system information about the use of common storage. For more information about the fields in this section, see [Table 71 on page 163](#).

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 72 on page 164](#).

Note:

1. The report can be incomplete for some jobs, this will be indicated by messages ERB617I, ERB618I, or ERB619I. They explain that CSA/ECSA or SQA/ESQA data needed for RMF reporting was not completely gathered. VSM common storage (CSA, ECSA, SQA, and ESQA) tracking was either not active or partially active since the job started.

The common storage data fields are reported in dark blue for those jobs that data was partially gathered.

If VSM common storage tracking was partially active, you can use the BREF command to select a range period when common storage data gathering was active.

If VSM common storage tracking was not active, contact your system programmer to activate VSM common storage tracking by issuing:

```
SET DIAG=01
```

The defaults in the SYS1.PARMLIB member DIAG01 are:

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

2. In the calculations used for this report, when CSA is converted to SQA but not allocated, the amount is still considered part of allocated CSA. Only when the converted CSA is allocated is it considered allocated SQA. Because CSA to SQA conversion can increase SQA to a value larger than defined at IPL, percent values of SQA can be greater than 100%. This also applies for ECSA to ESQA conversions.

How to request this report

To request the Common Storage report, select **3**, from the Primary Menu, then select **10** from the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)) or enter the following command:

```
STORC [job_class,service_class]
```

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

```
STORC T, TSOPRIME
```

Contents of the report

```

RMF 3.1   Common Storage                               Line 1 of 690
Command ===>                                         Scroll ===> CSR

Samples: 120      System: TRX1   Date: 04/25/2023   Time: 14.56.00   Range: 120   Sec

System Information
IPL Definitions
Peak Allocation Values
Average CSA to SQA Conversion
Average Use Summary
Available at End of Range

      ---- Percent ----      ----- Amount -----
      CSA ECSA SQA ESQA      CSA ECSA SQA ESQA
Peak Allocation Values      19   15  72   69      4456K 323M 1620K 140M
Average CSA to SQA Conversion      0    0      0      0      0    0
Average Use Summary      19   15  72   68      834K  49M 1163K   95M
Available at End of Range      81   85  28   32      3622K 274M 457K   45M

Unalloc CSA + SQA : 3460K   Unalloc RUCSA :  524K   Unalloc ERUCSA :   35M

Jobname  Act C  Service  ELAP  -- Percent Used -  ----- Amount Used -----
          ASID  Time   CSA ECSA SQA ESQA      CSA ECSA SQA ESQA
*STC                    5    3    6    6      229K  10M 92872 9079K
%MVS                    9   12  66   61      405K  38M 1072K   86M
%REMAIN                    0    0    0    0      0   5648  128  8368
*MASTER*  S SYSTEM  0001  20.2H  4    1    5    3      157K 2567K 75952 3757K
WLM        S SYSTEM  0011  20.2H  0    0    0    2      0    99K    0  2758K
NET        S SYSSTC  0095  20.1H  0    1    0    0      11448 2832K    0   824
XCFAS      S SYSTEM  0006  20.2H  0    0    0    1      0   1216   208 1135K

```

Figure 102. STORC Report

There is no graphic version of this report available.

Field descriptions

Table 71. Fields in the STORC Report - System Information Section	
Field Heading	Meaning
IPL Definitions	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 since IPL. 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 71. 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 range. 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 Important: • CSA/ECSA include the conversion values (SQA-to-CSA expansion) and do not include RUCSA/ERUCSA • SQA/ESQA include the conversion values
Available at End of Range	The amount and percentage of common storage (CSA, ECSA, SQA, and ESQA) available for allocation at the end of the specified range.
Unalloc CSA + SQA	The amount of unallocated common area below 16M (CSA + SQA) associated with a system key (0 - 7) available at the end of the specified range.
Unalloc RUCSA	The amount of unallocated common area below 16M (RUCSA) associated with a user key (8 - F) available at the end of the specified range.
Unalloc ERUCSA	The amount of unallocated common area above 16M (ERUCSA) associated with a user key (8 - F) available at the end of the specified range.

Table 72. 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 requested 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.
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.

Table 72. Fields in the STORC Report - Job Information Section (continued)	
Field Heading	Meaning
C	A one-character abbreviation for the job class as follows: A ASCH B Batch O OMVS S Started task T TSO ? Data is missing or invalid.
Service Class	The name of the service class that a specified job has been running in.
ASID	The unique system-assigned identifier for the address space in which the job is running.
ELAP time	The time elapsed between a job's starting time and end of range time. If a job's ELAP time is greater than 999.9 days, RMF displays asterisks (**.*D) in this field instead of the elapsed time. The field is blank when the start time of a job is zero. For example, the field is blank for batch initiators. 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 previously described, you can use the Utility to have the following values shown.

Table 73. 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.
Service Class Period	The WLM service class period that has been assigned to address space.

Report options




Figure 103. 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 41 on page 73](#) for an example.

STORCR - Common Storage Remaining Report

The Common Storage Remaining report identifies jobs that have ended but have not released all of their allocated common storage (CSA, ECSA, SQA, and ESQA) since IPL.

The report is a snapshot of the system at the end of the specified report interval.

The jobs are sorted by descending storage percentage; that is, for each job with the maximum of the four common storage percentages, the job with the highest maximum percentage is reported first. Cursor-sensitive control is not active on the report.

How to request this report

To request the Common Storage Remaining report, select **3** from the Primary Menu, then select **11** from the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)) or enter following command:

```
STORCR
```


Contents of the report

```

RMF 3.1 Common Storage Remaining Line 1 of 30
Command ==> Scroll ==> CSR
Samples: 119 System: SYSF Date: 04/25/2023 Time: 17.52.00 Range: 120 Sec

Jobname ID Job Ended Date Time Amount of Common Storage
Not Released at End of Job
CSA ECSA SQA ESQA
%REMAIN 0 1109K 1504 23952
CANFDSST STC04642 04/25/2023 15.59.39 0 513K 96 128
CANFDSST STC04661 04/25/2023 04.40.09 0 506K 0 128
CATALOG 04/25/2023 15.01.02 0 0 768 8552
CATALOG 04/25/2023 10.01.03 0 0 512 6624
IRRDPTAB STC04335 04/25/2023 13.36.06 0 86632 0 0
CATALOG 04/25/2023 13.35.32 0 0 128 1360
STARTMVS 04/25/2023 13.47.17 0 0 0 5336
SMFDUMP STC04334 04/25/2023 13.36.36 0 2416 0 0
BENK STC04637 04/25/2023 16.15.44 0 2304 0 0
TAGE TSU04619 04/25/2023 05.04.15 0 1024 0 0
BPXAS STC04881 04/25/2023 09.33.13 0 0 0 96
BPXAS STC04865 04/25/2023 07.23.36 0 0 0 96

```

Figure 104. STORCR Report

There is no graphic version of this report available.

Note: If you request the STORCR report and VSM tracking was stopped between IPL and end of range, an empty STORCR report can be displayed with the message 'No ended jobs found'.

Field descriptions

Table 74. Fields in the STORCR Report	
Field Heading	Meaning
Jobname	Name of the job that ended but did not release all of its common storage (CSA, ECSA, SQA, or ESQA). The %REMAIN summary line is always the first reported line and cannot be excluded from the report.
ID	The ID for the ended job. The ID can be used to identify the job in the system logs. The ID field is blank for some system-related address spaces.
Job Ended Date	The date the job ended. You can use the language options panel to customize the date format.
Job Ended Time	The time the job ended. You can use the language options panel to customize the time format.
Amount of Common Storage Not Released at End of Job	The amount of allocated common storage (CSA, ECSA, SQA, and ESQA) that was not released when the job ended. If data gathering was stopped between IPL and end of range, the field can be reported in dark blue.

There are no report options to specify for the STORCR report.

STORF - Storage Frames Report

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

How to request this report

To request the Storage Frames report, select **3** from the Primary Menu, then select **7** from the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)) or enter the following command:

```
STORF [job_class,service_class]
```

Contents of the report

```

RMF 3.1  Storage Frames                               Line 1 of 103
Command ===>                                         Scroll ===> CSR
Samples: 9      System: SYSF  Date: 04/25/2023  Time: 08.44.00  Range: 30  Sec

Jobname  Service  -- Frame Occup.-- - Active Frames - AUX  PGIN
C Class  Cr TOTAL  ACTV  IDLE  WSET FIXED  DIV SLOTS RATE
INIT     S SYSSTC   197K    0  197K    0  197K    0    53    0
IOSAS    S SYSTEM   133K   133K    0  133K   755    0    49    0
THR64GMD B JESLOW5   115K   115K    0  115K  115K    0    43    0
STGTHR32 B JESLOW5  77390  77390    0  77390   367    0  62388    0
STGTHR33 B JESLOW5  77390  77390    0  77390   367    0  59546    0
STGTHR31 B JESLOW5  77390  77390    0  77390   367    0  60125    0
STGTHR34 B JESLOW5  77389  77389    0  77389   367    0  64878    0
INIT     S SYSSTC   66778    0  66778    0  66627    0    24    0
THR64FMC B JESLOW5  65927  65927    0  65927  65720    0    60    0
THR64AMB B JESLOW5  65925  65925    0  65925  65720    0    62    0
THR64EMC B JESLOW5  65925  65925    0  65925  65720    0    62    0
THR64DMC B JESLOW5  65925  65925    0  65925  65720    0    62    0
THR64BMC B JESLOW5  65847  65847    0  65847  65612    0    18    0
THR64CMC B JESLOW5  65846  65846    0  65846  65612    0    18    0

```

Figure 105. STORF Report

There is no graphic version of this report available.

Field descriptions

Table 75. 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: A ASCH B Batch O OMVS S Started task T TSO ? Data is missing or invalid.
Service Class	The name of the service class that a specified job has been running in.
Cr	An S in this column indicates that WLM managed the address space as <i>storage critical</i> during the reporting interval.

Table 75. Fields in the STORF Report (continued)	
Field Heading	Meaning
Frame Occup.	This field shows the frame occupancy 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 storage frames the job used during the report interval. $\text{ACTV} = \frac{\sum \text{Central Storage Frames}}{\# \text{ Samples}}$ IDLE The average number of frames held by the job while it was idle. This value represents the average number of central storage frames the jobs used when it was idle during the report interval. $\text{IDLE} = \frac{\sum \text{Central Storage Idle Frames}}{\# \text{ Samples}}$
Active Frames	This field is broken into three categories. WSET See Table 69 on page 160 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. $\text{FIXED} = \frac{\sum \text{Fixed Frames}}{\# \text{ Samples}}$ DIV The number of central storage frames used by DIV. This count is accumulated only for jobs, not for service classes. A service class displayed on the report may not be the same as where the job was running when the DIV sample was taken. $\text{DIV} = \frac{\sum \text{DIV Frames}}{\# \text{ DIV Samples}}$
AUX SLOTS	Number of auxiliary slots for each address space.
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.

Report options

The STORF Report Options panel is similar to the DEV Report Options panel. See [Figure 43 on page 75](#) for an example. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed. See [Figure 41 on page 73](#) for an example.

STORM - Storage Memory Objects Report

This report provides information about the use of memory objects within the system. A memory object is a contiguous range of virtual addresses that is allocated by jobs in units of megabytes on a megabyte boundary.

The top section of the report provides overall system information about memory objects. The bottom section provides job-related information about the use of memory objects for jobs active during the specified report interval.

How to request this report

To request the Storage Memory Objects report, select a **3** from the Primary Menu, then select **7A** from the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)) or enter the following command:

```
STORM [job_class,service_class]
```

Contents of the report

```

Command ==>          RMF 3.1   Storage Memory Objects          Line 1 of 8
                               Scroll ==> CSR

Samples: 60          System: SYSF   Date: 04/25/2023   Time: 10.19.00   Range: 60   Sec

----MemObj-----  ---Frames---  -1MB Frames-  --1MB Fixed--  --2GB Fixed--
Fixed 1M          1  Shared   234    Total   1118    Max     512    Max     2
Fixed 2G          1  Common 19038   %Used   2.0    Common   0    %Used   50.0
Shared           78  %Used  10.1   %UsedP   0.9    %Used   2.0
Common           78

-----
Jobname  C  Service  --Memory Objects-  --1M Frames-  2G-Fr  -----Bytes-----
          C  Class   ASID  Total  Comm  Shr Fixed Pgable Fixed  Total  Comm  Shr
TCPIP    S  SYSSTC   0066   316    4    0    0    0    0  11.3G 2583M    0
JES2     S  SYSSTC   0052    14    1    0    0    0    0   850M 1024K    0
IOSAS    S  SYSTEM   0025    9    0    0    0    0    0   13.0M    0    0
COMM641M B  BATCHLOW 0057    9    1    0    30   0    0   43.0M 30.0M    0
PGBLE1MP B  BATCHLOW 0058    9    0    0    0    5    0   18.0M    0    0
PGBLE1MS B  BATCHLOW 0059    9    0    1    0    0    0   2061M    0 2048M
STOR2GN1 B  BATCHLOW 0060    9    0    0    0    0    0   2061M    0    0
GPMSERVE S  GPMSERVE 0101    6    0    0    0    0    0   10.0M    0    0

```

Figure 106. STORM Report

There is no graphic version of this report available.

Field descriptions

Table 76. Fields in the STORM Report - System Summary Section	
Field Heading	Meaning
MemObj	Fixed 1M Average number of fixed memory objects that are allocated in the system and can be backed in 1 MB frames. Fixed 2G Average number of fixed memory objects that are allocated in the system and are backed in 2 GB frames. Shared Average number of memory objects allocated in the high virtual shared storage of the system. Common Average total number of memory objects allocated in the high virtual common storage of the system. This value includes the memory objects that cannot be attributed to an address space. These fields include memory objects that can be backed in 1 MB frames (if Enhanced DAT Facility is installed), or 2 GB frames (if Enhanced DAT Facility 2 is installed).

Table 76. Fields in the STORM Report - System Summary Section (continued)	
Field Heading	Meaning
Frames	Shared Average number of high virtual shared storage pages backed in central storage (in units of 4 KB). Common Average number of high virtual common storage pages backed in central storage (in units of 4 KB). %Used Percentage of high virtual common storage used by the system.
1 MB Frames	Total Number of 1 MB frames in central storage. %Used Percentage of 1 MB frames in central storage that are used. %UsedP Percentage of 1 MB frames in central storage that are used by pageable 1 MB pages. These fields are only available if Enhanced DAT Facility is installed.
1 MB Fixed	Max Maximum number of 1 MB frames that can be used by fixed 1 MB pages. This value is specified in the LFAREA parameter. Common Average number of 1 MB high virtual common memory pages fixed in central storage. This value includes the pages that cannot be attributed to an address space or have not been freed during address space termination. %Used Percentage of the maximum number of 1 MB frames that are used by fixed 1 MB pages. These fields are only available if Enhanced DAT Facility is installed.
2 GB Fixed	Max Maximum number of 2 GB frames that can be used by fixed 2 GB pages. This value is specified in the LFAREA parameter. %Used Percentage of the maximum number of 2 GB frames that are used by fixed memory objects. These fields are only available if Enhanced DAT Facility 2 is installed.

Table 77. Fields in the STORM Report - Address Space Section	
Field Heading	Meaning
Jobname	The name of a job using memory objects
C	A one-character abbreviation for the job class as follows: A ASCH B Batch O OMVS S Started task T TSO ? Data is missing or invalid.
Service Class	The name of the service class that a specified job has been running in.
ASID	The decimal identifier of the address space in which the job is running.

Table 77. Fields in the STORM Report - Address Space Section (continued)	
Field Heading	Meaning
Memory Objects	Total Average number of memory objects allocated by this address space. Comm Average number of high virtual common memory objects allocated by this address space. Shr Average number of high virtual shared memory objects allocated by this address space.
1M Frames	Fixed Average number of 1 MB frames in the Large Frame Area owned by this address space. Frames that are used to satisfy 4 KB space requests on a constrained system are not included. Pgable Average number of 1 MB frames that are used by pageable and DREF memory objects owned by this address space. Pageable memory objects that have been fixed after allocation, are also included. Frames that are either used by common 1 MB pages or to satisfy 4 KB space requests on a constrained system are not included. These fields are only available if Enhanced DAT Facility is installed.
2G-Fr	Fixed Average number of 2 GB frames in the Large Frame Area owned by this address space. These fields are only available if Enhanced DAT Facility 2 is installed.
Bytes	Total Average amount of storage allocated from high virtual memory in memory objects with this address space as the owner. Comm Average amount of high virtual common storage allocated with this address space as the owner. Shr Average amount of storage allocated in shared memory objects with this address space as the owner.

Report options

The STORM Report Options panel is similar to the Device Report Options panel. See [Figure 43 on page 75](#) for an example. If you select YES for Jobs on the STORM Report Options panel, the Job Selection/Exclusion panel is displayed. See [Figure 41 on page 73](#) for an example.

STORR - Storage Resource Delays Report

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

How to request this report

To request the Storage Resource Delays report, select **3** from the Primary Menu, then select **8** from the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)) or enter the following command:

```
STORR
```

Contents of the report

Command ==>										RMF 3.1 Storage Resource Delays					Line 1 of 7 Scroll ==> PAGE				
Samples: 9			System: SYSF		Date: 04/25/2023			Time: 08.44.00		Range: 30		Sec							
----- Central Storage Summary -----																			
----- % Frames -----										Frames		System							
NUC	SQA	CSA	LPA	ACTV	IDLE	AVAIL	SHR		Online	UIC									
0	3	0	0	38	14	44	1		6291706	65534									
Total SQA + ESQA Overflow:					28K														
----- Page/Swap Activity -----																			
Volume	DEV		CU		ACT	CON	DSC	PND	Pend	SPACE	- AVG	Active	Users-						
Serial	Type		Type	PAV	%	%	%	%	Reasons	TYPE	TOTL	LOCL	SWAP	COMM					
S53PG7	33903		2105		4	0	0	0	None	LOCL	0.0	0.0	0.0	0.0					
S53PG8	33903		2105		4	0	0	0	None	LOCL	0.0	0.0	0.0	0.0					
S53PG9	33903		2105		4	0	0	0	None	LOCL	0.0	0.0	0.0	0.0					
S53PG3	33903		2105		4	0	0	0	None	LOCL	0.0	0.0	0.0	0.0					
S53PG2	33903		2105		6	0	0	0	None	LOCL	0.0	0.0	0.0	0.0					
S53PG1	33903		2105		9	0	0	0	None	COMM	0.0	0.0	0.0	0.0					

Figure 107. STORR Report

The report has two sections.

Central Storage Summary

This section includes general information about the use of central storage.

Page/Swap Activity

This section includes information about page/swap activity and paging delays.

If RMF cannot provide data in the Page/Swap Activity section for ACT %, CON %, DSC %, and PND %, dashes appear in these fields.

If RMF is unable to obtain valid hardware data for a sub-channel, it prints dashes (---) instead of DLY DB% and DLY CU%.

The graphic form of this report shows the average number of active users connected (CON), disconnected (DSC), pending (PND), and delayed for LOCL, SWAP, and COMM.

Field descriptions

Table 78. Fields in the STORR Report - Central Storage Summary Section	
Field Heading	Meaning
% 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 storage frames during the report interval. The categories are as follows: NUC Percentage of central storage frames allocated to the nucleus (NUC). SQA Percentage of central storage frames allocated to the system queue area (SQA). CSA Percentage of central storage frames allocated to the common storage area (CSA) and to the restricted use common service area (RUCSA). LPA Percentage of central storage frames allocated to the link pack area (LPA). ACTV 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. IDLE 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. AVAIL Percentage of available central frames. SHR Percentage of shared frames in central storage.
Frames Online	The number of central storage frames, excluding read-only and 2G LFAREA frames. Nucleus frames are included in this value.
System UIC	The system's unreferenced interval count indicates storage contention.
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. $\text{Total SQA + ESQA Overflow} = \frac{\text{Total Overflow}}{\text{\# Samples}}$ Total overflow Total overflow above and below 16M

Table 79. 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 report interval. If the number has been changed during the report interval, it is followed by an '**'. If the device is a HyperPAV base device, the number is followed by an 'H'. The value is the average number of HyperPAV volumes (base and alias) for that interval. $\text{Average \# of HPAV devices} = \frac{\text{Accumulated \# of HPAV devices}}{\text{Number of Samples}}$ This field appears only for parallel access volumes.

Table 79. Fields in the STORR Report - Page/Swap Activity Section (continued)	
Field Heading	Meaning
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: $\text{Active Time} = \text{PND Time} + \text{CON Time} + \text{DSC Time}$ 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.
CON %	The percent connect time: $\text{CON \%} = \frac{\text{Connect Time}}{\text{Range Time}} * 100$
DSC %	The percent disconnect time: $\text{DSC \%} = \frac{\text{Disconnect Time}}{\text{Range Time}} * 100$
PND %	The percent pending time: $\text{PND \%} = \frac{\text{Pending Time}}{\text{Range Time}} * 100$ Note: 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. $\text{DLY DB\%} = \frac{\text{Accumulated DB Delay Time}}{\text{Range Time}} * 100$ 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. $\text{DLY CMR\%} = \frac{\text{Accumulated Command Response Delay Time}}{\text{Range Time}} * 100$ Note: If either hardware data or volume related percentages are not available, this field is blank.

Table 79. Fields in the STORR Report - Page/Swap Activity Section (continued)

Field Heading	Meaning
SPACE TYPE	The space types for which the percentage of the volume's delay is reported. The types appear in the following order: LOCL User private area COMM Common area PLPA Pageable link pack area
AVG Active Users	The average number of jobs waiting for the data set. $\text{AVG Active Users} = \frac{\sum \text{Waiting Jobs}}{\# \text{ Samples}}$ 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 previously described, you can use the Utility to have the following delay percentages shown.

Table 80. Additional Fields in the STORR Report

Field Heading	Meaning
Percentage of using	The percentage of time the device was found being used by an address space.
Percentage of DLY-DB	The percentage of time during the report interval when the channel subsystem measured I/O request delay because the device was busy.
Percentage of DLY-CUB	The percentage of time during the report interval when there is I/O request delay because the control unit was busy.
Percentage of DLY-DPB	The percentage of time during the report interval when there is I/O request delay because the ES/Connection Director port was busy.
Delay reason percentage	The percentage of time the device was delayed.

STORS - Storage Delay Summary Report

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

How to request this report

To request the Storage Delay Summary report, select **3** from the Primary Menu, then select **9** from the Resource Report Selection Menu (shown in [Figure 8 on page 24](#)) or enter the following command:

```
STORS workload_group | service_class | report_class
```

Parameter report_class designates either a report class or tenant report class.

Contents of the report

```

RMF 3.1  Storage Delay Summary                               Line 1 of 19
Command ==>                                                Scroll ==> CSR

Samples: 119      System: SYSF  Date: 04/25/2023  Time: 10.52.00  Range: 120  Sec

----- Central Storage Summary -----
----- % Frames -----
NUC  SQA  CSA  LPA  ACTV  IDLE  AVAIL  SHR  Online  UIC
  0    3    0    0   38   14   44    1  6291706  65534

Group  T  -- Users --  - Average Number Delayed For-  - Average Frames-  PGIN
          TOTL  ACTV  ANY COMM  LOCL  SWAP  OTR  OTHR  ACTV  IDLE  FIXED  RATE

BATCH   W    4    0    0    0    0    0    0    0  11325    0   352  0.0
BTCHDEF S    3    0    0    0    0    0    0    0  10998    0   298  0.0
OMVSKERN S    1    0    0    0    0    0    0    0    327    0    54  0.0
OMVS    W    2    0    0    0    0    0    0    0  13746    0   304  0.0
OE      S    2    0    0    0    0    0    0    0  13746    0   304  0.0
STC     W   16    0    0    0    0    0    0    0  19207   373  1153  0.0
GPMSEVE S    1    0    0    0    0    0    0    0   8666    0   110  0.0
STCDEF  S   15    0    0    0    0    0    0    0  10541   373  1043  0.0

```

Figure 108. STORS Report

The top section on the report provides overall system information and is the same as the Central Storage Summary section of the STORR report. The bottom section of the report provides summary lines for service classes, report classes, and workload groups.

A graphic report shows the average number of users delayed for COMM, LOCL, SWAP, OTR, and OTHR.

Field descriptions

Table 81. Fields in the STORS Report	
Field Heading	Meaning
Central Storage Summary	Fields in this section are described in Table 78 on page 174 .
Group	The name of the group, including: • Workload group names • Service class names • Report class names
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

Table 81. Fields in the STORS Report (continued)

Field Heading	Meaning
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 216 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 69 on page 160).
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 75 on page 168 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 previously described, you can use the Utility to have the delays in Table 82 on page 178 shown in the Storage Delay Summary report.

Table 82. 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

```

Command ==>
RMF STORS Report Options
Line 1 of 4
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 109. 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 tenant report classes as well as their periods (if you also specified YES for Period) are displayed. Otherwise, no report classes and tenant 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

SYSENQ - Sysplex Enqueue Delays Report

The SYSENQ report is similar to the ENQR report (see “ENQR - Enqueue Resource Delays Report” on page 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.

How to request this report

To request the SYSENQ report, select **S** on the Primary Menu, and then select **4** on the Sysplex Report menu (shown in Figure 5 on page 22), or enter the following command:

SYSENQ

Contents of the report




Figure 110. SYSENQ Report

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

Field descriptions

Table 83. 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. $\text{Delayed \%} = \frac{\text{\# Delay Samples}}{\text{\# Samples}} * 100$ 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 z/OS system name where the job is running on.
-- Delayed -- ST	The status indicates whether the waiting job wants exclusive (EW) or shared (SW) use of the resource.
-- Holding -- %	The percent of the range that a specific job was holding the resource while the named job was delayed. $\text{Holding \%} = \frac{\text{\# Holding Samples}}{\text{\# Samples}} * 100$ 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 z/OS 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

The System Information (SYSINFO) report presents an overview of the system, its workload, the average response time for a transaction in a specific service class, report class, or workload group, and the total number of jobs using resources or delayed for resources.

How to request this report

To request the System Information report, select **1** from the Primary Menu, then select **2** from the Overview Report menu (shown in [Figure 6 on page 23](#)) or enter the following command:

```
SYSINFO workload_group | service_class | report_class
```

Parameter report_class designates either a report class or tenant report class.

Contents of the report

```

Command ==>
RMF 3.1 System Information
Line 1 of 23
Scroll ==> CSR

Samples: 120      System: SYSE Date: 11/30/22 Time: 11.44.00 Range: 100 Sec

Partition: SYSE      3906 Model 786      Appl%: 19 Policy: BASEPOL
CPs Online: 3.0      Avg CPU Util%: 21    EAppl%: 20 Date: 11/30/22
CBPs Online: 3.0      Avg MVS Util%: 31    Appl% CBP: 21 Time: 12.10.01
IIPs Online: 1.0      Appl% IIP: 0.3

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      20 168 7      0.29 1.4 0.0 0.7 0.0 0.0 0.0 0.0 0.0 0.0
*TSO          1 0      0.28 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0
*BATCH        19 13 7      0.00 1.3 0.0 0.7 0.0 0.0 0.0 0.0 0.0
*STC          64 146 0      0.01 0.1 0.0 0.0 0.0 0.0 0.0 0.0 0.0
*ASCH         0 0      0.00 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0
*OMVS         4 0      0.01 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0
*ENCLAVE      50 4 N/A      N/A 0.0 N/A 0.0 N/A 0.0 N/A N/A N/A
SYSTEM        W 64 147 0 0.000 0.00 0.1 0.0 0.0 0.0 0.0 0.0 0.0
SYSOTHER      S 0 0 0.000 0.00 0.0 0.0 0.0 0.0 0.0 0.0 0.0
SYSSTC        S 50 120 0 0.000 0.00 0.0 0.0 0.0 0.0 0.0 0.0 0.0
SYSTEM        S 70 27 0 0.000 0.00 0.0 0.0 0.0 0.0 0.0 0.0 0.0
STCLOW        W 19 21 7 17.32 0.28 1.4 0.0 0.7 0.0 0.0 0.0 0.0
BRMFSVZ1      S 14 3 2 0.000 0.00 0.3 0.0 0.0 0.0 0.0 0.0 0.0
DISCRETN      S 22 13 5 0.000 0.00 1.1 0.0 0.7 0.0 0.0 0.0 0.0
TSOCLASS      S 1 0 17.32 0.28 0.0 0.0 0.0 0.0 0.0 0.0 0.0 0.0

```

Figure 111. SYSINFO Report

The SYSINFO report has two sections. The top section provides you with an overview of the system. It identifies the measured system, the policy name, the policy activation date and time. It also includes information about processor usage during the report interval. In an LPAR environment, the header contains an extra line showing the z/OS view of CPU utilization and the LPAR partition name in which the Monitor III data gatherer is running. For the different aspects of CPU utilization refer to “[CPU - CPU Activity report](#)” on page 347.

The bottom section summarizes information about the total system (*SYSTEM), job classes (*TSO, *BATCH, *STC, *ASCH, or *OMVS), enclaves (*ENCLAVE), workload groups, service classes, and report classes. The *SYSTEM summary line represents the system total values as summarized information from all other summary lines.

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

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

The proportion of the active users in each using or delay category indicates the proportion of the average response time that is spent in that category. The graphic form of this report shows the average number of active users for each type of delay.

Note:

1. Report class data lines contain information for the transaction response time (RESP Time) and transaction rate (TRANS/SEC) fields. The rest of the fields are blank.
2. The transaction response time (RESP Time) field in all summary data lines is also blank.
3. There is no graphic support for report class lines.

Field descriptions

Table 84. Fields in the SYSINFO Report	
Field Heading	Meaning
Partition	Partition name.
CPs Online	The number of general purpose processors (standard CPs) online during the range period.
CBPs Online	The number of zCBP processors online during the range period.
IIPs Online	The number of zIIPs online during the range period. If the LOADxx PROCVIEW CORE parameter is in effect, the reported value designates the number of online threads.
Processor	Processor family and model (or N/A — if model information is not available). If the processor does not support the Concurrent Processor Upgrade function, the hexadecimal version number is displayed.
Avg CPU Util%	The average utilization percentage for all general purpose processors (CPs) during the report interval (LPAR view of the CPU utilization): $\text{Avg CPU Util\%} = \frac{\text{Sum of LPAR CPU Times}}{\text{Sum of Online Times}} * 100$ The LPAR CPU Time for one general purpose processor is calculated depending on the status of the logical processor: Wait Completion NO LPAR CPU Time = PR/SM Dispatch Time Wait Completion YES LPAR CPU Time = PR/SM Dispatch Time - Wait Time Dedicated LPAR CPU Time = Online Time - Wait Time '***' indicates missing or invalid data.
Avg MVS Util%	z/OS view of CPU utilization which is the percentage of the time that the general purpose processors (CPs) were busy: $\text{Avg MVS Util\%} = \frac{\text{Time Range} - \text{Sum of Wait Times}}{\text{Time Range}} * 100$ The time range is the sum of the times the general purpose processors were online. With HiperDispatch mode active, it is the sum of the times the processors were online but not parked. For more information about the z/OS view of CPU utilization refer to “CPU - CPU Activity report” on page 347.
Appl%	Percentage of the maximum general purpose processor capacity used by all address spaces during the report interval. This value is divided by the number of logical processors or cores that have been active during this interval.
EAppl%	Percentage of the maximum general purpose processor capacity used by all address spaces and enclaves during the report interval. This value is divided by the number of logical processors or cores that have been active during this interval.
Appl% CBP	Percentage of the maximum zCBP capacity used by all address spaces during the report interval. This value is divided by the number of logical zCBP processors or cores that have been active during this interval.
Appl% IIP	Percentage of the maximum zIIP capacity used by all address spaces during the report interval. This value is divided by the number of logical zIIP processors or cores that have been active during this interval.
Policy Date Time	The name and the activation date and time of the service policy in effect during collection of the reported data. This, however, does not imply that the complete policy definition is shown on this report.

Table 84. 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
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 average number of total users. ACT The average number of active users. See the definition of Users/Active under “WFEX - Workflow/Exceptions Report” on page 216 for more details.
RESP Time	The average response time (in milliseconds) for all transactions that ended during the report interval. The response time value is the elapsed time (sum of the execution time and the queued times) for an average ended transaction. More than 99999 milliseconds are shown with • K - times one thousand (10³) • M - times one million (10⁶) • G - times one billion (10⁹) 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 long compared to the value of the report interval. To increase the accuracy of your data, try increasing the value of the report interval to a value higher than or equal to the response time.
TRANS /SEC	The number of transactions per second. When used with the number of active users in the report, this field gives you an overview of how fast the system can handle the amount of work for a given group. The number of completed transactions between cycles is accumulated for each sample. $\text{TRANS /SEC} = \frac{\text{Completed Transaction Count}}{\text{Range Time}}$ When the enclave transaction rate is greater than 0, the TRANS /SEC field displays in reverse turquoise color. When you select the TRANS /SEC field, the cursor-sensitive control displays a pop-up with the following information about the enclave transaction: • Enclave transaction average execution time • Enclave transaction rate value

Table 84. Fields in the SYSINFO Report (continued)	
Field Heading	Meaning
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.
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 previously, you can use the Utility to have the following values shown.

Table 85. Additional Fields in the SYSINFO Report	
Field Heading	Meaning
SYS AFCVC	The percentage of central storage frames the job used during the report interval.
SYS ADJVC	The average number of users experiencing delay when requesting service from JES.
SYS ADHVC	The average number of users experiencing delay when requesting service from HSM.
SYS ADXVC	The average number of users experiencing delay when requesting service from XCF.
SYS ADNVC	The average number of users experiencing delay because of an operator mount request.
SYS ADMVC	The average number of users experiencing delay because of an operator message request.
SYS CPUVC	Percentage of the maximum general purpose processor capacity spent on behalf of a class or group.
SYS EAPVC	Percentage of the maximum general purpose processor capacity consumed within a class or group (including enclave time).
SYS SRBVC	Percentage of the maximum general purpose processor capacity spent by SRB work on behalf of a class or group.
SYS TCBVC	Percentage of the maximum general purpose processor capacity used by non-enclave TCB work that executed within a class or group.
SYS IFAVC	Percentage of the maximum zAAP processor capacity used within a class or group.
SYS SUPVC	Percentage of the maximum zIIP processor capacity used within a class or group.

Table 85. Additional Fields in the SYSINFO Report (continued)	
Field Heading	Meaning
SYSCBPVC	Percentage of the maximum zCBP processor capacity used within a class or group.
SYSCPVC	Percentage of the maximum general purpose processor capacity used by non-enclave TCB work that executed within a class or group.
SYSIFCVC	Percentage of the maximum general purpose processor capacity used by zAAP eligible work that executed within a class or group.
SYSSUCVC	Percentage of the maximum general purpose processor capacity used by zIIP eligible work that executed within a class or group.
SYSVELVC	Execution velocity. This value is calculated as CPU using, divided by the sum of CPU using and total delays gathered by WLM. The delays gathered by WLM include CPU delay and storage delay only.
SYSPDPVC	CPU time, in seconds, that transactions of a class or group were running at a promoted dispatching priority during the report interval.
SYSCVAVC	Whether CPU reconfiguration changes occurred during the reporting interval (YES or NO).
SYSECTVC	Enclave transaction rate per second.
SYSECEVC	Average enclave transaction execution time in millisecond.
SYSCRIVC	An indication of whether WLM-managed the address space has one of the following values during the report interval: I Service Class Period implicitly CPU critical I* First period of Service Class implicitly CPU critical S Storage critical SC Both CPU and Storage critical SI Both implicitly CPU and storage critical (Service Class Period) SI* Both first period implicitly CPU and storage critical (Service Class)
SYSAIIVC	WLM batch initiator management is AI-infused or not. The values can be: YES or NO

Report options

The Report Options panel is exactly the same as for the SYSSUM report and STORS report, shown in Figure 109 on page 179. Selections made on either options panel affect all three reports.

SYSRG - Resource Group Activity Report

The Resource Group Activity (SYSRG) report provides service administrators and performance analysts with detailed information on the resource groups and tenant resource groups that are defined in the WLM policy. The report shows the actual processor and memory consumption per resource group for each system in the sysplex as well as the capacity limits defined for each resource group.

If work in a resource group is consuming resources above the specified maximum capacity, the system throttles the associated work to slow down the rate of resource consumption by capping the amount of service that can be consumed. If a minimum processing capacity is set for a resource group and the work in the resource group is not meeting its goals, the system attempts to provide the defined minimum amount of processor resource to that resource group.

The report allows you to directly compare the consumption values of a resource group with their defined capacity limits. This can help you to better understand which resource groups are potential candidates for resource group management by the system.

How to request this report

To request the Resource Group Activity report, select **S** from the Primary Menu, then a **1A** on the Sysplex Report menu, (shown in [Figure 5 on page 22](#)) or enter one of the following commands using the format:

```
SYSRG resource_group
```

The parameter resource_group designates either a resource group or tenant resource group.

Contents of the report

RMF 3.1			Resource Groups - SYSDPLEX					Line 1 of 30		
WLM Samples: 480			Systems: 3		Date: 04/25/2023		Time: 13.07.00		Range: 120 Sec	
Name	Type	System	---CPU consumption---			---CPU capacity---			---Memory---	
			#CPs	MSU	SU/sec	Min	Max	Unit	Usage	Limit
RGPIX1	RG	*ALL	0.05	5	3891	1.00	9999	#CPs		95P
		SYSE	0.05	5	3891				1860K	
TRGRMFN1	TRG	*ALL	3.37	696	290K		50000	MSU		1111G
		SYSD	1.06	218	90803				7636K	
		SYSE	0.91	188	78307				7596K	
		SYSF	1.40	290	121K				7768K	
TRGRMFN2	TRG	*ALL	1.15	238	98916		99999	SU/sec		1024M
		SYSD	0.65	134	55893				5948K	
		SYSE	0.02	4	1462				5904K	
		SYSF	0.48	100	41562				7472K	
TRGRMFN3	TRG	*ALL	1.04	216	89775		3.33	#CPs		3333G
		SYSD	0.46	94	39236				1992K	
		SYSE	0.31	64	26472				1960K	
		SYSF	0.28	58	24068				2676K	

Figure 112. SYSRG Resource Group Activity report

The report shows all defined resource groups and tenant resource groups with their processor and memory limits and their actual consumption values at system and sysplex granularity.

Cursor-sensitive control on the SYSRG report

Using cursor-sensitive control on the resource group name in the table leads to a pop-up panel with a detailed breakdown of processor consumption data on service class or tenant report class level. The content of the pop-up panel depends on the resource group type.

Resource Group Activity Details			
Resource Group Name: RGPIX1			
Description: Resource Group RGPIX1			
Specialty Processor Consumption Included: No			
Definition of CPU capacity limits: #CPs			
Minimum CPU capacity: 0.03			
Maximum CPU capacity: 9999			
Service	----- CPU consumption -----		
Class	#CPs	MSU	SU/sec
*ALL	0.05	5	3891
WKLDSVC1	0.05	5	3891

Figure 113. Pop-up panel for a resource group

Tenant Resource Group Activity Details

Tenant Resource Group Name: TRGRMFN3
 Description: Tenant Resource Group for PG1

Specialty Processor Consumption Included: Yes
 Definition of CPU capacity limit: #CPs
 Maximum CPU capacity: 3.33

Tenant Report ----- CPU consumption -----

Class	#CPs	MSU	SU/sec
*ALL	1.04	216	89775
PG1TRCN3	0.74	154	63878
PG1TRCN4	0.30	62	25897

Figure 114. Pop-up panel for a tenant group

Field descriptions

Table 86. Fields in the SYSRG Report	
Field Heading	Meaning
Name	Name of the resource group or tenant resource group.
Type	Type of resource group being reported on. RG Resource group TRG Tenant resource group
System	Name of a MVS system that belongs to the sysplex. If the MVS system name is not available, the four character SMF system identifier is shown.
CPU consumption	Processor consumption in terms of #CPs number of general purpose processors MSU millions of service units per hour SU/sec unweighted CPU and SRB service units per second The scope of the reported value can be a resource group a particular system in the sysplex a service class that is associated with a resource group a tenant report class that is associated with a tenant resource group
CPU capacity	Resource group capacity limits. Min Minimum amount of service that the resource group should receive if demand exists. Max Maximum amount of service that the resource group should be allowed to consume. Unit The method of how the resource group's capacity is defined: SU/sec: in unweighted CPU and SRB service units per second across the sysplex %LPAR: as percentage of the LPAR share on each system in the sysplex #CPs: as number of general purpose processors on each system in the sysplex MSU: as MSU/h capacity across the sysplex
Memory Usage	Amount of memory in bytes used by the resource group on this system.
Memory Limit	Resource group memory limit in bytes for each system in the sysplex. Memory consumption of work running in address spaces associated with the resource group is limited to this size.
Description	Description of resource group or tenant resource group.

Table 86. Fields in the SYSRG Report (continued)	
Field Heading	Meaning
Specialty Processor Consumption Included	Yes Indicates that WLM capping considers service units consumed on general purpose and specialty processors. In such a case, the reported CPU consumption reflects the actual general purpose and specialty processor consumption. No indicates that only general purpose processor consumption is reported.
Service Class	Name of a service class that is associated with this resource group.
Tenant Report Class	Name of a tenant report class that is associated with this tenant resource group.

Report options

```

RMF Resource Group Activity Report Options: SYSRG

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

Name      ==> ALL      ALL or one of the available resource groups below
Type      ==> ALL      Resource group type (RG, TRG or ALL)
Detail    ==> NO       Show single system data (YES or NO)
Inactive  ==> NO       Show inactive resource groups (YES or NO)

                        Available Resource Groups
RGROUP01  REGTSO      TRGCLD      TRGRP01      TRGRP02      TRGRP03      TRGRP04
TRGRP05   TRGRP06

```

Figure 115. Pop-up panel for resource group activity report options

Name

Either **ALL** or the name of one of the resource groups being available in the sysplex as shown in section **Available Resource Groups**

Type

To select a specific resource group type in the SYSRG report, you can request

ALL

to show all resource groups as well as all tenant resource groups

RG

to show all resource groups

TRG

to show all tenant resource groups

Detail

With this option, you can select the level of detail in the SYSRG report:

YES

The report contains data for the sysplex and all single systems.

NO

The report contains data for the sysplex only.

Inactive

Specification about display of inactive resource groups in the SYSRG report:

YES

Include data lines in the report, even if the lines do not contain any activity data.

NO

Do not include data lines, if they do not contain any activity data.

Available Resource Groups

The list of all resource groups which are currently defined in the WLM Policy.

SYSRTD - Response Time Distribution Report

The Response Time Distribution (SYSRTD) report enables the service administrator and performance analyst to analyze the distribution of response time to see whether a response time goal was met and, if not, how close it came to failing. This report can also be used to "fine-tune" response time goals.

Use the bottom part of the report to see bottlenecks related to a specific system. Then you can use single-system reports for more detailed analysis.

How to request this report

To request the Response Time Distribution report, select **S** from the Primary Menu, then a **2** on the Sysplex Report menu, (shown in [Figure 5 on page 22](#)) or enter one of the following commands:

```
SYSRTD service_class, period
SYSRTD report_class, period
```

Parameter report_class designates either a report class or tenant report class.

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

```
SYSRTD POSMULTI, 1
```

Contents of the report

```

RMF 3.1 Response Time - RMFPLEX1
Command ==>
Line 1 of 5
Scroll ==> HALF

WLM Samples: 100 Systems: 8 Date: 04/25/2023 Time: 10.31.40 Range: 100 sec

Class : POSMULTI Period: 1 Goal: 0.500 sec for 90%

Trx # 50% 60% 70% 80% 90% 100% 110% 120% 130% 140% 150% 200% 400% >400%
Trx % 2692 50 169 246 167 46 22 17 8 7 1 18 1178 298
Trx % 57.1 1.00 3.30 4.70 3.20 0.90 0.40 0.30 0.20 0.10 0.00 0.30 22.7 5.70

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

*ALL 208.0 379.0 587.0 124.7 9 4 54
MVS1 all 61.0 311.0 372.0 84.3 7 8 74
MVS2 all 149.0 984.0 1134 40.4 2 0 42
MVS3 part 0.000 0.000 0.000 0.000 0 0 0
MVS5 none

```

Figure 116. 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 table shows the distribution of response time for all systems in a sysplex which have data available in the selected period.
- A table shows how each system contributed to the overall response time.

Depending on the goal of the period, different data may be available. A report for a service class period with a response time goal is shown in the previous figure. Here, the response time distribution is shown in the table 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 table is shown, and the top section of the report only shows the service class name and the service class period.


```

RMF 3.1 Response Time - RMFPLEX1           Line 1 of 5
Command ===>                               Scroll ===> HALF
WLM Samples: 100      Systems: 7  Date: 04/25/2023 Time: 10.31.40 Range: 100   sec
Class: ALLBATCH      Period: 2

No response time distribution values available,
because there was no response time goal specified
for this service class period.

System  Data  --Avg. Resp. Time--  Trx  --Subsystem Data--  --Exec Data--
          WAIT  EXECUT ACTUAL  Rate  Actv  Ready Delay  Ex Vel  Delay
*ALL                1974K 1734K 3708K  39.2                45    38
MVS6    all         5381  482K  487K  11.8                82    16
MVS7    all         217K 5328K 5544K   0.1                91     1
MVS8    part        268K 2970K 3240K   3.4                86    45
MVS9    all        391.0  1187  1227   1.9                79     3

```

Figure 117. SYSRTD Report - Without Response Time Data

The response time distribution table is not shown for heterogeneous report class periods. Please refer to [“Performance data” on page 195](#) for an explanation of homogeneous and heterogeneous report class periods.

How to read the response time distribution table

To provide a picture of how a performance group was performing, response time distributions are provided for both service classes and report classes.

These distributions consist of 14 buckets of information.

The header contains the value of the particular bucket, which is a percentage of the specified goal. One bucket always maps exactly to the specified goal, with a value of 100%.

Each bucket has the number of transactions that completed in the amount of time that is represented by that bucket. This number is reported in the first row of the response time distribution table, the second row of the table shows the percentage of transactions that are contained in each bucket.

In Figure 116 on page 190, each of the 14 buckets represents a percentage of the specified 0.5-second goal. For instance, bucket 2 represents all transactions that completed in 50% to 60% of the goal, or 250-300 milliseconds, while bucket 10 contains the number of transactions that completed in 130% to 140% of the goal, or 650-700 milliseconds. Notice that bucket 6 falls exactly on the goal (100% of goal, or 0.5 seconds). This bucket captures all those transactions that complete in 400-500 milliseconds.

The two end buckets (buckets 1 and 14) have special meaning. Bucket 1 (labeled 50%) contains the total number of transactions that completed in up to 50% of the goal. Bucket 14 (labeled >400%) contains the number of transactions that completed in greater than 4 times the goal.

Scrollable part of report

The bottom section of this report is scrollable. It shows a list of all systems that have workload activity data gathered for the service class period during the report interval.

The first row in the scrollable area is a summary line. To indicate this, the *System* column displays the word ***ALL**. The *Data* column remains empty, and all other columns contain the respective time value or delay percentage for the sysplex. The *Response time* columns for example show the same values as the row on the Sysplex Summary report for that service class period.

For each system, a row is shown where important response time data is provided. This information is intended to assist in tracking possible bottlenecks down to a specific system, where the analysis can be continued using the detailed reports for single systems.

Data reported

The report is for one service class period. Depending on the type of service class, the different parts of the report may contain data or remain empty. Here is a list of what kind of data you can expect under which circumstances:

- Response Time Distribution
Available only if a response time goal was specified
- Response Time Data
Almost always available (possibly not for STC)
- Subsystem Data
Available only for a subsystem transaction class
- Execution Data
Available only if it is NOT a subsystem transaction class

Cursor-sensitive control on the SYSRTD Report

In the non-scrollable area on the top of the report, which may show the response time distribution table or a message that the data for that table 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 63 on page 106) 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 87. Fields in the SYSRTD Report	
Field Heading	Meaning
50%, 60%, 70%, 80%, 90%, 100%, 110%, 120% , 130%, 140%, 150%, 200%, 400%, >400%	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 (<= 50%) covers the gap from 0 to half the goal • buckets 2 to 11 (<= 60% to <= 150%) cover the gap between half the goal to 1.5 times the goal evenly divided • bucket 12 (<= 200%) covers two times the goal • bucket 13 (<= 400%) covers four times the goal • bucket 14 (> 400%) covers the gap from four times the goal to infinity
Trx #	Number of transactions in this bucket.
Trx %	Percentage of transactions associated to the bucket.
Response time	Response time distribution. For a description, refer to “How to read the response time distribution table” on page 191.
System	The four-character SMF system identifier.

Table 87. Fields in the SYSRTD Report (continued)	
Field Heading	Meaning
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 193 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 205 for a description. The difference to the Work Manager Delay report is, that here the begin-to-end and the execution phase are combined within one row.
Execution Data Ex Vel (=Execution Velocity)	This is the same as the <i>Exec Vel - Actual</i> on the Sysplex Summary report with the only difference that the value here is only calculated for a single system (except *ALL). Refer to “SYSSUM - Sysplex Summary report” on page 193 for a description.
Execution Data Delay	This is the general execution delay used for the execution velocity calculation. Note that in a service class more than one transaction can be delayed at the same point of time. For example, if two transactions on average are delayed each time WLM takes a measurement sample, a value of 200 will be displayed.

Report options

The Report Options panel shows the RMF default options. It is the same as for the Group Response Time (GROUP) report, shown in [Figure 64 on page 112](#), only the header line is different in showing the respective report name.

Note: The list of available service classes will be shown only if one of the sysplex reports SYSSUM, SYSRTD, or SYSWKM has been displayed at least once.

SYSSUM - Sysplex Summary report

The Sysplex Summary (SYSSUM) report allows the service administrator and performance analyst to see at a glance whether service goals are being satisfied by:

- Showing a performance status line showing the performance status of the sysplex covering a time range of up to 80 refresh intervals.
- Showing the actual throughput being achieved by all workloads on one report
- Displaying goals not met in red or yellow
- Calculating the Performance Index for each service class period

The report provides an overview of workload groups, service classes, service class periods, report classes, and report class periods. It allows "summarizing" of actual values for every group using threshold values, and includes a goal versus actual comparison for each period. You can compare different goals by using the performance index.

Furthermore, the response time for all groups is calculated independently of any specified goals, and a transaction rate is provided to enable you to evaluate the importance of the figures shown.

To facilitate detection of goals that were not met, the line of that service class period, as well as the related workload group and service class, is displayed in red or yellow.

Finally, options are available to select specific groups, or limit the report to groups that have exceeded their goal by a certain amount.

How to request this report

To request the Sysplex Summary report, select **S** from the Primary Menu, then a **1** on the Sysplex Report menu (shown in [Figure 5 on page 22](#)), or enter the following command:

```
SYSSUM workload_group | service_class | report_class
```

Parameter report_class designates either a report class or tenant report class.

Contents of the report




Figure 118. SYSSUM Report - GO Mode

The report can be logically broken into three sections.

The Performance Status line

In GO mode, a colored status row at the top of the screen gives an overview over the sysplex during the last ranges. For every range, a one-character field will be added to the right-hand side of the performance status line. For each range that has been reported, the one-character field is marked in one of the following ways:

■ (green)

If all goals have been met during that interval (the performance index is less than, or equal to, 1 for all periods)

■ (yellow)

If a warning level was reached during that interval (the performance index was greater than 1 for periods with an importance equal to 3, 4, or 5)

X (red)

If goals have been exceeded seriously during that interval (the performance index was greater than 1 for periods with an importance equal to 1 or 2)

Blank


If you changed the mode from GO to STOP during some intervals

The **Refresh** value that you can specify on the Session Options panel will define the how often the status line will be updated.

If you define a refresh value that is less than the MINTIME, the refresh value will be reset to equal the MINTIME.

Example

For example, if you specify a Refresh value of 100 seconds and a Range of 200 seconds, the status line will be updated every 100 seconds, showing information from the latest 200 seconds.



If you end GO mode, and do not change the Refresh or range values, historical data will be saved, so that when you start GO mode again, the status line will continue where it left off. If you do change the refresh or range value, the status line will start from the beginning again. See [Figure 118 on page 194](#) for an example.

In STOP mode, the row is reduced to a single colored field in the center of the report that shows the overall status of the displayed range. For example:




Figure 119. 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 TSOEPTA and TSOEPTB. There is one common report class TSOREPCL. This definition done in the WLM application creates a heterogeneous report class.

The average response time column shows the average time that a transaction spent waiting in a queue and was active in the system. This gives you an indication of where a possible response time bottleneck may be located.

A performance index is introduced to allow a better comparison between different goals. See [Table 88 on page 197](#) for information on how to calculate the performance index.

The scrollable area is ordered by workload group. Each workload group is followed by a list of its service classes. Every service class is followed by a detailed comparison of actual values versus goals for each service class period. The workload groups, and the service classes below each workload group, are sorted alphabetically.

In detail, the rows show the following:

- For each workload group (indicated by the type W), one line is shown containing the actual values achieved for the whole group. This can be seen as a summary line for that group.
- For each service class (indicated by the type S), one line is shown containing the actual values achieved for the whole class. As for workload groups, it can be seen as a summary line for that class.
- For each service class period (indicated by the period number in the type column), one line is shown containing the defined goals accompanied by the values actually achieved. For a service class with one period, the data of that period is shown.

While the execution velocity goal is a percentage that can easily be compared with an actual value, the WLM response time goals can be specified in two different ways:

- A response time together with a percentile:
In this case, the actual value is a percentage indicating the percentage of transactions that ended within the time specified in the response time goal.
- An average response time:
In this case, the average response time value is shown as actual value that can be compared against the goal.

Cursor-sensitive control on the SYSSUM Report

Cursor-sensitive control on this report lets you navigate to detailed reports that offer a possibility to make single system selections.

Cursor-sensitive control of the sysplex field and of the *Systems* field in the report header leads to the Data Index screen.

Using cursor-sensitive control in the *Type* column:

- On a workload group abbreviation, the report is redisplayed containing only workload group entries.
- On service class or report class abbreviation, the processing is analogous to the workload group abbreviation.

Cursor-sensitive control on a field in the *Importance* column gives you a filtered report. What is displayed on the filtered report depends on the **Type** value that you can specify on the Report Options panel.

- If Type is **ALL**, a workload group and all its service classes are displayed if one service class period has the importance you selected using cursor-sensitive control

- If Type is **W**, a workload group is displayed if it contains a service class period with the importance you selected using cursor-sensitive control
- If Type is **S**, a service class is displayed if it contains a service class period with the importance you selected using cursor-sensitive control

Cursor-sensitive control in all other columns (with the exception of the WAIT Time column):

- 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.

Using cursor-sensitive control on the WAIT Time column will display the pop-up panel with a detailed breakdown of the different wait reasons and their average duration (see [Figure 62 on page 106](#)).

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 120 on page 197](#).




Figure 120. SYSSUM Report for Workload Groups

Field descriptions

Table 88. Fields in the SYSSUM Report	
Field Heading	Meaning
Refresh	The refresh value specified on the Session Options panel. The report is updated according to this value.
Service Definition	Name of the service definition in effect during collection of the workload activity data.
Installed at	The date and time the active service definition was installed.
Active Policy	Name of the service policy in effect during collection of the workload activity data.
Activated at	The date and time the current service policy was activated.
Name	Name of the group being reported on. This field can contain a workload group name, a service class name, or a report class name. For a line with a period, the field remains empty, because there is always a summary line further up containing the name of the class to which the period belongs.

Table 88. Fields in the SYSSUM Report (continued)

Field Heading	Meaning
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
Goals versus Actuals	For a service or report class period, these columns show the goal, if specified, and the actual values corresponding to the goal. That means, if an execution velocity goal was specified, the <i>Response Time</i> goal and actual columns remain empty. If a response time goal with a percentile was specified, the <i>Actual</i> column shows a percentage which corresponds to the response time specified in the goal. The actual average response time can be found in the columns for <i>Avg. Resp. Time</i>. For a report class period, this data will be shown only if the period is homogeneous.
Execution Velocity Goal	The target execution velocity for ended transactions that has been in effect for the period during the reported range. This field shows N/A for system service classes, since they do not have a user defined goal.
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 14 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 <i>Execution Velocity - Actual</i> 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.

Table 88. Fields in the SYSSUM Report (continued)	
Field Heading	Meaning
Time	The time units shown in the columns <i>Goal versus Actuals - Response Time</i> and <i>Avg. Resp. Time</i> are all milliseconds. More than 99999 milliseconds are scaled with K (times one thousand - 10³), M (times one million - 10⁶) or G (times one billion - 10⁹).
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.
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 is blank for heterogeneous report classes periods and report classes. It is also blank for system service classes, since they do not have a user defined goal. The <i>Performance Index</i> field can also be shown in dark blue. This has the same reason as for the <i>Execution Velocity - Actual</i> column, described previously. 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: - Calculate the number of transactions N that correspond to the goal: $N = \frac{\sum \text{Transactions} * \text{Goal Percentage}}{100}$ - Add up all transactions until a bucket M is reached where the sum is greater than N. - The "actual" response time in the previously shown formula for the performance index 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 in this example show:

- Number of buckets
- Response time distribution
- Number of transactions (within that bucket)
- Response time associated with that bucket

Bucket:	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Distr :	<50%	60%	70%	80%	90%	100%	110%	120%	130%	140%	150%	200%	400%	>400%
TRX :	10	10	20	20	20	20	20	10	10	10	5	5		
Time :	1.0	1.2	1.4	1.6	1.8	2.0	2.2	2.4	2.6	2.8	3.0	4.0	8.0	>8.0

The three steps of calculation:

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 89. Fields in the SYSSUM Report - Continuation	
Field Heading	Meaning
Trans Ended Rate	The number of transactions ended per second.

Table 89. Fields in the SYSSUM Report - Continuation (continued)	
Field Heading	Meaning
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 <i>Avg. Resp. Time</i> columns, zeros will show up for "server" service classes in most cases, because their "transactions" are address spaces, and response times are available only for ended transactions. So there are only numbers, when one of the address spaces in that service class ends, or is RESET via operator command.
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 <i>Avg. Resp. Time</i> columns, zeros will show up for "server" service classes in most cases, because their "transactions" are address spaces, and response times are available only for ended transactions. So there are only numbers, when one of the address spaces in that service class ends, or is RESET via operator command.
Avg. Resp. Time ACTUAL Time	In general, this is the sum of the previously described wait and execution times, but does not include ineligible time. For CICS transaction service classes, you may see the average EXECUT time greater than the average ACTUAL time, when you would normally expect EXECUT to be less than or equal to ACTUAL. This is because these two fields report on a different set of transactions. EXECUT time can include transactions which originated on a remote system as well as transactions originating locally. ACTUAL time includes response times for only transactions originating locally. If the remote transaction tends to be longer than the local transaction, EXECUT could be greater than ACTUAL. 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 <i>Avg. Resp. Time</i> columns, zeros will show up for "server" service classes in most cases, because their "transactions" are address spaces, and response times are available only for ended transactions. So there are only numbers, when one of the address spaces in that service class ends, or is RESET via operator command.

Monitor III Utility fields

You can use the Monitor III Utility to customize the SYSSUM report. In addition to the values previously described, you can use the Utility to have the following information shown.

Table 90. 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. If the resource group is associated with a report class, this is the name of a tenant resource group and the associated report class is a tenant report class.
Capacity Min	Minimum CPU capacity limit specified for the resource group in the active WLM policy. Note: N/A is reported if no minimum CPU capacity limit is specified for the resource group in the active WLM policy.
Capacity Max	Maximum CPU capacity limit specified for the resource group in the active WLM policy. Note: N/A is reported if no maximum CPU capacity limit is specified for the resource group in the active WLM policy.
Capacity Actual	Actual capacity, in unweighted CPU service units per second, as consumed within that resource group.
Memory Actual	Amount of memory (in GB) used by all address spaces on this system that are associated with this resource group.
SUMECTR	Enclave transaction rate per second.
SUMECTRE	Average enclave transaction execution time in millisecond.
SUMCRIT	An indication of whether WLM managed the address space has one of the following values during the report interval: I Service Class Period implicitly CPU critical I* First period of Service Class implicitly CPU critical C CPU critical S Storage critical SC Both CPU and Storage critical SI Both implicitly CPU and storage critical (Service Class Period) SI* Both first period implicitly CPU and storage critical (Service Class)
SUMAIINF	WLM batch initiator management is AI-infused or not. The values can be: YES or NO

Report options

The Report Options panel is exactly the same as for the SYSINFO report and STORS report, shown in Figure 109 on page 179. The only difference is, that the group names shown are accumulated from the whole sysplex and not only from a single system.

Note: The list of available service classes will be shown only if one of the sysplex reports SYSSUM, SYSRTD, or SYSWKM has been displayed at least once.

SYSTREND - System Trend Report

The System Trend (SYSTREND) report presents the last 20 reporting ranges for the system summary line (*SYSTEM) or any other selected workload line from the System Information (SYSINFO) report. It can be used:

- To analyze how delay situations develop in the system
- To analyze how long delay situations last in the system
- To understand the change in system utilization
- To identify peak utilization
- As a system summary report

You can also use this report as a starting point to analyze system or workload delay. For a more detailed analysis, select the SYSINFO or respective delay report for any reported range using cursor-sensitive control.

How to request this report

To request this report, select **U** from the Primary menu, then **ST** from the User menu.

You need to enter a system name on the User menu.

Note: If no workload or an invalid workload name is specified on the User Selection menu, the report is created for the *SYSTEM line as shown on the System Information (SYSINFO) report.

Contents of the report

The SYSTREND report has two parts.

- The top part provides information about the start date and time of the first and last reported range, the total reported range and the total and average number of samples used to create the report.
- The bottom part shows for each line the CPU utilization for the system and the respective SRB and TCB percentage, total and active users, and average number of jobs using resources or delayed because of resources for the selected workload.

This information is extracted from the SYSINFO reports for the displayed report interval and the meaning is the same as for the SYSINFO report.

The graphic form of the report shows the average number of active users for each type of delay for the selected workload.

Note: On the SYSTREND report, the RESP Time, TRANS/SEC and VEC Util columns (shown on the SYSINFO report) have been replaced by the TCB% and SRB% columns contained in the ISPF table of the SYSINFO report.




Figure 121. SYSTREND Report

Field descriptions

Table 91. Fields in the SYSTREND Report	
Field Heading	Meaning
Latest:	Begin date and time of the first reported range on the report.
Range/Line:	Reported range per displayed line on the report.
Earliest:	Begin date and time of the last reported range on the report.
Total Range:	Total reported range on the report, expressed in seconds and HH.MM.SS.
Time	The begin time of the reported range.
CPU %	The average CPU utilization percentage for all processors is also displayed on the SYSINFO report. See Table 84 on page 183 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 84 on page 183](#)).

Cursor-sensitive control

Cursor-sensitive control allows you to navigate to the SYSINFO or a delay report for a selected reporting range.

Table 92. 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

Table 92. SYSTREND Report - Cursor-sensitive Control for Navigation (continued)

Report Column where Cursor-sensitive Control is Used	Displayed Report
-AVG USG- PROC, -Avg. Act. Users Delayed For- PROC	PROC
-AVG USG- DEV, -Avg. Act. Users Delayed For- DEV	DEV
-Avg. Act. Users Delayed For- STOR	STOR
-Avg. Act. Users Delayed For- SUBS	DELAY
-Avg. Act. Users Delayed For- OPER	DELAY
-Avg. Act. Users Delayed For- ENQ	ENQ

Note: If the SYSTREND report is recreated using cursor-sensitive control on the Time column, pressing PF3 on the new SYSTREND report will return you to the Primary Menu. In all other cases, pressing PF3 from the report you have selected will return you to the SYSTREND report.

SYSWKM - Work Manager Delays Report

The Work Manager Delays (SYSWKM) report shows details for resource-manager or work-manager oriented subsystems and is intended as a basis on which to start tuning. Using the real-time data, you can use this report to track problems as they happen.

The report shows the average transaction response time and how the various transaction states contribute to it. Furthermore, it lists the address spaces that have been used by the transactions. [Figure 122 on page 206](#) and [Figure 123 on page 206](#) 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-Usq and Veloc columns give an indication of how much work is actually being done, and should be as high as possible.

How to request this report

To request the Work Manager Delays report, select **S** from the Primary Menu, then a **3** on the Sysplex Report menu (shown in [Figure 5 on page 22](#)) or enter the following command using the format:

```
SYSWKM service_class,period
SYSWKM report_class,period
```

Parameter report_class designates either a report class or tenant report class.

As this report is available for homogeneous report classes only, you cannot specify a heterogeneous report class with the reportclass parameter.

For example, to get a Work Manager Delays report for the service class POSMULTI and service class period 1, enter:

```
SYSWKM POSMULTI, 1
```

Contents of the report

RMF 3.1 Work Manager Delays - RMFPLEX1										Line 1 of 3															
Command ==>										Scroll ==> HALF															
WLM Samples: 100 Systems: 2 Date: 04/25/2023 Time: 16.03.00 Range: 100 Sec																									
Class:		POSMULTI		Period: 1		Avg. Resp. time: 587.0 ms				for 12473 TRX.															
Goal:		500.0 ms		average		Avg. Exec. time: 379.0 ms				for 12389 TRX.															
Actual:		587.0 ms		average		Abnormally ended:				0 TRX.															
-----Response time breakdown (in %)-													-Switched--												
Sub P		-----		-----		-----				-----		-----													
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 122. SYSWKM Report for Subsystem CICS

The panel shows an example for a CICS system using IMS as database, where all measurement values that theoretically could be provided are actually available. However, in a CICS system with regions spread over several z/OS images and with different CICS releases installed, data may be available or missing in nearly every column, depending on how varied the installation is and how the CICS releases involved differ.

RMF 3.1 Work Manager Delays - RMFPLEX1										Line 1 of 6			
Command ==>										Scroll ==> HALF			
WLM Samples: 100 Systems: 2 Date: 04/25/2023 Time: 16.03.00 Range: 100 Sec													
Class: CRDTAUTH		Period: 1		Avg. Resp. time: 587.0 ms for					12473 TRX.				
Goal: 500.0 ms for		80%		Avg. Exec. time: 311.0 ms for					12389 TRX.				
Actual: 500.0 ms for		91%		Abnormally ended:					0 TRX.				
-----Response time breakdown (in %)-													
Sub P		-----Switched--											
Type		Tot		Act		Rdy		Idle		Delayed by		Time (%)	
												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 123. 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 CICSLOW.

- 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:

Goal: 0.500 sec for 80%	Avg. Resp. time: 0.586 sec for 12473 TRX.
	Avg. Exec. time: 0.311 sec for 12389 TRX.

Here the goal is met, even though the average response time is 0.586 seconds.

Response time breakdown

In the middle section, a response time breakdown for the various transaction states is shown, split into total time and execution time.

The unit of the response time values shown can be switched between percentages and seconds using the Report Options panel or use cursor-sensitive control anywhere in the middle section of report.

If you select seconds and the value does not fit, then ******* will be shown in that report field. In this case, changing to percentage will provide a better representation of the figures.

Address spaces serving

In the bottom section, the **Address Spaces Serving this Service Class** during the report interval are listed in a scrollable area.

- For each address space, the jobname (together with an indication about how WLM is managing a server region), address space id, and system id are shown to allow you to track the address space to the specific z/OS image.
- The service class shown is the class the serving address space belongs to. This helps you to relate this data to the Sysplex Summary report.

- The service percentage (**Service**) shows the percentage of service given to the reported service class. For example, if the address space serves only this class, then 100% is shown. If the address space gives equal service to three different service classes, then 33% is shown.
- Then, for each address space the execution velocity, the processor and I/O using percentage are shown. This gives you a hint of the "health" of the address space.
- Finally, for each address space, a capping percentage and a quiesce percentage are shown. The capping column shows the WLM percentage for capping. If the address space was delayed for other reasons as well, the actual capping delay may be much smaller. The actual capping delay is added to the single system Processor Delays (PROC) report.

The quiesce percentage normally shows either 0 or 100, because the address space is either quiesced by the operator with the RESET command or not. However, for a time range where the quiesce state was changed, a percentage between 0 and 100 is possible, indicating how long the address space was quiesced during the report interval.

This part of the report is empty if you call the report for a report class.

Cursor-sensitive control on the SYSWKM Report

Cursor-sensitive control used on the response time fields in the sub-header lines in the top portion of the report shows you the Response Time Distribution report for that service class period.

Cursor-sensitive control used on the response time breakdown fields in the middle portion of the report switches between the units that can be selected for the data. This "toggling" does not change the unit selected on the Report Options panel.

- If the current unit is *seconds*, cursor-sensitive control switches the unit to percentage.
- If the current unit is *percentage*, cursor-sensitive control switches the unit to seconds.

Cursor-sensitive control on the server address space section in the scrollable bottom portion of the report is active on the following fields:

- Cursor-sensitive control on column *Jobname* and *ASID* leads to the JOB Delay report of the respective system.
- Cursor-sensitive control on column *System* leads to the SYSINFO report of the respective system.
- Cursor-sensitive control on column *Service Class* and *Service* leads to the GROUP report of the respective system.
- Cursor-sensitive control on column *Proc-Usg* leads to the PROC report of the respective system.
- Cursor-sensitive control on column *I/O-Usg* leads to the DEV report of the respective system.
- Cursor-sensitive control on column *Velocity* leads to the Delay report of the respective system.

Field descriptions

All fields of the report are described in detail in the following field description table:

Table 93. 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 the average response time in milliseconds, 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.

Table 93. Fields in the SYSWKM Report (continued)	
Field Heading	Meaning
Average execution time	The average execution time of all ended transactions belonging to that period, possibly spread over several systems.
For <i>nnnnn</i> TRX.	<i>nnnnn</i> is the total number of transactions that completed their execution phase during this report interval.
Abnormally ended	The number of abnormally ended transactions from all reported systems. This value is not included in the number of total completed transactions.
Subsystem Type	A 4-character identification for the subsystem for which the data was attributed to, as shown for example in the WLM administrative application.
Phase (P)	B 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 93. 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 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 z/OS 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. BPMI Time spent waiting for I/O resulting from a Db2 buffer pool miss.
Switched Time (%)	Percentage of time that transactions spent routed to another region for processing. This percentage also refers to the <i>Average Response Time</i> shown in the sub-header. For a begin-to-end phase, the sum of these percentages should approximately equal the value shown in the <i>CONV</i> column. For an execution phase, these percentages, as well as the figure in the <i>CONV</i> column, are expected to be zero. LOC Percentage of time that transactions spent switched on this z/OS image. Subsystems might set this state when they function ship a transaction to another component within the same z/OS image. SYS Percentage of time that transactions spent switched to another z/OS image in the sysplex. Subsystems might set this state when they function ship a transaction to another component on another z/OS image within the sysplex. 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 93. Fields in the SYSWKM Report (continued)	
Field Heading	Meaning
Address Spaces Serving this Service Class <i>srvcls</i>	<i>Srvcls</i> is the name of the <i>Class</i> 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	The capping delay percentage of the address space. Using WLM, the minimum and maximum capacity values for a resource group can be used to restrict the amount of processor capacity that a collection of address spaces is allowed to consume, or with discretionary management, if the work for which the job is running is overachieving its goal, this work may be capped in order to divert its resources to run discretionary work (see also section 'Using Discretionary Goals' in 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    CICSPTS   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 124. 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. Tenant report classes are included in the list of report classes. These names are obtained from the current data.

Note: The classes will be shown only if one of the sysplex reports SYSSUM, SYSRTD, or SYSWKM has been displayed at least once.

The RMF *FIND* command works on the scrollable area.

USAGE - Monitor III Job USAGE Report

The Monitor III Job USAGE Report is provided as a complement to the Monitor III Job Delay Report. The USAGE report allows you to identify at a glance the jobs that are consuming the most resources within various resource categories. The report gives you information about job resource consumption in terms of I/O, processor, and storage related key metrics.

Additionally, the report provides a dedicated section with QSCAN usage statistics. This enables you to identify the jobs that have issued either the GQSCAN or the ISGQUERY REQINFO=QSCAN service.

Hence, those jobs that are issuing QSCAN requests improperly or too frequently can be detected more easily.

How to request this report

To request the USAGE (Job Usage) report, select **1** from the Primary Menu, and then select **4A** on the Overview Report menu (shown in [Figure 6 on page 23](#)) or enter the following command using the format:

```
USAGE [ job_class, service_class ]
```

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

```
USAGE T, TSOPRIME
```

Contents of the report

[Figure 125 on page 213](#) shows a sample Job Usage report.

```

RMF 3.1    Job Oriented Usage                                Line 1 of 14
Command ===>                                                Scroll ===> CSR
Samples: 60      System: SYSF   Date: 04/25/2023   Time: 15.35.00   Range: 60      Sec
Jobname  Service --- I/O --- --- CPU --- - Storage - ----- QScan ----
          CX Class Conn  EXCP Total  TCB Total Fixed Total Resct Time
BHBE     T  TSODEF  0.399 16.42  0.22  0.22   743    0    0    0    0
XCFAS    S  SYSTEM  0.264  4.30  0.02  0.02 13443 2943  0    0    0
*MASTER* S  SYSTEM  0.041  0.27  0.05  0.01  5472 3643  0    0    0
BJAGHTM  B0 BTCHDEF  0.037 18.50  0.06  0.06   499   58   8    0  810
CATALOG  S  SYSTEM  0.037  1.58  0.02  0.02  1418  176  0    0    0
RMFGAT   S0 SYSSTC  0.033  0.12  0.23  0.23 14505  152  0    0    0
JES2     S  SYSSTC  0.019  1.85  0.06  0.06  9390  766  0    0    0
GRS      S  SYSTEM  0.018  0.00  0.03  0.03 14680  454  3    0  551
CONSOLE  S  SYSTEM  0.000  0.03  0.01  0.01  3921  157  0    0    0
GPMSEVER S0 GPMSEVER 0.000  0.02  0.02  0.02  3018  101  0    0    0
WLM      S  SYSTEM  0.000  0.00  0.10  0.10 29486  272  0    0    0
RRS      S  STCDEF  0.000  0.00  0.01  0.00  2813  178  0    0    0
TCPIP    S0 SYSSTC  0.000  0.00  0.16  0.00  8038  150  0    0    0
RMF      S  SYSSTC  0.000  0.00  0.01  0.01  7281  122  2    0  443

```

Figure 125. Job Usage report

Field descriptions

Table 94. Fields in the Job Usage report	
Field Heading	Meaning
Jobname	Name of the job.
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 for this address space.
I/O Conn	Device connect time, in seconds, for this address space in the report interval.
I/O EXCP	Number of EXCP operations per second for this address space in the report interval.
CPU Total	Amount of total processor time, in seconds, for this address space in the report interval as sum of TCB time, global and local SRB time, and preemptable or client SRB time.
CPU TCB	Amount of TCB processor time, in seconds, for this address space in the report interval.

Table 94. Fields in the Job Usage report (continued)

Field Heading	Meaning
Storage Total	The number of active and idle frames for this address space, averaged over the report interval.
Storage Fixed	Number of fixed frames for this address space, averaged over the report interval.
QScan Total	Total number of QScan requests for this address space, including START and RESUME, but not QUIT requests.
QScan Resct	Average number of resources returned by QScan requests for this address space.
QScan Time	Average QScan request time, in microseconds, for this address space.

Monitor III Utility fields

You can use the Monitor III Utility to customize the Job Usage report. In addition to the information previously shown, you can use the Utility to have the following values shown:

Table 95. Additional fields in the Job Usage report

Field Name	Meaning
JUSPASI	Address space ID (decimal)
JUSPCLA	Abbreviation for the job class as follows: A ASCH B Batch O OMVS S Started task T TSO
JUSPCLP	The number of the service class period for this address space.
JUSPDPP	Dispatching priority for this address space.
JUSPTAT	The time that has elapsed since the current transaction in this address space was started, in the form <i>hhh:mm</i>
JUSPTRT	The time that has elapsed since the current transaction in this address space became resident, in the form <i>hhh:mm</i> or <i>hh:mm:ss</i>
JUSPTCT	Number of transactions for this address space since address space creation.
JUSPFRXH	Number of fixed frames above 2GB for this address space averaged over the report interval.
JUSPFRXA	Number of fixed frames between 16MB and 2GB for this address space averaged over the report interval.
JUSPFRXB	Number of fixed frames below 16MB for this address space averaged over the report interval.

<i>Table 95. Additional fields in the Job Usage report (continued)</i>	
Field Name	Meaning
JUSPDCTT	Accumulated device connect time, in seconds, for this address space since address space creation.
JUSPEXCT	Number of EXCP operations total since address space creation.
JUSPEXCD	Number of EXCP operations in this interval.
JUSPCPUT	Accumulated amount of processor time, in seconds, for this address space since address space creation as sum of TCB time, global and local SRB time, and preemptable or client SRB time.
JUSPTCBT	Accumulated amount of TCB processor time, in seconds, for this address space since address space creation.
JUSPQSPR	Number of specific QSCAN requests for this address space, that are either GQSCAN requests specified by QNAME and RNAME or ISGQUERY requests specifying a search by ENQTOKEN.
JUSPQRSD	Standard deviation for the number of resources returned by QSCAN requests for this address space.
JUSPQTSD	Standard deviation for the QSCAN request time for this address space.
JUSPRPB	YES or NO . Indicates whether the address space had Recovery Process Boost active during the interval.

Cursor-sensitive control on the USAGE Report

Table 96 on page 215 shows which report is displayed for each cursor-sensitive field.

<i>Table 96. Cursor sensitivity on Job Usage report</i>	
Report column	Displayed report
Jobname	DELAYJ Report for selected job
CX	USAGE Report filtered by job class
Service Class	USAGE Report filtered by service class
I/O Conn	DEV Report
I/O EXCP	DEV Report
CPU Total	PROCU Report
CPU TCB	PROCU Report
Storage Total	STORF Report
Storage Fixed	STORF Report

WFEX - Workflow/Exceptions Report

The Workflow/Exceptions (WFEX) report presents information about system activity and system resources.

The top part of the report shows you speedometers in graphic mode, or workflow indicators in tabular mode.

The color of a workflow indicator tells you how well the jobs are performing depending on the exception criteria specified. (Usually, red indicates a problem, yellow indicates caution, and turquoise indicates that a job or volume is missing from the system configuration.)

The speedometer needle points to the relative speed of the job or resource in the system. The shaded part to the left of the needle represents the proportion of a user's time spent doing useful work. The part to the right of the needle represents the proportion of a user's time spent delayed.

You can define the exception criteria on the Definition and Criteria panel.

A line in the Exceptions section of the report corresponds to each colored speedometer or workload indicator. The line has the same color and the same Name as the speedometer or workload indicator, and gives details about the exception.

Speed (Workflow)

Under **Speed (Workflow)**, a high workflow percentage or speed indicates that a job has the resources it needs to process, and that it is moving through the system at a relatively high speed.

If the Criteria Set you defined for a workflow indicator is met, it will be displayed in the corresponding color.

A low value under Speed indicates that a job has few of the resources it needs and is contending with other jobs for system resources and may indicate a problem.

For resources (for example *PROC and *DEV), a high value under Speed indicates that jobs are moving through the system with little resource contention. A low Speed represents a large queue of work requests from users and may indicate a problem.

Exceptions

In the **Exceptions** section of the report, a line that is displayed in yellow or red indicates a job or resource exceeds the exception criteria defined by you on the **WFEX Report Options: Definition and Criteria** panel. Use cursor-sensitive control to investigate exceptions further.

If **Not avail** appears on your report, the job you selected on the Definition panel was not running during the report interval. If **No work** appears, the job or group was idle (not requesting system resources) during the report interval.

How to request this report

To request the Workflow/Exceptions, select **1** from the Primary Menu, then select **1** from the Overview Report menu (shown in [Figure 6 on page 23](#)) or enter the following command:

```
WFEX
```

Contents of the report

```

RMF 3.1 Workflow/Exceptions                                     Line 1 of 12
Command ==>                                                    Scroll ==> HALF

Samples:   100 System: MVS1 Date: 04/25/2023 Time: 08.53.20 Range: 100 Sec

----- Speed (Workflow) -----
Name      Speed of 100 = Maximum, 0 = Stopped      Average CPU Util: 81 %
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
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 126. 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 126 on page 217 is an example of the Workflow/Exceptions report. For information about the WFEX Report Options panels, see [“Report options” on page 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 230](#).

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 10](#).

A low workflow percentage indicates that a job has few of the resources it needs and is contending with other jobs for system resources. A high workflow percentage indicates that a job has the resources it needs to execute, and that it is moving through the system at a relatively high speed.

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

Workflow of resources (processors or devices) represents how well the system is serving users. The speed at which each resource performs the work of all user's is expressed as a value from 0% to 100%. A low resource workflow percentage represents a large queue of work requests from users. A high workflow percentage represents little resource contention.

If **Not avail** appears on your report, the job that you selected on the Definition and Criteria panel was not running during the report interval. If **No work** appears, the job or job group was idle (not requesting system resources) during the report interval.

Exceptions are suppressed without notification when using criteria on historic RMF gatherer records that do not provide the corresponding data. However, if there is at least one valid criteria set defined, and the thresholds are met, the exception is displayed.

Field descriptions — Speed Section

Table 97. Fields in the Speed Section of the WFEX Report	
Field Heading	Meaning
Average CPU Util	The average utilization percentage for all general purpose processors (CPs) during the report interval. For details, see Table 84 on page 183 .
Name	The one to ten character identifier of a workflow indicator. It can be a job, job group, or resource (processor or device). You can specify Name on the Label field of the Definition and Criteria panel or leave it blank and use the default name generated by RMF. If an indicator changes color, there is a corresponding line in the Exceptions section of the report with the same name and color giving more information about the exception.
Users/Active	The average number of users and the average number of active users in an address space or group of address spaces. The average number of active users is a measure of system workload. A user in a system is either ACTIVE, IDLE, or unknown during a report interval. An ACTIVE user is using a resource or is delayed by a resource. An IDLE user is in terminal wait, timer wait, or is waiting for JES job selection. A user that is not in either of these states is unknown.
Speed	For jobs and job groups, Speed is a measure of the speed at which jobs are moving through the system in relation to the maximum speed at which the jobs could move through the system. A low workflow percentage indicates that a job has few of the resources it needs and is contending with other jobs for system resources. A high workflow percentage indicates that a job has the resources it needs to execute, and that it is moving through the system at a relatively high speed. For resources (processors or devices), Speed represents how well the system is serving the users. A low resource workflow percentage represents a large queue of work requests from users. A high workflow percentage represents little resource contention.

Field descriptions — Exceptions Section

The Exceptions section of the report shows the exceptional situations a job or job group encountered in relation to the hardware and software resources. The exceptions are those specified on the Workflow/Exceptions Report Options panels.

Table 98. 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.

Table 98. Fields in the Exceptions Section of the WFEX Report (continued)

Field Heading	Meaning
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 98. 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 then the name of the job using the processor most often when the delayed job wanted to use it, appears to the right of the hyphen. If one or more enclaves contributed most to the processor delay, then *ENCLAVE will be shown under Reason. DEV then the volume serial number of the device that the job was most often delayed for, appears to the right of the hyphen. ENQ then the major name of the serially reusable resource causing the greatest percentage of delay, appears to the right of the hyphen. STOR then 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 then either Message or Mount or Quiesce appears to the right of the hyphen. Message indicates that the operator did not respond to a message. Mount indicates that the operator did not mount a tape. Quiesce indicates that the address space was quiesced by the operator. SUBS then 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 98. 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. 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 13.
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.

Report options

```

RMF WFEX Report Options: Action Panel
Command ==>
Line 1 of 23
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
--      -
--      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 127. 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 230](#).

On the Report Options panels, you can specify workflow indicators, exception conditions, or a combination of both. However, on the **Speed** section of the report, a limit of 14 workflow indicators can be displayed. You can also specify the threshold values and color highlighting criteria for the indicators.

RMF displays the Action panel. The Action panel shows a list of the Report Options currently in effect. On the Action panel, you can enter codes in the **Action** column to change, add, and delete workflow and exception indicators, and vary where the indicators will appear in the report.

To add, view, or change criteria, enter the Add (**AD**) or Select (**S**) code in the **Action** column. RMF then displays the Definition and Criteria panel.

To reset the Report Options to the RMF default values, if customization is set to NO in your Session Options, enter the RESET command.

Table 99. 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, 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
----- --  ----  ----          ----- --  ----  ----          ----- --  ----  ----
----- --  ----  ----          ----- --  ----  ----          ----- --  ----  ----
----- --  ----  ----          ----- --  ----  ----          ----- --  ----  ----
----- --  ----  ----          ----- --  ----  ----          ----- --  ----  ----
----- --  ----  ----          ----- --  ----  ----          ----- --  ----  ----
----- --  ----  ----          ----- --  ----  ----          ----- --  ----  ----
----- --  ----  ----          ----- --  ----  ----          ----- --  ----  ----

```

Figure 128. 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.

Table 100. 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 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. 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. 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 100. 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 129 on page 226 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.
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.

Table 100. Fields in the WFEX Definition and Criteria Panel (continued)	
Field Heading	Meaning
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 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
Y	PAGE08	T	SP00L1	T	DRV021	T	D13IOG
		T	410PRM	T	D13MCP	T	D94RM1
		T	410SR0				

Figure 129. WFEX Definition of UNAVAIL Panel

On this panel, you modify the report by defining or changing exception conditions. The above sample shows a panel for the class DEV.

Specify Label, Alert, and Text information on the top half of the panel.

Note: Class, Qualifier, and Indicator cannot be changed on this panel.

Select or type in the volume and color highlighting in the bottom half of the panel. Whenever the selected volume is not mounted, it will appear as an exception in the report highlighted in your chosen color.

To exit this panel, you must either:

- Select or type in at least one volume and press the END key.
- Use the CANCEL command to cancel any volume selections.

On this panel, you can enter one of the following codes in the **A** column to color highlight exception lines.

- T** Turquoise highlighting
- Y** Yellow highlighting
- R** Red highlighting

You can also type in any volume on a blank line at the top of the list of volumes.

If you want to remove a volume from this list, just specify a **blank** in column **A**.

Criteria Names Selection panel

On any of the six Criteria Names Selection panels, you can modify the report by selecting or changing exception conditions.

Each selection panel shows a complete list of exceptions allowed for the specified Class, Qualifier, and Indicator. You can select up to 7 criteria names. The selected criteria names are shown in the corresponding criteria set when you return to the Definition and Criteria panel.

The six Criteria Names Selection panels are:

- Criteria selection panel for jobs
- Criteria selection panel for service classes, ASCH, and OMVS
- Criteria selection panel for processors
- Criteria selection panel for report performance groups
- Criteria selection panel for storage
- Criteria selection panel for devices

To get to the Criteria Names Selection panel, use cursor-sensitive control on a blank Name field on the Definition and Criteria panel.

All of the conditions (Names) defined within a Criteria Set must be met in order for the color highlighting to appear on the Workflow/Exceptions report. The following figures show sample Criteria Names Selection panels, all panels are scrollable.

```

RMF WFEX Criteria Names for Class: JOB
Command ==>
Select (S) a maximum of 7 items to use in a criteria set. Press END.

More: +
- AAU   Average active users in group
- AAUS  Avg active users in system
- ADU   Average delayed users
- ASTO% Active storage %
- AUU   Average using users
- AVAIL Number of jobs in group
- COMM% Common storage delay %
- CPU%  TCB+SRB %
- CPUS% CPU utilization %
- CSA%  CSA storage %
- DEV%  Device delay %
- DLY%  Overall delay %
- ENQ%  Enqueue delay %
- ESMAG Exp storage migration age
- ESMRT Exp storage migration rate
- ESPRT Exp storage page movement rate
- ESQO% ESQA overflow %
- EWSET Expanded storage WSET frames
- HIPR% Hiperspace storage delay %
- HSM%  HSM delay %
- ISTO% Idle storage %
- JCSA% Job CSA use %
- JECS% Job ECSA use %
- JES%  JES delay %
- JESQ% Job ESQA use %
- JSQA% Job SQA use %
- LOCL% Local storage delay %
- LPA%  LPA storage %
- MAXU% Maximum allowed user %
- MNT%  Tape mount delay %
- MSG%  Operator message delay %
- ONLF  Online real storage frames
- ONLXF Online expanded storage frames
- OPER% Operator delay %
- OREPL Outstanding replies
- OTR%  Out/ready delay %
- PINRT Page-in rate
- PROC% Processor delay %
- QUI%  Quiesced delay %
- RATE  Transaction rate
- SCSA% System CSA use %
- SECS% System ECSA use %
- SESQ% System ESQA use %
- SQA%  SQA storage %
- SQA0% SQA overflow %
- SSQA% System SQA use %
- STOR% Storage delay %
- SUBS% SUBS delay %
- SWAP% Swap delay %
- TET   Transaction elapsed time
- TSQA0 Total SQA overflow
- UIC   Unreferenced interval count
- USG%  Overall using %
- USGD% Device using %
- USGP% Processor using %
- VIO%  VIO storage delay %
- WFL%  Workflow/Speed %
- WFLG% Group workflow %
- XCF%  XCF delay %
- XMEM% Cross-memory storage delay %

```

Figure 130. 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
- AAUS Avg active users in system
- ADU  Average delayed users
- ASTO% Active storage %
- AT   Execute response time
- AUU  Average using users
- AVAIL Number of jobs in group
- COMM% Common storage delay %
- CPU%  TCB+SRB %
- CPUS% CPU utilization %
- CSA%  CSA storage %
- DEV%  Device delay %
- DLY%  Overall delay %
- ENQ%  Enqueue delay %
- ESMAG Exp storage migration age
- ESMRT Exp storage migration rate
- ESPRT Exp storage page movement rate
- ESQO% ESQA overflow %
- EWSET Expanded storage WSET frames
- HIPR% Hiperspace storage delay %
- HSM%  HSM delay %
- ISTO% Idle storage %
- JCSA% Job CSA use %
- JECS% Job ECSA use %
- JES%  JES delay %
- JESQ% Job ESQA use %
- JSQA% Job SQA use %
- LOCL% Local storage delay %
- LPA%  LPA storage %
- MAXU% Maximum allowed user %
- MNT%  Tape mount delay %
- MSG%  Operator message delay %
- ONLF  Online real storage frames
- ONLXF Online expanded storage frames
- OPER% Operator delay %
- OREPL Outstanding replies
- OUTR% Out/ready delay %
- PINRT Page-in rate
- PROC% Processor delay %
- QT   Queue time (ASCH)
- QUI% Quiesced delay %
- RATE Transaction rate
- RT   Response time total
- SCSA% System CSA use %
- SECS% System ECSA use %
- SESQ% System ESQA use %
- SQA%  SQA storage %
- SQA0% SQA overflow %
- SSQA% System SQA use %
- STOR% Storage delay %
- SUBS% SUBS delay %
- SWAP% Swap delay %
- TET   Transaction elapsed time
- TSQA0 Total SQA overflow
- UIC   Unreferenced interval count
- USG%  Overall using %
- USGD% Device using %
- USGP% Processor using %
- VIO%  VIO storage delay %
- WFL%  Workflow/Speed %
- WFLG% Group workflow %
- XCF%  XCF delay %
- XMEM% Cross-memory storage delay %

```

Figure 131. 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
- AAUS Avg active users in system
- ADU  Average delayed users
- ASTO% Active storage %
- AUU  Average using users
- AVAIL Number of CPUs
- CPUS% CPU utilization %
- CSA%  CSA storage %
- ESMAG Exp storage migration age
- ESMRT Exp storage migration rate
- ESPRT Exp storage page movement rate
- ESQO% ESQA overflow %
- ISTO% Idle storage %
- LPA%  LPA storage %
- ONLF  Online real storage frames
- ONLXF Online expanded storage frames
- SCSA% System CSA use %
- SECS% System ECSA use %
- SESQ% System ESQA use %
- SQA%  SQA storage %
- SQA0% SQA overflow %
- SSQA% System SQA use %
- TSQA0 Total SQA overflow
- UIC   Unreferenced interval count
- WFL%  Workflow/Speed %

```

Figure 132. WFEX Criteria Names for Class: Processor

```

RMF WFEX Criteria Names for Class: STOR
Command ==>

Select (S) a maximum of 7 items to use in a criteria set. Press END.

- ASTO% Active storage %           - ONLXF Online expanded storage frames
- CSA%  CSA storage %             - SCSA% System CSA use %
- ESMAG Exp storage migration age  - SECS% System ECSA use %
- ESMRT Exp storage migration rate - SESQ% System ESQA use %
- ESPRT Exp storage page movement rate - SQA% SQA storage %
- ESQO% ESQA overflow %           - SQA0% SQA overflow %
- ISTO% Idle storage %            - SSQA% System SQA use %
- LPA%  LPA storage %             - TSQA0 Total SQA overflow
- ONLF  Online real storage frames - UIC  Unreferenced interval count

```

Figure 133. WFEX Criteria Names for Class: Storage

```

RMF WFEX Criteria Names for Class: DEV
Command ==>

Select (S) a maximum of 7 items to use in a criteria set. Press END.

- AAU  Average active users in group - ESQO% ESQA overflow %
- AAUS Avg active users in system    - ISTO% Idle storage %
- ACT% Device active time %          - LPA%  LPA storage %
- ADU  Average delayed users         - ONLF  Online real storage frames
- ASTO% Active storage %             - ONLXF Online expanded storage frames
- AUU  Average using users           - PND%  Device pending time %
- AVAIL Number of devices            - SCSA% System CSA use %
- CON% Device connect time %         - SECS% System ECSA use %
- CPUS% CPU utilization %            - SESQ% System ESQA use %
- CSA%  CSA storage %               - SQA%  SQA storage %
- DAR  Device activity rate          - SQA0% SQA overflow %
- DRT  Device response time          - SSQA% System SQA use %
- DSC% Device disconnect time %      - TSQA0 Total SQA overflow
- ESMAG Exp storage migration age     - UIC  Unreferenced interval count
- ESMRT Exp storage migration rate    - WFL%  Workflow/Speed %
- ESPRT Exp storage page movement rate

```

Figure 134. WFEX Criteria Names for Class: Device

To select a criterion name, type an **S** next to the **Name(s)** you want to use in a criteria set and press the END key.

Each **Name** represents a condition that RMF checks for as an exceptional value.

You can select a maximum of seven names for a criteria set.

To delete a criterion name, you can either:

- Blank out the **S** next to the criterion name you want to delete on the Criteria Names Selection panel.
- Blank out the threshold values in the **Yel** and **Red** fields on the Definition and Criteria panel. When you press END, RMF removes the whole criterion line.

Automatic customization

To automatically customize RMF option sets, specify YES for Customization on the Session Options panel.

Automatic customization ensures that the option set used matches the service policy name and system ID that was in effect when the data was gathered. If the service policy name or the system ID changes between reports, an option set with the same name is made active. The new option set is listed under Current Option Set on the Option Set Selection Menu (invoked via command OPTSET).

If no option set exists with the same name as the service policy name and system ID associated with the data, a new option set with that name is generated from the current option set and made active.

Automatic customization is turned off if you enter NO in the Customization field on the Session Options panel or if you change the Current Option Set on the Option Set Selection menu to an option set that does not match the service policy name and system ID of the data at which you are looking.

Every time a new option set is selected or created, RMF saves all the changes from the old option set before making the new option set current.

Workflow/Exceptions graphic Report




Figure 135. WFEX Graphic Report

The Workflow/Exception graphic report illustrates workflow in speedometers.

The speedometer needle points to the relative speed of the job or resource in the system. The solid/colored part to the left of the needle represents the proportion of a user's time spent doing useful work. The part to the right of the needle represents the proportion of a user's time spent delayed. If the part to the right of the needle is colored either yellow or red, then one or more exception criteria were met. A line in the Exceptions section of the report corresponds to each yellow or red speedometer. The line has the same name and the same color as the speedometer, and gives details about the exception.

Field descriptions – Graphic WFEX Report

Table 101. 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 101. 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.

XCF - Cross-System Coupling Facility Delays Report

The XCF Delays report lets you investigate situations where executing jobs are delayed when requesting service from XCF.

How to request this report

To request the XCF Delays report, select **4** from the Primary Menu, then select **3** from the Subsystem Report menu (shown in [Figure 9 on page 25](#)) or enter the following command:

```
XCF [job_class,service_class]
```

Contents of the report




Figure 136. XCF Report

The graphic form of this report shows the percentage of each user's time spent waiting for XCF services. RMF reports the overall delay (DLY %) and the four paths contributing most to delay (Main Delay Paths) due to XCF signalling traffic. RMF lists all delayed jobs by descending delay percentages.

None appears as the path number for pending jobs without an associated device number.

Possible causes for high XCF delay value might be caused by one or more of the following:

- Path capacity exceeded.
- Other applications are stressing the path.
- XCF delays on the receiving system.
- Some data paths are unavailable or offline.

Note: Any delay value shown in the report represents a delay of a message being sent. All messages are sent asynchronously. Whether the application can truly 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 102. 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.
DLY %	Delay the waiting job (address space) is experiencing because of delay for XCF during the report interval, expressed as a percentage. $\text{DLY \%} = \frac{\text{\# Delay Samples}}{\text{\# Samples}} * 100$ 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 43 on page 75](#) for an example. If you select YES for Jobs on the Report Options panel, the Job Selection/Exclusion panel is displayed. See [Figure 41 on page 73](#) for an example.

XCFGROUP - XCF Group Data

The XCF Group Data report provides information about the groups and members, their statuses and request flow.

How to request the XCF Group Data report

You can request the XCF Group Data report by using one of the following options:

- Select **S** from the RMF Monitor III Primary Menu, and then select **22** on the RMF Sysplex Report Selection Menu.
- Enter the following command:

```
XCFGROUP
```

Contents of the XCF Group Data report

```

RMF V3R1   XCF Group Data   - SYSPLEX1           Line 1 of 10
Command ===>                               Scroll ===> CSR

Samples: 100   Systems: 3   Date: 01/01/23   Time: 00.00.00   Range: 100   Sec

Group Name   Member Name   Status   System Name   Job Name   ---Requests---
                                     Outbound   Inbound

XSDSTI1     *ALL                               960       960
XSDSTI1     XSDSTI1S1A   Active   SYS1          DS1AIRLM   264       264
XSDSTI1     XSDSTI1S1B   Active   SYS1          DS1BIRLM   456       456
XSDSTI1     XSDSTI1S2C   Active   SYS2          DS2CIRLM   240       240
XSDSTI2     *ALL                               832       832
XSDSTI2     XSDSTI2S2A   Active   SYS2          DS2AIRLM   208       208
XSDSTI2     XSDSTI2S2B   Active   SYS2          DS2BIRLM   408       408
XSDSTI2     XSDSTI2S3C   Active   SYS3          DS3CIRLM   216       216
XSDSTI3     *ALL                               0         0
XSDSTI3     XSDSTI3S3A   Active   SYS4          DS3AIRLM   0         0

```

Figure 137. XCF Group Data Report

Note: There is no graphic version of this report available.

Field descriptions of the XCF Group Data report

Table 103. Fields in the XCF Group Data Report	
Field Heading	Meaning
Group Name	The XCF group name to which the member belongs to.
Member Name	The XCF member that sent or received requests. The *ALL line displays the total number of inbound and outbound requests for this XCF group.
Status	The status of the XCF member: - Active - The member is active and can use the XCF services - Missing - Information about the status of this member is missing - Created - The member is known to XCF, but it cannot use the XCF services - Quiesced - The member is disassociated from XCF services - Failed - The member's associated task, job step task, address space, or system has terminated - monitorRemoved - Monitoring has been removed for this member - sysTermination - XCF system containing member is terminating
System Name	The system name where this member resides.
Job Name	The job name where the member belongs to.
Outbound Requests	The number of messages sent by the member on the local system.
Inbound Requests	The number of messages received by the member on the local system.

XCFOVW - XCF Overview

The XCF Overview report provides a quick overview with basic information about systems in the sysplex.

How to request the XCF Overview report

You can request the XCF Overview report by using one of the following options:

- Select **S** from the RMF Monitor III Primary Menu, and then select **19** on the RMF Sysplex Report Selection Menu.
- Enter the following command:

XCF0VW

Contents of the XCF Overview report

```

Command ==>          RMF V3R1   XCF Overview   - SYSplex1          Line 1 of 4
                                Scroll ==> CSR

Samples: 100      Systems: 4      Date: 01/01/23   Time: 00.00.00   Range: 100   Sec

System   SMF   Partition z/OS      -----XCF Interval----- XCF      RMF
Name     ID    Name      Level      Monitoring Operator   Status   Master

SYS1     SYS1  SYS1      z/OS 02.05.00   165.00   168.00   Active   No
SYS2     SYS2  SYS2      z/OS 02.05.00   165.00   168.00   Active   No
SYS3     SYS3  SYS3      z/OS 02.04.00   165.00   168.00   Active   No
SYS4     SYS4  SYS4      z/OS 02.05.00   165.00   168.00   Active   Yes

```

Figure 138. XCF Overview Report

Note: There is no graphic version of this report available.

Field descriptions of the XCF Overview report

Table 104. Fields in the XCF Overview Report	
Field Heading	Meaning
System Name	Name of the z/OS image in the sysplex.
SMF ID	SMF identifier.
Partition Name	Name of the logical partition.
z/OS Level	The version level of the operating system.
XCF Monitoring Interval	Specifies the XCF failure detection interval in seconds. This value is based on the INTERVAL parameter specified in SYS1.PARMLIB(COUPLExx).
XCF Operator Interval	Specifies the amount of time at which XCF on another system is to notify the operator about a failure. This value is based on the OPNOTIFY parameter specified in SYS1.PARMLIB(COUPLExx).
XCF Status	Specifies the XCF status of the system: - Active - the system is active in the sysplex - Missing - the status update from this system is missing - Removed - the system is partitioning out of the sysplex - Local - single system, no sysplex - Cleanup - the system is still in the process of cleanup - Unknown - the status of the system is not known
RMF Master	This system is the RMF master gatherer in the sysplex.

XCFPATH - XCF Path statistics

The XCF Path statistics report provides information about XCF paths, their status and signal flow.

How to request the XCF Path statistics report

You can request the XCF Path statistics report by using one of the following options:

- Select **S** from the RMF Monitor III Primary Menu, and then select **21** on the RMF Sysplex Report Selection Menu.
- Enter the following command:

```
XCFPATH
```

Contents of the XCF Path statistics report

```

Command ==>
RMF V3R1   XCF Path Statistics - RSPLEX0      Line 1 of 168
Scroll ==> CSR

Samples: 100   Systems: 7   Date: 02/17/23   Time: 10.36.40   Range: 100   Sec

Systems      CF Structure/      TYP ST Trans      Signals      Signals D
              CTC Devices              Class      Sent      Received

SYS1:SYS2    IXC_DEF8K(0030)      LST WR _XCFMGD      55          0 0
SYS1:SYS2    IXC_DEFAULT(0045)    LST WR _XCFMGD     364          0 0
SYS2:SYS1    IXC_DEFAULT(0044)    LST WR _XCFMGD     576          0 0
SYS2:SYS1    IXC_DEF8K(0031)     LST WR _XCFMGD     241          0 0
SYS1:SYS2    IXC_DEFAULT(0044)    LST WR DEFAULT      0         576 I
SYS1:SYS2    IXC_DEF8K(0031)     LST WR DEF8K         0         241 I
SYS2:SYS1    IXC_DEFAULT(0045)    LST WR DEFAULT      0         364 I
SYS2:SYS1    IXC_DEF8K(0030)     LST WR DEF8K         0          55 I

```

Figure 139. XCF Path statistics Report

Note: There is no graphic version of this report available.

Field descriptions of the XCF Path statistics report

Table 105. Fields in the XCF Path statistics Report	
Field Heading	Meaning
Systems	For outbound requests: The first name is the system from which the signals are sent. The second name is the system on which the signals are received. For inbound requests: The first name is the system on which the signals are received. The second name is the system from which the signals are sent.
CF Structure / CTC Devices	CF Structure or CTC Device Pair. - CTC - The device number pair being used as path - STR - The coupling facility structure name - LST - The coupling facility structure name and list number
Status	Path type. - CTC - Channel to Channel - STR - Coupling Facility Structure - LST - List within Coupling Facility Structure
ST	The status of the signaling path. - ST - starting - RS - restarting - WR - working - PP - stopped - WC - waitingForComp (waiting for completion) - NO - notOperational (not operational) - FL - failed - RB - rebuilding - QG - quiescing - QD - quiesced
Trans Class	Transport Class. The name of the transport class XCF uses for message transfer.
Signals Sent	The total number of outbound signals.
Signals Received	The total number of inbound signals.

Table 105. Fields in the XCF Path statistics Report (continued)

Field Heading	Meaning
D	Path direction. • O - outbound • I - inbound

XCFSYS - XCF Systems data

The XCF Systems data report provides information and statistics about transport classes and XCF message flow.

How to request the XCF Systems data report

You can request the XCF Systems data report by using one of the following options:

- Select **5** from the RMF Monitor III Primary Menu, and then select **20** on the RMF Sysplex Report Selection Menu.
- Enter the following command:

```
XCFSYS
```

Contents of the XCF Systems data report

```

Command ==>
RMF V3R1   XCF Systems Data - SYSPLEX1
Line 1 of 14
Scroll ==> CSR

Samples: 100   Systems: 2   Date: 01/01/22   Time: 00.00.00   Range: 100   Sec

Systems      Transport Class  Signals Sent  Signals Received  FIT %  SML %  BIG %  DEG %  DIR
SYS1:SYS1    XCFALL             177           0  95.5   0.0   4.5  100.0  L
SYS1:SYS1    DEFAULT            0             0   0.0   0.0   0.0   0.0  L
SYS1:SYS1    DEF8K              0             0   0.0   0.0   0.0   0.0  L
SYS1:SYS2    *ALL               0            1018  0.0   0.0   0.0   0.0  I
SYS1:SYS2    XCFALL             452           0  88.3   0.0  11.7  100.0  0
SYS1:SYS2    DEFAULT            0             0   0.0   0.0   0.0   0.0  0
SYS1:SYS2    DEF8K              0             0   0.0   0.0   0.0   0.0  0
SYS2:SYS2    XCFALL             4411          0  95.6   0.0   4.4  100.0  L
SYS2:SYS2    DEFAULT            0             0   0.0   0.0   0.0   0.0  L
SYS2:SYS2    DEF8K              0             0   0.0   0.0   0.0   0.0  L
SYS2:SYS1    *ALL               0             453  0.0   0.0   0.0   0.0  I
SYS2:SYS1    XCFALL             1020          0  63.2   0.0  36.8  100.0  0
SYS2:SYS1    DEFAULT            0             0   0.0   0.0   0.0   0.0  0
SYS2:SYS1    DEF8K              0             0   0.0   0.0   0.0   0.0  0

```

Figure 140. XCF Systems Data Report

There is no graphic version of this report available.

Field descriptions of the XCF Systems data report

Table 106. Fields in the XCF Systems data Report

Field Heading	Meaning
Systems	For outbound requests: The first name is the system from which the signals are sent. The second name is the system on which the signals are received. For inbound requests: The first name is the system on which the signals are received. The second name is the system from which the signals are sent.
Transport Class	Transport Class. The name of the transport class that XCF uses for message transfer.
Signals Sent	The total number of outbound signals.

Table 106. Fields in the XCF Systems data Report (continued)	
Field Heading	Meaning
Signals Received	The total number of inbound signals.
FIT%	The percentage of messages sent whose length fit the buffer size that supports the defined transport class length.
SML%	The percentage of messages sent whose length was smaller than the buffer size that supports the defined transport class length.
BIG%	The percentage of messages sent whose length exceeds the buffer size that supports the defined transport class length.
DEG%	The percentage of messages sent whose length exceeds the buffer size for which the signalling service was optimized.
DIR	Path direction. • L - local traffic • O - outbound • I - inbound

ZFSFS - zFS File System

The zFS File System Report measures zFS activity on the basis of single file systems. With this information, you can monitor DASD performance to ensure that there are no volumes or channels working near the limit of their capacity (space and workload, for example, I/O rates and response times).

How to request this report

To request the ZFS File System report, select **S** on the Primary Menu, and then select **14** on the Sysplex Report menu (shown in [Figure 5 on page 22](#)), or enter one of these commands:

```
ZFSFS
ZFF
```

Report Options

The Report Options panel for the zFS File System report allows you to specify options for this report.


```

                                RMF zFS File System Report Options                Line 1 of 285
Command ==>                                Scroll ==> CSR

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

Name   ==> OMVS.CB8B.JAVATEST.OUTPUT.ZFS
                                ALL or one of the available zFS file systems below
Detail ==> YES                    Show single system data (YES or NO) in ZFSFS report

                                Available zFS File Systems
APIRWW.DB2
NETVIEW.V6R1M0C.ZFS
OMVS.CB8A.JAVATEST.OUTPUT.ZFS
OMVS.CB8A.JAVATEST.ZFS
OMVS.CB8B.JAVATEST.OUTPUT.ZFS
OMVS.CB8B.JAVATEST.TESTSUITE.ZFS
OMVS.CB8B.JAVATEST.ZFS
OMVS.CB8C.JAVATEST.OUTPUT.ZFS
OMVS.CB8C.JAVATEST.TESTSUITE.ZFS
OMVS.CB8C.JAVATEST.ZFS
OMVS.CB8D.JAVATEST.OUTPUT.ZFS
OMVS.CB8D.JAVATEST.ZFS
OMVS.CB8E.JAVATEST.OUTPUT.ZFS
OMVS.CB8E.JAVATEST.TESTSUITE.ZFS
OMVS.CB8E.JAVATEST.ZFS

```

Figure 141. ZFSFS - Report Options

Name

Specify either ALL or the name of one of the zFS file systems available in the sysplex as shown in the field **Available File Systems**, which provides a list of all zFS file systems that are currently defined to the sysplex.

You can use an asterisk (*) as the last character of the file system name as a wild card. When a wild card is used, all file systems whose names start with the specified character sequence before the asterisk are reported on, no matter which characters follow.

Detail

Specify the desired level of detail in the zFS File System report:

NO

The report contains summary data for the sysplex only.

YES

The report contains data for the sysplex and all single systems.

If the list of file names is too long to fit on the first page, this report options panel can be scrolled up and down using function keys F7 and F8.

Contents of the report

```

RMF 3.1    zFS File System - UTCPLXCB
Command ==>
Line 1 of 570
Scroll ==> CSR

Samples: 120    Systems: 8    Date: 04/25/2023    Time: 09.04.00    Range: 120    Sec

----- File System Name -----
System      Owner      Mode      Size  Usg%      I/O  Resp  Read  XCF
Rate        Time      %        Rate

APIRWW.DB2
*ALL        CB86      RW S      1320M  1.0      <0.01  0.422  100  <0.01
NETVIEW.V6R1M0C.ZFS
*ALL        CB86      RW S      12M  49.8      <0.01  0.501  100  <0.01
OMVS.CB8A.JAVATEST.OUTPUT.ZFS
*ALL        CB8A      RW S      14G   0.2      1250  0.005  99.9  <0.01
OMVS.CB8A.JAVATEST.ZFS
*ALL        CB8A      RW S      4922M  0.7      21.98  0.004  100  0.025
OMVS.CB8B.JAVATEST.OUTPUT.ZFS
*ALL        CB8B      RW S      24G  15.6      1794  0.002  100  <0.01
OMVS.CB8B.JAVATEST.TESTSUIT.ZFS
*ALL        CB8B      RW S      989M  30.6      <0.01  1.418  100  <0.01
OMVS.CB8B.JAVATEST.ZFS
*ALL        CB8B      RW S      5868M  20.3      28.54  0.003  100  0.100
OMVS.CB8C.JAVATEST.OUTPUT.ZFS
*ALL        CB8C      RW S      9849M  0.3      <0.01  0.787  100  <0.01
OMVS.CB8C.JAVATEST.TESTSUIT.ZFS
*ALL        CB8C      RW S      989M  1.0      <0.01  1.102  100  <0.01
OMVS.CB8C.JAVATEST.ZFS

```

Figure 142. ZFSFS Report

Table 107 on page 241 describes the fields in this report.

If a file system name is specified and Detail is set to YES, then the performance values returned by each single system in the sysplex are reported underneath the *ALL summary row for this file system. A name value of ALL will report details for all file systems.

```

RMF 3.1    zFS File System - UTCPLXCB
Command ==>
Line 1 of 10
Scroll ==> CSR

Samples: 120    Systems: 8    Date: 04/25/2023    Time: 09.04.00    Range: 120    Sec

----- File System Name -----
System      Owner      Mode      Size  Usg%      I/O  Resp  Read  XCF
Rate        Time      %        Rate

OMVS.CB8B.JAVATEST.OUTPUT.ZFS
*ALL        CB8B      RW S      24G  15.6      1794  0.002  100  <0.01
CB8A        CB8B      RW S      0.0      0.000  0.000  0.0  0.000
CB8B        CB8B      RW S      24G  15.6      1794  0.002  100  0.000
CB8C        CB8C      CB8B      RW S      0.0      0.000  0.000  0.0  0.000
CB8D        CB8B      RW S      0.0      0.000  0.000  0.0  0.000
CB8E        CB8B      RW S      0.0      0.000  0.000  0.0  0.000
CB86        CB8B      RW S      0.0      0.000  0.000  0.0  0.000
CB88        CB8B      RW S      0.0      <0.01  0.000  100  <0.01
CB89        CB8B      RW S      0.0      0.000  0.000  0.0  0.000

```

Figure 143. ZFSFS Report - File system details

From the zFS File System Report, you can navigate to a variety of detail information using cursor-sensitive control. If you place the cursor on any of the lines with file system values, a pop-up window appears showing the details for this file system. [Figure 144 on page 241](#) shows an example of a pop-up panel with File System Details.

```

                                zFS File System Details

File System Name : OMVS.CB8B.JAVATEST.OUTPUT.ZFS
Mount
Point : /CB8B/javatest/output

System : CB8B                      Owner : CB8B                      Mode : RW S

----- Read -----
--- Appl --- --- XCF --- Aggr
Rate  Resp  Rate  Resp  Rate
Time  Time
1793  0.002  0.000  0.000  170.7

----- Write -----
--- Appl --- --- XCF --- Aggr
Rate  Resp  Rate  Resp  Rate
Time  Time
0.758  0.047  0.000  0.000  682.7

Vnodes          :    33      USS held vnodes      :    5
Open objects     :     0      Tokens                :    4
User cache 4k pages :     1      Metadata cache 8k pages :   90

ENOSPC errors    :     0      Disk I/O error        :    0
XCF comm. failures :     0      Cancelled operations :    0

Press Enter to return to the Report panel.

```

Figure 144. ZFSFS Report - zFS File System Details

Table 108 on page 242 describes the fields in this report.

Field descriptions

This topic describes the fields in the zFS File System report.

Table 107. zFS File System report field descriptions	
Field Heading	Meaning
File System Name	File system name.
System	Name of the system connected to the file system. In the first data line for a file system, the name is '*ALL' to indicate that this line shows the SYSPLEX view of the data rather than a single system view.
Owner	Name of owning system.
Mode	Mount mode of the file system. Possible values are: NM Not mounted. QS Not available because the aggregate is quiesced. RO Mounted in read-only mode. RW Mounted in read-write mode. The mount mode is followed by an S if the file system is using zFS sysplex sharing (RWSHARE).
Size	Maximum logical size of the file system (in Bytes).
Usg%	Percentage of currently used space by the file system.
I/O Rate	The rate of read and write requests per second (directory and file) made by applications to this file system.

Table 107. zFS File System report field descriptions (continued)

Field Heading	Meaning
Resp Time	Average response time in milliseconds for read and write requests made by applications to this file system.
Read%	Percentage of read operations contained in 'I/O Rate'.
XCF Rate	The rate of read and write XCF calls per second to the server.

Table 108. zFS File System report field descriptions - zFS File Systems Details panel

Field Heading	Meaning
File System Name	File system name.
Mount Point	Mount point of the file system.
System	Name of the system connected to the file system. If the name is *ALL, this indicates that this line shows the SYSPLEX view of the data rather than a single system view.
Owner	Name of owning system.
Mode	Mount mode of the file system. Possible values are: NM Not mounted. QS Not available because the aggregate is quiesced. RO Mounted in read-write mode. RW Mounted in read-write mode. The mount mode is followed by an S if the file system is using zFS sysplex sharing (RWSHARE).
Application read rate	The rate of read requests per second (directory and file) made by applications to this file system.
Application read resp time	The average response time for read requests made by applications to this file system (in milliseconds).
XCF read rate	The rate of read XCF calls per second to the server
XCF read resp time	The average response time required for a read XCF call to the server (in milliseconds).
Aggregate read rate	Read data transfer rate in bytes/second for the aggregate.
Application write rate	The rate of write requests per second (directory and file) made by applications to this file system.
Application write resp time	The average response time for write requests made by applications to this file system (in milliseconds).
XCF write rate	The rate of write XCF calls per second to the server.
XCF write resp time	The average response time required for a write XCF call to the server (in milliseconds).

Table 108. zFS File System report field descriptions - zFS File Systems Details panel (continued)

Field Heading	Meaning
Aggregate write rate	Write data transfer rate in bytes/second for the aggregate. The write data transfer rate also includes write activity from zFS daemons for file systems mounted in R/W mode.
Vnodes	Number of vnodes in memory for the file system.
USS held vnodes	Number of vnodes that the logical file system layer, also known as USS, holds for the file system.
Open objects	The number of files and directories currently open.
Tokens	The number of tokens held for objects in the file system from the token manager.
User cache 4K pages	The number of 4K pages in the user file cache for this file system.
Metadata cache 8K pages	The number of 8K pages in the metadata cache for this file system.
ENOSPC errors	The number of ENOSPC errors seen by applications for this file system on this system.
Disk I/O errors	The number of disk I/O errors for disk I/Os performed by this system.
XCF comm. failures	The number of XCF communication timeouts or errors seen by XCF requests made for this file system on this system.
Cancelled operations	The number of times a task was asynchronously ABENDED (CANCEL) or EOMd (FORCE) while running an operation in this file system.

ZFSKN - zFS Kernel report

The zFS Kernel report provides a variety of measurements counting the calls made to zFS from z/OS UNIX and the average response time of zFS requests. This information gives the basic measure of zFS performance and can be used to determine the appropriate tuning options needed to make best use of the z/OS File System (zFS).

How to request this report

To request the ZFS Kernel report, select **S** on the Primary Menu, and then select **15** on the Sysplex Report menu (shown in [Figure 5 on page 22](#)), or enter one of these commands:

```
ZFSKN
ZFK
```

Contents of the report

Command ==>		RMF 3.1	zFS Kernel	- UTCPLXCB	Line 1 of 8 Scroll ==> CSR	
Samples: 120	Systems: 8	Date: 04/25/2023	Time: 09.04.00	Range: 120	Sec	
System Name	- Request Rate - Local Remote	--- XCF Rate --- Local Remote	- Response Time - Local Remote			
CB8A	1356 1305	0.025 82.57	8.000 955.0			
CB8B	2552 120.8	0.000 0.400	13057 1027			
CB8C	11.33 0.175	0.000 0.000	22.00 106.0			
CB8D	22315 1245	0.000 0.592	148.1 3115			
CB8E	1747 2099	0.000 75.46	186.0 1214			
CB86	29605 418.1	0.000 34.98	354.0 4921			
CB88	1921 260.7	0.000 43.13	183.0 1738			
CB89	517.3 2590	0.000 75.33	97.00 1184			

Figure 145. zFS Kernel Report

Field descriptions

This topic describes the fields in the zFS Kernel report.

Table 109. zFS Kernel report field descriptions	
Field Heading	Meaning
System Name	Name of the system running zFS. In the context of requests against file systems, this is the name of the requesting system.
Request Rate	Rate of zFS requests during the report interval for file systems which are locally and remotely owned. A file system is locally owned if the requesting system is also the owner of the file system. It is remotely owned if the owner of the file system is not the requesting system.
XCF Rate	Rate of zFS requests during the report interval requiring data from another system by XCF, both for locally and remotely owned file systems.
Response Time	Average time in milliseconds required for the completion of the zFS requests during the report interval for locally and remotely owned file systems.

ZFSOVW - zFS Overview Report

To use a zFS file system within a z/OS UNIX file system hierarchy to its full capacity, it is necessary to apply appropriate tuning options. The zFS performance especially depends on a suitable tailoring of its cache sizes to reduce I/O rates and path lengths. The performance can also be improved by adapting available disk space.

This report provides a summary of zFS activity, request and DASD statistics on the current system and thus helps to control and tune the zFS environment. For example, you can use the HIT% values in the Cache Activity section as an indication of whether the current cache sizes are sufficient.

How to request this report

To request the zFS Overview report, select **S** from the Primary Menu, then select **13** from the Sysplex Report Selection Menu (shown in [Figure 5 on page 22](#)) or enter one of the following commands:

```
ZFS0VW
ZFO
```

Contents of the report

This topic shows the content of the zFS Overview report.

RMF 3.1 zFS Overview - SYSDPLEX Line 1 of 8											
Command ==> Scroll ==> CSR											
Samples: 120 Systems: 8 Date: 04/25/2023 Time: 09.04.00 Range: 120 Sec											
System	-----Wait%-----			-----Cache Activity-----				-----Metadata-----			
	I/O	Lock	Sleep	---User---	Rate	Hit%	--Vnode---	Rate	Hit%	Rate	Hit%
CB8A	100	1.4	100	954.6	99.5		1350	94.8		236.5	96.0
CB8B	5.2	0.2	0.0	1226	97.8		1967	86.7		4427	96.1
CB8C	0.0	----	0.1	61.32	100		649.7	100		62.88	98.3
CB8D	0.1	0.1	0.2	893.1	100		15393	100		1102	97.1
CB8E	100	0.0	7.2	1205	100		2287	97.2		195.1	96.6
CB86	9.4	0.1	0.1	2326	34.5		18660	91.8		42145	91.8
CB88	5.5	0.0	2.5	1242	100		676.6	97.0		96.36	97.8
CB89	100	0.3	46.0	386.7	100		2373	97.4		273.9	97.3
F1=HELP F2=SPLIT F3=END F4=RETURN F5=RFIND F6=TOGGLE											
F7=UP F8=DOWN F9=SWAP F10=BREF F11=FREF F12=RETRIEVE											

Figure 146. zFS Overview Report

Field descriptions

Table 110. zFS Overview report field descriptions	
Field Heading	Meaning
System	Name of the system running zFS.
Wait%	These Wait percentages are reported: I/O Percentage of time that zFS requests had to wait for I/O completion. Lock Percentage of time that zFS requests had to wait for locks. Sleep Percentage of time that zFS requests had to wait for events. Dashes (----) in these fields indicate that RMF is unable to calculate a reasonable value.
Cache Activity section	

Table 110. zFS Overview report field descriptions (continued)

Field Heading	Meaning
User	The user file cache is for caching regular user files that are larger than 7K. The measured statistics have these meanings: Rate Total number of read and write requests per second made to the user file cache. Hit% Percentage of read and write requests to the user file cache that completed without accessing the DASDs.
Vnode	The vnode cache is used to hold virtual inodes. An inode is a data structure related to a file in the file system, holding information about the file's user and group ownership, access mode and type. The measured statistics have these meanings: Rate Number of read and write requests per second made to the vnode cache. Hit% Percentage of read and write requests to the vnode cache that completed without accessing the DASDs.
Metadata	The metadata cache is used for file system metadata and for files smaller than 7K. It resides in the primary z/FS address space. The measured statistics have these meanings: Rate Number of read and write requests per second made to the metadata cache. Hit% Percentage of read and write requests to the metadata cache that completed without accessing the DASDs.

Navigating to details in the zFS Overview report

From the **zFS Overview report**, you can navigate to a variety of detail information using cursor-sensitive control.

The detail information is provided in panels shown in [Figure 147 on page 247](#) through [Figure 150 on page 250](#). You reach these panels as follows:

- From the **Wait% -I/O** field, you can reach the **I/O Details by Type** panel ([Figure 147 on page 247](#)).
- From any value in the **Cache Activity - User** section, you can reach the **User Cache Details** panel ([Figure 148 on page 247](#)).
- From any value in the **Cache Activity - Vnode** section, you can reach the **Vnode Cache Details** panel ([Figure 149 on page 249](#)).
- From any value in the **Cache Activity - Metadata** section, you can reach the **Metadata Cache Details** panel ([Figure 150 on page 250](#)).

zFS Overview - I/O Details by Type

The **zFS Overview - I/O Details by Type** report displays a breakdown of I/O requests into the following types:

- I/O for file system metadata
- I/O for log data
- I/O for user file data




Figure 147. I/O Details by Type

Table 111. Fields in the zFS Overview report - I/O details	
Field Heading	Meaning
Count	Total number of I/O requests of the indicated type.
Waits	Number of zFS requests waiting for an I/O completion of the indicated I/O type.
Canc1	Number of cancelled zFS requests during an I/O request of the indicated type, for example, a user tried to delete a file during a pending I/O to this file's metadata.
Merge	Number of merges of two I/O requests into a single request because of better performance.
Type	Type of the I/O request (I/O for metadata, log data or user file data).

zFS Overview - User Cache Details

The user file cache is for caching regular user files that are larger than 7K. The **zFS Overview - User Cache Details** report displays the following details of the user file cache activity:




Figure 148. User Cache Details

Table 112. Fields in the zFS Overview Report - User Cache Details

Field Heading	Meaning
System	Name of the system running zFS.
Size	Total size of the user file cache.
Total Pages	Total number of pages in the user file cache.
Free Pages	Total number of free pages in the user file cache.
Segments	Total number of allocated segments in the user file cache.
Storage fixed	Shows whether the size of the user file cache storage is fixed. If the zFS parameter <code>user_cache_size</code> is set to 'fixed', then zFS reserves real storage for use by zFS only. The 'fixed option' helps to improve performance during data access and can be applied if you have enough real storage available.
Read Rate	Number of read requests per second made to the user file cache.
Read Hit%	Percentage of read requests to the user file cache that completed without accessing the DASD.
Read Dly%	Percentage of delayed read requests to the user file cache. A read request is delayed if it must wait for pending I/O, for example, because the file is in a pending read state due to asynchronous read ahead from DASD to the user file cache.
Async Read Rate	Number of read aheads per second.
Write Rate	Number of write requests per second made to the user file cache.
Write Hit%	Percentage of write requests to the user file cache that completed without accessing the DASD.
Write Dly%	Percentage of delayed write requests to the user file cache. The following reasons are counted as write request delays: Write wait Write must wait for pending I/O. Write faulted Write to a file needs to perform a read from DASD. If a write-only updates a part of a file's page, and this page is not in the user file cache, then the page must be read from DASD before the new data is written to the cache.
Scheduled Write Rate	Number of scheduled writes per second.
Read%	Percentage of read requests, based on the sum of read and write requests.
Dly%	Percentage of delayed requests, with the following events counted as delays: • Read wait: a read request must wait for a pending I/O operation. • Write wait: a write request must wait because of a pending I/O operation. • Write faulted: a write request to a file in the user file cache needs to perform a read operation from DASD before writing, because the required page of that file is currently not in the cache.
Page Reclaim Writes	Total number of page reclaim writes. A page reclaim write action writes one segment of a file from the user file cache to DASD. Page reclaim writes are performed to reclaim space in the user file cache. If page reclaim writes occur too often in relation to the write rate, then the user file cache may be too small.
Fsyncs	Total number of requests for file synchronization (fsync) between user file cache and DASD.

zFS Overview - Vnode Cache Details

The vnode cache is used to hold virtual inodes. An inode is a data structure related to a file in the file system, holding information about the file's user and group ownership, access mode, and type. The **zFS Overview - Vnode Cache Details** report displays the following details of the vnode cache activity:




Figure 149. Vnode Cache Details

Table 113. Fields in the zFS Overview Report - Vnode Cache Details	
Field Heading	Meaning
System	Name of the system running zFS.
Size	Number of vnodes that will be initially cached by zFS.
Vnodes Total	Number of currently allocated vnodes in the vnode cache. If more vnodes are requested than are currently available, then zFS dynamically allocates more vnodes.
Vnodes Size	Size of a vnode data structure in bytes.
Vnodes Ext.#	Number of extended vnodes.
Vnodes Ext. Size	Size of an extended vnode data structure in bytes.
Vnodes Open	Number of currently open vnodes.
Vnodes Held	Number of vnodes currently held in zFS by USS.
Requests Total	Number of requests to the vnode cache.
Requests Rate	Number of requests per second made to the vnode cache.
Requests Hit%	Percentage of requests to the vnode data that found the target vnode data structures in the vnode cache. High hit rates indicate a favorable zFS environment, because each miss involves initialization of vnode data structures in the vnode cache.
Requests Alloc	Number of requests to create new vnodes (for operations such as create or mkdir).
Requests Delete	Number of requests to delete vnodes (for operations such as remove or failed creates or mkdirs).

zFS Overview - Metadata Cache Details

The metadata cache is used to contain all file system metadata; this metadata includes all directory contents, file status information, and file system structures. Additionally, it also caches data for files smaller than 7 K. It resides in the primary zFS address space.

The **zFS Overview - Metadata Cache Details** report displays these details of the metadata cache:




Figure 150. Metadata Cache Details

Table 114. Fields in the zFS Overview Report - Metadata Cache/Backing Cache Details	
Field Heading	Meaning
System	Name of the system running zFS.
Size	Total size of the metadata cache.
Buffers	Total number of buffers in the metadata cache. The buffer size is 8K.
Storage fixed	Shows whether the size of the metadata cache storage is fixed. If the zFS parameter meta_cache_size is set to 'fixed', then zFS reserves real storage for use by zFS only. The 'fixed option' helps to improve performance during data access and can be applied if there is enough real memory available.
Total	Number of requests made to the metadata cache.
Rate	Number of requests per second made to the metadata cache.
Hit%	Percentage of requests to the metadata cache completing without accessing the DASD.
Updates	Number of updates made to buffers in the metadata cache.
Partial writes	Number of times that only half of an 8K metadata block needed to be written.

Chapter 3. Snapshot reporting with Monitor II

This topic describes the following information about Monitor II reports:

- [“Monitor II sessions” on page 251](#)
- [“Structure of Monitor II reports” on page 252](#)
- [“ARD/ARDJ - Address Space Resource Data report” on page 256](#)
- [“ASD/ASDJ - Address Space State Data report” on page 260](#)
- [“ASRM/ASRMJ - Address Space SRM Data report” on page 264](#)
- [“CHANNEL - Channel Path Activity report” on page 267](#)
- [“DEV/DEVV - Device Activity report” on page 270](#)
- [“HFS - Hierarchical File System Statistics report” on page 276](#)
- [“ILOCK - IRLM Long Lock Detection report” on page 279](#)
- [“IOQUEUE - I/O Queuing Activity report” on page 281](#)
- [“LLI - Library List report” on page 285](#)
- [“OPT - OPT Settings report” on page 288](#)
- [“PGSP - Page Data Set Activity report” on page 290](#)
- [“SDS - Sysplex Data Server report” on page 292](#)
- [“SENQ - System Enqueue Contention report” on page 293](#)
- [“SENQR - System Enqueue Reserve report” on page 297](#)
- [“SPAG - Paging Activity report” on page 299](#)
- [“SRCS - Central Storage/Processor/SRM report” on page 301](#)

Monitor II sessions

You can **display** a Monitor II report during:

- An ISPF display session

This session is started with the command

```
RMF
```

This leads to the RMF Primary menu, then you select **2** to get the Monitor II ISPF session.

- A TSO/E display session

This session is started with the TSO/E command

```
RMFMON
```

- A background session

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

```
MODIFY RMF,START AB
```

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

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

In all sessions, you can get the same reports. There is just a small difference in the syntax used to call them:

- Display Session

The reports are called via *commands* according to TSO/E syntax rules:

Example: ASD T,A

- Background Session

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

Example: ASD(T,A)

This chapter shows report examples from an ISPF session, the report format of a RMFMON session is very similar, and the meaning of all report fields is the same in all versions of a report.

Structure of Monitor II reports

This chapter presents sample reports and the meaning and contents of each field in each report. The sample reports show the display screen contents from an ISPF session for each report.

When the reports are printed, the contents are identical to the report contents shown on the screen with some differences in the layout of the printed output.

Contents of the Monitor II report header

A Monitor II report header looks different, depending on whether you use the ISPF interface or the TSO/E interface.

If you are using an ISPF session

Each report consists of

- A header line identifying the report
- A line for commands and scroll amount field
- A status line for CPU, UIC, and PR. This line also contains the SMF system ID and the current setting of the report mode (Total or Delta).
- A variable number of data lines.

If you are using a TSO/E session

Each report consists of

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

Figure 151 on page 253 shows the different report areas for a TSO/E session. For a description of each area, see [Table 115 on page 253](#).




Figure 151. Header of a Monitor II TSO/E session report

When you begin a session, the cursor appears in the **input area**. During the session you issue all display commands from this area. Other areas indicated in the figure are described in [Table 115 on page 253](#).

Table 115. Monitor II Display Session Areas	
Area	Function
Report title	The type of measurement data.
F	Indicates more pages.
Input Area	Issue all commands from here.
Status/Message Area	This area contains: CPU Current average processor utilization. This information depends on the activity of Monitor I. If Monitor I CPU gathering is active, the header line shows two views separated by a slash (/): • The MVS view of the CPU utilization • The LPAR view of the CPU utilization If Monitor I CPU gathering is not active, the header line shows: • The SRM view of the CPU utilization. • '***' due to missing CPU measurement data for the LPAR view For more information about the different views of CPU utilization refer to “CPU - CPU Activity report” on page 347. UIC The current system unreferenced interval count. Values greater than 9999 are displayed as nnK to indicate a multiple of 1000. The maximum value is 65K. PR The rate of page-ins per second excluding swap-ins, VIO (virtual input/output), reclaims, and hiperspaces. System The SMF system ID of this system.
Report Name Area	The report name.
Mode Area	The current setting for the report mode (either D for delta or T for total) and hardcopy mode (either H for hardcopy, or blank).

Table 115. Monitor II Display Session Areas (continued)	
Area	Function
Header Area	Consists of two lines of column headings that identify the data fields included in the report.
Output Area	Contains the report data.

Different formats of Monitor II reports

Monitor II offers two types of reports:

- **Table Reports** - Example: ASD Report

Table reports have a variable number of data lines.

- **Row Reports** - Example: ASDJ Report

Row reports have only one line of data. When you request a row report repeatedly, each request adds one line of data to the display. You can use the repetitive requests to build a table of information.

Different modes of Monitor II reports

Monitor II offers two modes for the session reports:

- **Total** mode

A total mode report shows the cumulative total since the beginning of the Monitor I interval.

- **Delta** mode

A delta report mode shows the change in the activity since the previous request for the report.

Monitor II display session reports

For a Monitor II display session, RMF creates a single output data set for each session. All **printed output** resulting from either hardcopy mode or the non-ISPF Print command is sent to the same output data set.

You need to allocate this data set before starting the display session:

```
ALLOC F(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 *z/OS Resource Measurement Facility User's Guide* for 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 152 on page 255 shows an example of printed output from a display session.


```

RMF Monitor II HARDCOPY LOG

SESSION NAME TSO

z/OS 3.1          SYSTEM ID SYS1      DATE 04/25/2023
                   RPT VERSION 3.1 RMF  TIME 17.12.25

+++++++
++ MIG= 729K CPU= 41 UIC= 254 PR= 0 SYSTEM=SYS1 ++++++ TOTAL MODE ++++++ ARD ++++++
++ 0 SYSTEM=SYS1 ++++++ 17.12.28 ++++++

17:12:28 DEV  FF  FF PRIV LSQA X C SRM TCB  CPU  EXCP SWAP LPA CSA NVI V&H
JOBNAME  CONN 16M 2G  FF  CSF M R  ABS  TIME  TIME  RATE  RATE  RT  RT  RT  RT
*MASTER* 11432 0 493 1261 116 0.0 7973.9 17944 -----
PCAUTH 0.000 0 94 2 30 X 0.0 0.04 0.05 -----
RASP 0.000 --- 30 --- --- X 0.0 0.02 27.48 -----
TRACE 0.000 0 145 3 49 X 0.0 0.02 0.03 -----
XCFAS 1470 0 2407 2457 2087 X 0.0 245.4 796.20 -----
GRS 29514 17 17 38 33 X 0.0 2135.2 10346 -----
SMXC 0.000 0 47 2 20 0.0 44.81 51.38 -----
SYSBMAS 0.000 0 13 5 19 0.0 0.01 0.02 -----
DUMPSRV 278.1 0 100 2 37 0.0 19.30 61.10 -----
CONSOLE 1148 0 44 2 31 X 0.0 1189.7 1365.2 -----
ANTMAIN 0.085 0 86 2 23 X 0.0 0.02 0.05 -----
ALLOCAS 0.000 0 78 2 21 X 0.0 0.02 0.03 -----
SMF 324.0 0 49 2 36 X 0.0 1.32 43.32 -----
VLF 0.956 1 117 32 34 X 0.0 231.96 235.34 -----
LLA 215.2 10 9 12 56 X 0.0 80.48 96.14 -----
DFRMM 265.3 0 32 2 61 X 0.0 80.95 98.33 -----
FTPSRV01 0.783 1 1 32 S 0.0 7.22 7.64 -----

```

Figure 152. Format of printed reports from a Monitor II display session

The report was printed as a result of a HARDCOPY ON or HARDCOPY command. This command puts the session in hardcopy mode and causes all reports to be displayed and a printable version to be sent to the output data set.

Monitor II background session reports

For a Monitor II background session, RMF creates a single output data set for each report requested. All **interval reports** for that measurement activity are written to a single output data set. Thus, if you request three measurements for a session with five reporting intervals, RMF creates three data sets and writes five reports to each data set.

The printed output generated for each **table report** requested is preceded by a line of plus signs (+). Fields within the line of plus signs indicate the option and any operands specified for the report and whether the report is in delta mode or total mode. Each iteration of the report is separated by a line of plus signs. A field within the line of plus signs indicates the time of the report. The column headings are repeated for each iteration of the report.

The printed output generated for each **row report** is preceded by a single line of plus signs (+). Fields within the line of plus signs indicate the option and any operands specified for the report and whether the session is in delta or total mode. Headings appear only once. There is no delimiter between successive iterations of a row report, and a field within the report line indicates the time of the report.

Figure 153 on page 256 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	SYSTC	1	NS	76	33	0	33	0	11	0	0	+0	+0	+0		
TRACE	SYSTC	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	SYSTC	1	NS	71	170	0	170	0	30	0	150	+0	+0	+0		
RMF33	SYSTC	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	SYSTC	1	NS	FC	1350	0	1350	0	33	0	150	+0	+0	+0		
BOYLEMM	BATCH	2	IN	78	85	0	133	0	0	4	101	+0	+0	+0		
SMF	SYSTEM	1	NS	FF	94	0	94	0	11	0	150	+0	+0	+0		
DFHSM	SYSTC	1	NS	74	610	0	610	0	527	1	150	+0	+0	+0		
VTAM	SYSTC	1	NS	FD	678	0	678	0	34	0	0	+0	+0	+0		
SOS	SYSTC	1	IN	79	76	0	567	0	0	1.1K	150	+0	+0	+0		
AMSAQFT	SYSTC	1	NS	72	54	0	54	0	33	0	150	+0	+0	+0		
JES2	SYSTC	1	NS	FE	900	0	900	0	24	0	0	+0	+0	+0		
CATALOG	SYSTEM	1	NS	FF	1552	0	1552	0	11	0	150	+0	+0	+0		
+++++++ 14.36.02 ++++++																
14:36:02	S	C	R	DP	CS	ESF	ESF	TAR	WS	TX	WRK	CPU	I/O	STM		
JOBNAME	SRVCLASS	P	L	LS	PR	F	+CS	WSS	IN	SC	RV	RV	RV	RV		
HUBERF	TSO	2	IN	78	498	0	597	0	0	3	0	+0	+0	+0		
ZAPPERD	TSO	2	IN	72	110	0	121	0	0	1	150	+0	+0	+0		
PTRACYB	TSO	2	IN	78	174	0	229	0	0	3	0	+0	+0	+0		
IRLMPROC	SYSTC	1	NS	FC	61	0	61	0	33	0	0	+0	+0	+0		

Figure 153. Format of printed reports from a Monitor II background session

ARD/ARDJ - Address Space Resource Data report

The ARD and ARDJ reports 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. The information provided in these reports includes, for example, information on processor time, paging, and central storage.

The ARD report enables you to determine which jobs are creating performance problems.

When you identify a problem job, you can request an ARDJ report for that particular job. This enables you to focus your reporting on a known problem area.

How to request the ARD report

Different methods are used to request the ARD and ARDJ reports.

How to request an ARD report

- In ISPF, specify **1** on the Address Space Report Selection menu.
- In TSO/E, use **PF1** to select the ARD report.
- Command interface:

Display session

```
ARD [class,status,,boost]
```

Background session

```
ARD [(class,status,,boost)]
```

How to request an ARDJ report

- In ISPF, specify **4** on the Address Space Report Selection menu.
- Command interface:

Display session

```
ARDJ jobname
```

Background session

```
ARDJ (jobname)
```

Contents of the report

The information shown in an ARD and an ARDJ report is identical, except the content of the first column, which is:

JOBNAME for the ARD report

TIME for the ARDJ report

In the ARD report of Figure 154 on page 257, the number of data lines in the report depends on the number of address space identifiers in the system that meet your selection criteria. The shown report is a sample for a system running in z/Architecture.

```

RMF - ARD Address Space Resource Data
Command ===>
Line 1 of 85
Scroll ===> HALF

CPU= 9      UIC= 255 PR= 13      System= SYS1 Total

14:51:59 DEV  FF  FF PRIV LSQA X C SRM TCB  CPU  EXCP SWAP LPA CSA NVI V&H
JOBNAME  CONN 16M 2G  FF  CSF M R ABS TIME TIME RATE RATE RT RT RT RT
*MASTER* 1476 0 493 1261 110 0.0 149.0 544.5 0.29 0.00 0.0 0.0 0.0 0.0
PCAUTH 0.000 0 94 0 112 X 0.0 0.00 0.00 0.00 0.00 0.0 0.0 0.0 0.0
RASP 0.000 0 30 207 44 X 0.0 0.00 1.90 0.00 0.00 0.0 0.0 0.0 0.0
TRACE 0.067 0 145 1 173 X 0.0 0.00 0.00 0.00 0.00 0.0 0.0 0.0 0.0
DUMPSRV 14.69 0 33 0 76 0.0 1.01 1.73 0.00 0.00 0.0 0.0 0.0 0.0
XCFAS 1470 0 2407 2457 2087 X 0.0 245.4 796.2 2.51 0.00 0.0 0.0 0.0 0.0
GRS 0.000 0 47 52 1047 X S 0.0 235.0 339.8 0.00 0.00 0.0 0.0 0.0 0.0
SMXC 2.400 0 13 0 50 0.0 9.44 11.10 0.00 0.00 0.0 0.0 0.0 0.0
SYSBMAS 0.000 0 100 104 31 0.0 4.74 4.83 0.00 0.00 0.0 0.0 0.0 0.0
CONSOLE 25.23 0 44 19 101 X 0.0 61.85 66.37 0.07 0.00 0.0 0.0 0.0 0.0
WLM 0.493 0 86 52 503 X 0.0 2130 2241 0.00 0.00 0.0 0.0 0.0 0.0
ANTMAIN 0.934 0 78 3 163 X X 0.0 2.60 3.05 0.00 0.00 0.0 0.0 0.0 0.0
ANTAS000 0.621 0 49 2 100 X 0.0 0.07 0.08 0.00 0.00 0.0 0.0 0.0 0.0
OMVS 168.5 0 117 85 647 X 0.0 22.57 26.50 0.00 0.00 0.0 0.0 0.0 0.0
IEFSCHAS 0.000 0 9 0 34 X 0.0 0.00 0.00 0.00 0.00 0.0 0.0 0.0 0.0
JESXCF 1.629 0 32 4 78 X 0.0 24.64 42.03 0.00 0.00 0.0 0.0 0.0 0.0

```

Figure 154. ARD Report in z/Architecture

Field descriptions

Table 116. Fields in the ARD and ARDJ reports	
Field Heading	Meaning
hh:mm:ss JOBNAME (ARD report)	The time the report was requested, and the one to eight character jobname associated with the address space.
jobname TIME (ARDJ report)	The name of the job you requested. This column contains the time the report was requested.
DEV CONN	The device connect time (in seconds) used by the job. If it is greater than 99999 seconds, then it is expressed in hours (a four-digit number with a floating decimal point followed by H). If the device connect time exceeds 76.4 hours, three asterisks will appear in the field.
FF 16M	The number of fixed pages below 16 megabytes for the job. If the data gathered is not valid, dashes will appear in this and the following fields.

Table 116. Fields in the ARD and ARDJ reports (continued)	
Field Heading	Meaning
FF 2G	The number of fixed pages for the job between 16 megabytes and 2 gigabytes.
PRIV FF	The number of private non-LSQA pages for the job.
LSQA CSF	The number of private LSQA fixed pages for the job in central storage. This value includes fixed frames and private DREF pages.
X M	The cross memory address space indicator. When the field contains X, the line of data describes a cross memory address space; that is, an address space accessed primarily from other address spaces by means of cross memory functions. If it is not a cross memory address space, the field is blank.
C R	An indication whether WLM managed the address space as <i>storage critical</i> and/or <i>CPU critical</i> during the reporting interval. S Storage critical C CPU critical X Both storage and CPU critical
SRM ABS	The total SRM service absorption rate for the job. This field is reported only for address spaces that are currently in central storage. If no data is reported, dashes will appear in the field.
TCB TIME	The number of seconds of TCB processor time used by the current job step.
CPU TIME	The amount of processor (TCB + SRB) time, in seconds, for the current job step. When a valid delta value cannot be computed because the job has changed steps between requests for the report, this field contains dashes when delta mode is in effect.
EXCP RATE	The EXCP rate. This field always contains the rate since the last report request.
The following fields always contain a value that reflects the change since the last report request. They are reported only for address spaces that are currently in central storage.	
SWAP RATE	The page rate (the sum of pages in and pages out) for the job.
LPA RT	The common LPA page-in rate for the current transaction.
CSA RT	The common CSA page-in rate for the current transaction.
NVI RT	The private non-VIO page rate (the sum of pages in and pages out) for the current transaction.
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.
Only boosted   ==> NO     Specify YES to show only boosted address spaces.
                        Only applicable to ARD and ASD reports.
```

Figure 155. 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.

Only boosted

Allows you to specify the boost status of the address spaces shown on the report. The default is **NO** to include address spaces regardless of their boost status. **YES** causes only boosted address spaces to be shown. The option is not applicable to ASRM report. Boosted address spaces will be highlighted in reversed turquoise in ISPF.

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 156. ARDJ, ASDJ, and ASRMJ Report Options Panel

Jobname

The jobname must:

- Be one to eight characters long
- Consist of the characters A-Z, 1-9, and the special characters @, #, and \$

The exception to these rules is *MASTER*, which is also a valid jobname.

ASD/ASDJ - Address Space State Data report

The ASD/ASDJ report gives an overview of the current state of each address space in the system or each address space that meets the selection criteria that you specify when you request the report. Basically, the report tells you where each address space is and what it is doing.

You can use the ASD report, for example, to determine which jobs are using large amounts of central storage or which jobs are being swapped excessively and why the swapping is occurring.

Once a problem job has been identified, you can request an ASDJ report for that particular job. This enables you to focus your reporting on a known problem area.

If you have a workload delaying your application, you can check the workloads dispatching priority (**DP PR**) on the ASD report, and change it if necessary.

How to request the ASD report

Different methods are used to request the ASD and ASDJ reports.

How to request an ASD report

- In ISPF, specify **2** on the Address Space Report Selection menu.
- In TSO/E, use **PF2** to select the ASD report.
- Command interface:

Display session

```
ASD [class,status,,boost]
```

Background session

```
ASD [(class,status,,boost)]
```

How to request an ASDJ report

- In ISPF, specify **4** on the Address Space Report Selection menu.
- Command interface:

Display session

```
ASDJ jobname
```

Background session

```
ASDJ (jobname)
```

Contents of the report

The information shown in an ASD and an ASDJ report is identical except the heading for the first column which is:

JOBNAME for the ASD report
TIME for the ASDJ report

In the ASD report of [Figure 157 on page 261](#), 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 158 on page 261](#), 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 - ASD Address Space State Data
Line 1 of 53
Command ==> Scroll ==> HALF

CPU= 3/ 3 UIC= 65K PR= 0 System= SYS1 Total

11:08:36 S C R DP CS CS TAR X PIN TX SWAP WSM
JOBNAME SRVCLASS P L LS PR F TAR WSS M RT SC RV RV

*MASTER* SYSTEM 1 NS FF 6225 0 ---- 0 0
PCAUTH SYSTEM 1 NS FF 141 0 X ---- 0 0
RASP SYSTEM 1 NS FF 360 0 X ---- 0 0
TRACE SYSTEM 1 NS FF 875 0 X ---- 0 0
DUMPSRV SYSTEM 1 NS FF 862 0 ---- 0 0

```

Figure 157. ASD Report

```

RMF - ASDJ Address Space State Data
Line 1 of 7
Command ==> Scroll ==> HALF

CPU= 2/ 2 UIC= 65K PR= 0 System= SYS1 Total

RMFGAT S C R DP CS CS TAR X PIN TX SWAP WSM
TIME SRVCLASS P L LS PR F TAR WSS M RT SC RV RV

11:03:44 SYSSTC 1 NS FE 46.5K 0 X ---- 0 0

```

Figure 158. ASDJ Report

Field descriptions

Table 117. Fields in the ASD and ASDJ Report	
Field Heading	Meaning
hh:mm:ss JOBNAME	The time the ASD report was requested, in the form hh:mm:ss, and the one to eight character jobname associated with the address space.
SRVCLASS	The service class name.
S P	Service class period.

Table 117. 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 (for example, 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

Table 117. Fields in the ASD and ASDJ Report (continued)

Field Heading	Meaning
R LS	The field showing the reason for the last swap out associated with the job. This field is blank when the current location of the job is NS, IN, or PR. The possible contents and their meanings are: TO Terminal output wait TI Terminal input wait LW Long wait XS Auxiliary storage shortage RS Central storage shortage DW Detected wait MP Memory Pool shortage 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
DP PR	The dispatching priority for the job.
CS F	The number of central storage frames assigned to the job. If a job is swapped out of central storage the number represents the number of central storage frames assigned to the job before the swap out occurred.
CS TAR	The SRM central storage target value. If no special monitoring is requested, this field is blank.
TAR WSS	The target working set size for the job (in number of pages).
X M	The cross memory address space indicator. When the field contains X, the line of data describes a cross memory address space; that is, an address space accessed primarily from other address spaces by means of cross memory functions. If it is not a cross memory address space, this field is blank.

Table 117. 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: $\text{PIN} = \frac{\# \text{ Page-ins}}{\text{Transaction Residency Time}}$ - For all other address spaces, the calculation is: $\text{PIN} = \frac{\# \text{ Page-ins}}{\text{Seconds of CPU Time}}$ 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. Dashes (----) in this field indicate 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.
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 259 for a description.

ASRM/ASRMJ - Address Space SRM Data report

The ASRM/ASRMJ report gives an overview of the system resources that are used by each address space in the system or each address space that meets the selection criteria that you specify when you request the report. The report gives, for example, information on processor service, storage service, and I/O service.

The report enables you to determine which jobs are using which services and whether certain jobs are creating performance problems by making excessive use of system services.

How to request this report

Different methods are used to request the ASRM and ASRMJ reports.

How to request an ASRM report

- In ISPF, specify **3** on the Address Space Report Selection menu.
- In TSO/E, use **PF3** to select the ASRM report.

- Command interface:

Display session

```
ASRM [class,status]
```

Background session

```
ASRM [(class,status)]
```

How to request an ASRMJ report

- In ISPF, specify **6** on the Address Space Report Selection menu.
- Command interface:

Display session

```
ASRMJ jobname
```

Background session

```
ASRMJ (jobname)
```

Contents of the report

The information shown in an ASRM and an ASRMJ report is identical except the contents of the first column which is:

JOBNAME for the ASRM report

TIME for the ASRMJ report

In the ASRM report shown in [Figure 159 on page 265](#) 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 JOBNAME	SRVCLASS	S P	TRANS ACTIVE	TRANS CUR RES	TX CT	TX SC	TX CPU	TX MSO	TX IOC	TX SRB	SESS 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 159. ASRM Report

```

RMF - ASRMJ Address Space SRM Data
Command ==>
Line 1 of 14
Scroll ==> HALF

CPU= 37/ 35 UIC=2540 PR= 0
System= SYS1 Total

BARE    S TRANS  TRANS  TX  TX  TX  TX  TX  TX  TX  SESS
TIME    SRVCLASS P ACTIVE CUR RES CT  SC  CPU  MSO  IOC  SRB  TOTAL
15:44:27 TSOPRD  1  00:00 00:00:00 ---  0    0    0    0    0    0  -----
15:44:29 TSOPRD  1  00:00 00:00:00 27  0    0    0    0    0    0 606809
15:44:31 TSOPRD  2  00:00 00:00:00 28  0   267  1296  365    0 617346
15:44:32 TSOPRD  3  00:00 00:00:01 28  0   422  2165  745    0 618750
15:44:36 TSOPRD  4  00:00 00:00:06 28  0  6491 48528 2255    0 672692
15:44:49 TSOPRD  1  00:00 00:00:00 ---  0    0    0    0    0    0  -----
15:44:55 TSOPRD  1  00:00 00:00:00 ---  0    0    0    0    0    0  -----
15:45:05 TSOPRD  1  00:00 00:00:00 ---  0    0    0    0    0    0  -----
15:45:18 TSOPRD  3  00:00 00:00:00 34  0   551  2215  130    0 767129
15:45:22 TSOPRD  4  00:00 00:00:03 35  0  2281 15853 1780    0 792687
15:45:26 TSOPRD  4  00:00 00:00:06 35  0  6998 52654 2275    0 834700
15:45:31 TSOPRD  1  00:00 00:00:00 ---  0    0    0    0    0    0  -----
15:45:32 TSOPRD  1  00:00 00:00:00 ---  0    0    0    0    0    0  -----
15:45:33 TSOPRD  1  00:00 00:00:00 ---  0    0    0    0    0    0  -----

```

Figure 160. ASRMJ Report

Field descriptions

Table 118. Fields in the ASRM and ASRMJ Report		
Field Heading	Meaning	
hh:mm:ss JOBNAME	The time the ASRM report was requested, in the form hh:mm:ss, and the one to eight character jobname associated with the address space.	
JOBNAME TIME	The job name for the ASRMJ report. This column contains the time the report was requested.	
SRVCLASS	The service class name.	
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 259 for a description.

CHANNEL - Channel Path Activity report

In general, the CHANNEL report gives you information about channel path activity for all channel paths in the system. The report contains data for every channel path that is online at the time you request the report.

Information about channel path activity, I/O device activity, and I/O request queuing information can be used to identify performance bottlenecks associated with the channel paths.

For all channels that are managed by **Dynamic Channel Path Management** (DCM), additional information is available. DCM allows an installation to identify channels that they wish to be managed dynamically. These channels are not assigned permanently to a specific control unit, but belong to a pool of channels. Based on workload requirements in the system, these channels are 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.

Important: The Monitor II report does not support detailed reporting for certain Channel Measurement Groups (CMGs). As a result, detailed data is not available for channel paths that belong to unsupported measurement groups. However, you can access detailed data for these channels through the Monitor III report.

How to request this report

- In ISPF, specify **1** on the I/O Report Selection Menu.
- In TSO/E, use **PF4** to select the CHANNEL report.
- Command interface:

Display session

```
CHANNEL
```

Background session

```
CHANNEL
```

Special considerations of report output

You can obtain the report whether or not a Monitor I session measuring channel path activity is active. However, the channel path type appears only when RMF is active.

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

Contents of the report

RMF - CHANNEL Channel Path Activity												Line 1 of 69			
Command ===>												Scroll ===> HALF			
CPU= 37/ 35 UIC=2540 PR= 0												System= CB88 Total			
08:01:56		Channel	Utilization(%)				Read(B/s)		Write(B/s)		FICON	OPS	zHPF	OPS	
ID	No	G	Type	S	Part	Tot	Bus	Part	Tot	Part	Tot	Rate	Actv	Rate	Actv
	4		*CNCSM		0.1	0.5									
	4		*FC_SM		0.0	0.0	0.0	0	0	0	0				
12			OSD	Y	0.0	0.0	0.0	2K	19K	0	0				
14			OSD	Y	0.0	0.0	0.0	5K	478K	458K	461K				
16			OSD	Y	0.4	1.3	0.0	493K	5M	3M	5M				
20			CTC_S	Y	0.0	0.0									
27			CNC_S	Y	0.0	0.0									
2B			CNC_S	Y	0.9	3.7									
2C			CNC_S	Y	0.2	0.6									
30		5	FC_S	Y	0.0	31.8	8.6	301	50M	71	227K	177	1	0	0
31		5	FC_S	Y	0.0	31.6	7.9	239	46M	220	235K	177	1	0	0
37		4	FC_S	Y	0.0	0.1	0.0	0	85K	0	26K	6	1	0	0
38		4	FC_S	Y	0.0	0.1	0.0	0	69K	0	36K	7	1	0	0
39		4	FC_S	Y	0.0	0.0	0.0	390	11K	0	15K	2	1	0	0
3A		4	FC_S	Y	0.0	0.0	0.0	0	13K	0	16K	3	1	0	0
3E		4	FC_S	Y	0.0	0.0	0.0	0	1K	0	568	0	1	0	0
7C			CNCSM	Y	0.4	1.8									
7D			CNCSM	Y	0.0	0.1									
81		3	FC_S	Y	2.3	20.2	5.3	3M	30M	671K	2M	998	3	214	1
82		5	FC_S	Y	0.1	0.9	0.3	147K	2M	282	89K	30	1	48	1
83		5	FC_S	Y	0.1	0.9	0.3	162K	2M	291	86K	30	1	48	1
84		4	FC_S	Y	0.0	0.0	0.0	56	223	0	0	1	1	0	0
85		3	FC_S	Y	0.3	13.4	1.5	45K	4M	43K	2M	842	1	356	1
8C		3	FC_S	Y	0.9	10.9	1.6	585K	7M	80K	1M	718	2	0	0
A6		5	FC_SM	Y	0.0	0.0	0.0	0	0	0	0	0	0	0	0
B6		5	FC_SM	Y	0.0	0.0	0.0	0	0	0	0	0	0	0	0
E0			IQD	Y						0	346K				
E1			IQD	Y						0	0				
E2			IQD	Y						0	0				
E3			IQD	Y						0	0				

Figure 161. CHANNEL Report

Field descriptions

Table 119. Fields in the CHANNEL Report	
Field Heading	Meaning
Channel ID	Hexadecimal channel path identifier (CHPID).
Channel No	For each channel type which is managed by DCM, a summary line is shown with the average values for all channels in this group. These summary lines are characterized by an * preceding the channel path type, and the number of channels of the group is displayed in column No.
Channel G	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 4 indicates a throughput of 2Gbit/sec on a FICON Express4 card or a FICON Express2 card.
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 119. Fields in the CHANNEL Report (continued)

Field Heading	Meaning
Note: 1. On a machine running in LPAR mode, but with only one LPAR defined, the <i>Part</i> columns for the <i>Read</i>, <i>Write</i>, and <i>Utilization</i> fields display a zero value for channels of type FC (FICON). 2. When Channel Path Measurement Facility (CPMF) is not available, for example, on z/OS systems running as a z/VM or alternate VM guest, RMF uses sampled data from SRM so that the reported channel utilization is only an approximate value. With increasing channel speed, the channel utilization value becomes more and more inaccurate. Therefore, in such cases, RMF does not provide accurate values of FICON channel utilization. Beginning with z990 processors, the channel data from SRM is no longer available. As a result, the channel utilization data on a z/OS system running as z/VM or alternate VM guest, is reported as '-----'.	
Utilization (%) Part	The channel path utilization percentage for an individual logical partition. RMF uses the values provided by CPMF. The calculation is: $\text{Part Utilization (\%)} = \frac{\text{Channel Path Busy Time}}{\text{Channel Path Elapsed Time}} * 100$ For channels like FICON, OSA Express, or OSA Direct Express, which are running in extended CPMF mode, the calculation is as follows: $\text{Part Utilization (\%)} = \frac{\text{LPAR \# of Channel Work Units}}{\text{Max \# of Channel Work Units} * \text{Channel Path Elapsed Time}} * 100$ For OSAEGbE, the value reflects the microprocessor utilization. For hipersockets, this value is not available.
Utilization (%) Tot	The channel path utilization percentage for the CPC during an interval. For processors earlier than z990 and shared channels in LPAR mode, where CPMF is not available, the calculation is: $\text{Total Utilization (\%)} = \frac{\text{\# SRM Observations of Channel Path Busy}}{\text{\# Samples}} * 100$ For unshared channels in LPAR mode, the value for total utilization is the same as partition utilization. For channels like FICON, OSA Express, or OSA Direct Express, which are running in extended CPMF mode, the calculation is as follows: $\text{Total Utilization (\%)} = \frac{\text{Total \# of Channel Work Units}}{\text{Max \# of Channel Work Units} * \text{Channel Path Elapsed Time}} * 100$ For OSAEGbE, the value reflects the microprocessor utilization. For hipersockets, this value is not available.
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. For hipersockets, this value is not available.
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 CPC. For hipersockets, this value is not available.

Table 119. Fields in the CHANNEL Report (continued)	
Field Heading	Meaning
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 CPC.
FICON OPS	Rate Number of native FICON operations per second. Actv The average number of native FICON operations that are concurrently active during the reporting interval.
zHPF OPS	Rate Number of zHPF (High Performance FICON) operations per second. Actv The average number of zHPF operations that are concurrently active during the reporting interval.

DEV/DEVV - Device Activity report

The Device Activity report gives information on I/O device use for all online devices you requested either by device class, by device number, or by volume serial number.

The Device Activity report, like the Monitor I session report, can help you to analyze device performance, to identify bottlenecks caused by a particular device, and to overcome obstacles that prevent efficient use of the resource.

Requesting the report during a display session enables you, for example, to track the device use on a real-time basis. You can get a timely picture of device use or track a specific critical device on a real-time basis, thus making it possible to take corrective action immediately.

To evaluate the data, you need to understand what a reporting period is and how it relates to the Monitor I interval.

Evaluating details of cumulative mode output

The I= field in the header of each report shows the percentage of the Monitor I interval that has elapsed when RMF generates the Monitor II session report.

$$I = \frac{\# \text{ Samples Taken} \times \text{Cycle Time}}{\text{Monitor I Interval Length}}$$

For a report that reflects the total device activity (delta mode is off), the reporting period is the time that has elapsed from the start of the Monitor I interval to the time when you requested the report. The maximum reporting period is one Monitor I interval. When a Monitor II report covers a complete Monitor I interval, the I= field contains an upper-case 'T' (I=T).

Evaluating details of delta mode output

The I= field in the header of the report equals the percentage of the interval that is represented by the data; thus, for your initial request, the I= field equals the percentage of the interval that expires between your initial request and the time you press the ENTER key. For all subsequent requests, the I= field equals the percentage of the interval that expires.

How to request this report

Different methods are used to request the DEV and DEVV reports.

How to request a DEV report

- In ISPF, specify **3** on the I/O Report Selection Menu.
- In TSO/E, use **PF6** to select the DEV report.
- Command interface:

Display session

```
DEV [type ]
```

Background session

```
DEV [(type)]
```

How to request a DEVV report

- In ISPF, specify **4** on the I/O Report Selection Menu.
- Command interface:

Display session

```
DEVV {VOLSER(volid)      }  
      {NUMBER(device-number) }
```

Background session

```
DEVV {(VOLSER(volid))    }  
      {(NUMBER(device-number)) }
```

Special considerations of report output

The report is based on both hardware measurements and data collected during a Monitor I session. Therefore, **a Monitor I session must be active** when you issue your request.

If no data is available, RMF issues a descriptive message.

Because the data comes from both hardware measurements and Monitor I session measurements, the data required to report some or all of the fields might be invalid or unavailable. A field based on data that is unavailable or invalid contains dashes (---).

The fields that RMF might not be able to report and the possible reasons for the unavailable or invalid data are:

Field Heading	Not reported when
ACTV RATE RESP TIME IOSQ TIME DB DELAY PEND TIME DISC TIME CONN TIME %DEV UTIL	The hardware measurements are not available 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.

Field Heading	Not reported when
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 377 for an explanation of this condition.

Contents of the report

Command ==>															RMF - DEV Device Activity										Line 1 of 1513 Scroll ==> CSR									
															CPU= 8/ 6 UIC=1190 PR= 0										System= SYSF Total									
14:42:48	I=85%	DEV													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																			
		SYC337	0C337	1.0H	0094	0.000	.000	.000	.00	.00	.000	.000	.000	0	0																			
		SYC338	0C338	1.0H	0094	0.000	.000	.000	.00	.00	.000	.000	.000	0	0																			
		SYC339	0C339	1.0H	0094	0.000	.000	.000	.00	.00	.000	.000	.000	0	0																			
		SYC33A	0C33A	1.0H	0094	0.000	.000	.000	.00	.00	.000	.000	.000	0	0																			
		SYC33B	0C33B	1.0H	0094	0.000	.000	.000	.00	.00	.000	.000	.000	0	0																			
		SYC33C	0C33C	1.0H	0094	0.000	.000	.000	.00	.00	.000	.000	.000	0	0																			
DB2		USC401	0C401	1.0H	0095	0.001	.256	.000	.00	.00	.128	.000	.128	0	0																			
DB2		USC402	0C402	1.0H	0095	0.001	.768	.000	.26	.00	.384	.000	.384	0	0																			
DB2		USC403	0C403	1.0H	0095	0.001	.128	.000	.00	.00	.128	.000	.000	0	0																			
DB2		USC404	0C404	1.0H	0095	0.033	1.40	.000	.00	.00	.201	.846	.349	0	0																			
DB2		USC405	0C405	1.0H	0095	0.033	1.33	.000	.00	.00	.177	.881	.275	0	0																			
DB2		USC406	0C406	1.0H	0095	0.033	1.24	.000	.00	.00	.211	.654	.369	0	0																			
DB2		USC407	0C407	1.0H	0095	0.001	.384	.000	.00	.00	.128	.000	.256	0	0																			
DB2		USC408	0C408	1.0H	0095	0.033	1.41	.000	.00	.00	.206	.881	.320	0	0																			
DB2		USC409	0C409	1.0H	0095	0.001	.256	.000	.00	.00	.128	.000	.128	0	0																			
DB2		USC40A	0C40A	1.0H	0095	0.033	1.47	.000	.00	.00	.216	.876	.374	0	0																			

Figure 162. DEV Report

By default, the DEV report is sorted by LCU, unless you specify the storage group (SG) option. The SG option causes the DEV report to be sorted by device numbers within storage groups.

Type can be either a device class, or one or more volume serial numbers, device numbers, or storage group numbers.

When you request the report during a display session, the data line for any device that is more than 30% utilized is highlighted.

COMMAND ==>															RMF - DEVV DEVICE ACTIVITY										LINE 1 OF 11				
																									SCROLL ==> PAGE				
															CPU= 1/ 1 UIC= 65K PR= 0										SYSTEM= TRX2 TOTAL				
															I=55% DEV										ACTV RESP IOSQ -DELAY- PEND DISC CONN %D %D				
															VOLSER NUM										PAV LCU RATE TIME TIME CMR DB TIME TIME TIME UT RV				
15:26:37															USC401	0C401	1	0095	0.038	.592	.000	.00	.00	.156	.000	.436	0	0	
15:26:39															USC401	0C401	1	0095	0.038	.592	.000	.00	.00	.156	.000	.436	0	0	
15:26:49															USC401	0C401	1	0095	0.038	.592	.000	.00	.00	.156	.000	.436	0	0	
15:27:02															USC401	0C401	1	0095	0.040	.591	.000	.00	.00	.154	.000	.436	0	0	
15:28:32															USC401	0C401	1	0095	0.038	.594	.000	.00	.00	.152	.000	.441	0	0	
15:30:58															USC401	0C401	1	0095	0.120	.530	.000	.00	.00	.164	.000	.365	0	0	
15:33:49															USC401	0C401	1	0095	0.048	.570	.000	.00	.00	.162	.000	.407	0	0	
15:34:32															USC401	0C401	1	0095	0.047	.580	.000	.00	.00	.167	.000	.413	0	0	
15:36:58															USC401	0C401	1	0095	0.045	.586	.000	.00	.00	.161	.000	.424	0	0	
15:37:48															USC401	0C401	1	0095	0.040	.586	.000	.00	.00	.161	.000	.424	0	0	
15:38:18															USC401	0C401	1	0095	0.042	.597	.000	.00	.00	.158	.000	.438	0	0	

Figure 163. 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 120. 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. $I = \frac{\# \text{ Samples Taken} * \text{Cycle Time}}{\text{Monitor I Interval Length}}$ 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 “Evaluating details of cumulative mode output” on page 270 and “Evaluating details of delta mode output” on page 270 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 five-digit hexadecimal device number that identifies a physical I/O device. The first digit represents the ID of the subchannel set to which the I/O device is physically configured.
PAV	The number of parallel access volumes (base and alias) which were available at the end of the reporting interval. If the number has changed during the reporting interval, it is followed by an '*'. If the device is a HyperPAV base device, the number is followed by an 'H', for example, 5.4H. The value is the average number of HyperPAV volumes (base and alias) in that interval. $\text{Average \# of HPAV devices} = \frac{\text{Accumulated \# of HPAV devices}}{\text{Number of Samples}}$
LCU	The hexadecimal identifier of the logical control unit (LCU) to which the device belongs. The set of devices associated with an LCU measurement are not the same on all processors because the definition of an LCU is model-dependent. 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 z/VM or alternate VM guest system environment.
ACTV RATE	The rate at which start subchannel (SSCH) instructions to the device completed successfully. $\text{ACTV RATE} = \frac{\# \text{ Successful SSCH Instructions}}{\text{Interval}}$
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. $\text{ACT TIME} = \frac{\text{Total Device Active Time}}{\text{Measurement Event Count}}$ $\text{RESP TIME} = \text{ACT TIME} + \text{IOSQ TIME}$ The active time is the sum of connect, disconnect, and pending time as described later.

Table 120. 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. $\text{IOSQ TIME} = \frac{\text{Total IOSQ Time}}{\text{Start Subchannel Count}}$
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. $\text{DELAY CMR} = \frac{\text{Initial Command Response Time}}{\text{Measurement Event Count}}$
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, this value is updated every 20 seconds. $\text{DELAY DB} = \frac{\text{Device Busy Delay Time}}{\text{Measurement Event Count}}$ 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. $\text{PEND TIME} = \frac{\text{Device Pending Time}}{\text{Measurement Event Count}}$
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. $\text{DISC TIME} = \frac{\text{Device Disconnect Time}}{\text{Measurement Event Count}}$ 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. $\text{CONN TIME} = \frac{\text{Device Connect Time}}{\text{Measurement Event Count}}$

Table 120. 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. $\% D UT = \left(\frac{CONN + DISC}{INT * PAV} + \frac{RSV}{N} \right) * 100$ 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. $\% D RV = \frac{\# \text{ Device-reserved Samples}}{\# \text{ Monitor I Samples}}$

Report options

RMF Monitor II - Device Activity Options

Command ==>

Specify one of the options below. For Volume, Device Number, and Storage Group, a single name, a list of names, and a range of names is valid.
To exit press END.

Device Class ==> DASD_

Specify one of the following classes:
DASD, TAPE, COMM, CHRDR, UNITR or GRAPH

Volume ==> _____

Ex: P500002:P50004,P50007

Device Number ==> _____

Ex: 0380:038F,10580

Storage Group ==> _____

Ex: MANF13:MANF20
Specify SG to display all
Storage Groups.

Figure 164. 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 is 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. A device number is a hexadecimal number [s]dddd where s represents an optional 1-digit subchannel set ID. If the subchannel set ID is omitted, data reporting for devices dddd configured to any subchannel set is requested.

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 165. DEVV Report Options Panel

Volume

Allows you to specify a volume serial number of a specific DASD or tape device.

The name is restricted to the characters A-Z, 0-9, @, # and \$.

Device Number

Allows you to specify a hexadecimal number [s]dddd where s represents an optional 1-digit subchannel set ID. If the subchannel set ID is omitted, data reporting for device dddd configured to subchannel set 0 is requested.

HFS - Hierarchical File System Statistics report

The HFS report provides data for capacity planning and for basic performance analysis and problem determination:

- A general understanding of the throughput recognized and achieved by HFS allows you to optimally use your resources.
- The ability to display performance statistics of HFS enables you to identify potential problems and bottlenecks within the HFS component and to take corrective actions.

How to request this report

- In ISPF, specify **5** on the I/O Report Selection Menu.
- Command interface:

Display session

```
HFS [hfsname]
```

Contents of the report

```

RMF - HFS File System Statistics                               Line 1 of 12
Command ==>                                                    Scroll ==> PAGE

CPU= 37/ 35 UIC=2540 PR=   0                               System= SYSA Delta

File System Name: OMVS.SYS1.ROOT
Mount Date: 09/30/2021   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
Sequential         1234    12.340
Random             123     1.230
  
```

Figure 166. HFS Report

Field descriptions

Table 121. Fields in the HFS Report	
Field Heading	Meaning
General Section	
File System Name	The name of the file system which has been selected for reporting.
Mount Date	Date when the file system has been mounted (mm/dd/yyyy).
Time	Time when the file system has been mounted (hh:mm:ss).
Elapsed Time	Delta mode: Time between two consecutive clickings on Enter. Total mode: Time since the file system has been mounted. The format can be in days and hours (6d 19h) or in hh:mm:ss.
Allocation - All values in megabytes	
System	Amount of storage allocated to this HFS.
Data	Amount of storage internally used within HFS for data files, directories and HFS internal structures like the attribute directory (AD).
Attr. Dir	Amount of storage used for the attribute directory (AD). This number is included in the "Data" field. The attribute directory is the internal HFS structure (index) which contains attribute information about individual file system objects as well as attributes of the file system itself.
Cached	Amount of data buffer storage cached by this file system.
Index Events	
New Level	Number how often HFS added a new level to its index structure. The index statistics are relative to all of the indices in the HFS data set. The attribute directory (AD) is one index (the largest) but each directory (including the root) is also an index.
Splits	Number how often an index page was split into two pages because new records were inserted. This gives an idea of how much insertion activity there has been for the index structure.
Joins	Number how often HFS was able to combine two index pages into one, because enough index records had been deleted in the two pages.
File I/O - all values are reported as 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.

Table 121. Fields in the HFS Report (continued)	
Field Heading	Meaning
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.
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 167. HFS Report Options Panel

Table 122. 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".

Table 122. Fields in the HFS Report Options Panel (continued)	
Field Heading	Meaning
HFS File System Name	The name of a file system which was found active. The file system names are sorted in alphabetical order.

ILOCK - IRLM Long Lock Detection report

Services of the IMS/VS Resource Lock Manager (IRLM) are used by IMS to serialize application program requests for data base records to ensure that two programs do not access the same record for update at the same time.

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

Excessive use of a resource on one instance can suspend the work on other systems. To avoid such locking situations or, in the worst case, a re-IPL, the report provides information for IMS operators to perform the necessary actions to eliminate the problem.

How to request this report

- In ISPF, specify **9** on the Resource Report Selection menu.
- Command interface:

Display session

```
ILOCK [ALL]
```

Special considerations

There is no data gathering component for this report. Instead, the retrieval of the IRLM data from the RMF SMF data buffer is done by the reporter. To have the data available in the SMF data buffer (SMF record type 79 subtype 15), it is necessary to specify this option explicitly, for example:

```
S RMF,,, (SMFBUF(RECTYPE(70:78,79(15))))
```

For details, see *z/OS Resource Measurement Facility 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 irldmid,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. See "Specifying access definitions" in *z/OS Resource Measurement Facility User's Guide* for details.

Contents of the report

```

Command ==>          RMF - ILOCK IRLM Long Lock Detection          Line 1 of 15
                                Scroll ==> HALF

                                CPU= 37/ 35 UIC=2540 PR=   0          System= RMF5 Total

State   Type   Lock_Name   PSB_Name   Elap_Time   CICS_ID
        IMS_ID Recovery_Token   PST#   Trx/Job   Wait_Time   DB/Area

-----
CF Structure ACOXLOCK          at 09/30/2021 13:02:10 Deadlock Cycle 00002EC7
-----
TOP      BMP      09C943CFA7800101D700000000000000 DFSSAMB1  00:06:04
BLOCKER  AC03      AC03      0000000300000000  0006  IRLMTOPZ
-----
TOP      BMP      09C3614505800101D700000000000000 DFSSAMB1  00:06:09
BLOCKER  AC01      AC01      0000000600000000  0006  IRLMTOPA
-----
WAITER   BMP      09C3614505800101D700000000000000 DFSSAMB2
AC02      AC02      0000000800000000  0007  IRLMWT A1  00:05:52  DI21PART
-----
WAITER   BMP      09C943CFA7800101D700000000000000 DFSSAMB7
AC02      AC02      0000000900000000  0008  IRLMWT Z2  00:05:42  DI21PART
-----

```

Figure 168. ILOCK Report

Field descriptions

Table 123. Fields in the ILOCK Report	
Field Heading	Meaning
CF Structure	The name of the coupling facility structure used by IRLM. The date/time field shows the time when the SMF record for the first displayed lock entry was written.
Deadlock Cycle	The hexadecimal deadlock cycle number generated by IRLM and passed to IMS. As IMS requests longlock data for two cycles, this value is used to uniquely identify when the data was gathered.
State	State distinguishes between a blocker (holder) and a waiter: 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), SYPST (fast path system service ITASK), 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.
Recovery_Token	Recovery Token - a 16-byte token used to uniquely identify a unit of work.

Table 123. Fields in the ILOCK Report (continued)	
Field Heading	Meaning
PST#	Partition Specification Table (PST) Number. As the IMS region is initially brought up, several PSTs are initialized, and each is assigned a unique ID (or PST number). The PST block is the primary block used to dispatch transactions in IMS, and the PST number is used to uniquely identify each transaction.
PSB_Name	Partition Specification Block Name - the name given to a PSB (Program Specification Block) at PSBGEN time. This block is used to define which segments a particular application can have access to.
Trx/Job	The name of the transaction in a BMP or MPP region, or the job name for all of the remaining region types.
Elap_Time	The field is available for blockers, it contains the elapsed time between the time the PST was scheduled, or the unit of work (UOW) was created, and the time the 79.15 record was created for this entry.
Wait_Time	The field is available for waiters, it contains the elapsed time between the time that IRLM processed the request for the resource and the time that the 79.15 record was created for this entry.
CICS_ID	CICS task identifier - an 8-digit ID generated by CICS and passed to IMS at the time a CICS application is scheduled.
DB/Area	A name given to a data base (IMS full-function) or an area (IMS fastpath) at DBDGEN time. It is used to uniquely identify the data base or area the lock is held on by this transaction.

IOQUEUE - I/O Queuing Activity report

The IOQUEUE report provides information, grouped by LCU (logical control unit), on the I/O configuration. The information includes contention rate, queue lengths, and percentages of time when one or more I/O components were busy. Information about the LCU is useful because the LCU is the focus of I/O configuration and path management measurements for a related group of I/O devices.

For all channels that are managed by Dynamic Channel Path Management (DCM), additional information is available. DCM allows an installation to identify channels which they wish to be managed dynamically. These channels are not assigned permanently to a specific control unit, but belong to a pool of channels. Based on workload requirements in the system, these channels are assigned dynamically by DCM. For each LCU with DCM managed channels, a summary line displays the minimum and maximum number of connected DCM managed channels, the number of defined DCM managed channels and accumulated activity data.

PAV base mode is the mode when alias devices are assigned to one PAV base device. An I/O for a PAV base device is executed using aliases assigned to that PAV base device.

HyperPAV mode is the mode when a pool of alias devices is assigned to one LCU. An I/O for a PAV base device can be executed using any alias device of that pool.

SuperPAV mode is the mode, when a pool of alias devices is assigned to one LCU and multiple LCUs are grouped into one Alias Management Group (AMG). An I/O for a PAV base device can be executed using any alias device of these multiple alias pools. The favored way is to use the alias device assigned to the same LCU (home LCU) the PAV base device is assigned to.

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

Using the information given in the report

This report can tell you about the cause of performance problems associated with channel paths and devices. You could, for example, find the reason for an unusually long pending time reported on the device report. Check the relationship between the percentage of requests deferred for device busy and control unit busy for the LCU on the IOQUEUE report.

To help you determine the best way to fix a performance problem related to an LCU, you can request the report during a display session. This tracks the I/O queuing on a real-time basis.

Evaluating details of cumulative mode output

The I= field in the heading of each report shows the percentage of the Monitor I interval that has elapsed when RMF generates the Monitor II session report.

$$I = \frac{\# \text{ Samples Taken} * \text{Cycle Time}}{\text{Monitor I Interval Length}} * 100$$

For a report that reflects the total device activity (DELTA mode off), the reporting period is the time that has elapsed from the start of the Monitor I interval to the time when you requested the report. The maximum reporting period is one Monitor I interval. When a Monitor II report covers a complete Monitor I interval, the I= field contains an upper-case 'T' (I=T).

How to request this report

- In ISPF, specify **2** on the I/O Report Selection Menu.
- Command interface:

Display session

```
IOQUEUE [type]
```

Background session

```
IOQUEUE [(type)]
```

Special considerations of report output

The report depends on data that the Monitor I session collects. To get this report, the Monitor I I/O Queuing Activity report must be active. The Monitor I gatherer gets a new set of model dependent data every second or every cycle, whichever time period is greater.

If the hardware measurements are not available, the channel measurement facility is not available. If there is a failure in the diagnose interface, RMF does not provide model-dependent data generated by the hardware for the following fields:

- CONTENTION RATE
- DELAY Q LENGH
- CHPID TAKEN
- %CU BUSY

If the data is not reliable (indicated by a successive invalid sample count greater than zero), RMF does not provide model-dependent data generated by the hardware for the following fields:

- ACTIV RATE
- AVG Q LENGH
- %REQ DEFER

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

When an LCU has no activity during the interval, it is omitted from the report. If channel paths were brought online or taken offline during the interval, data is still formatted, but only for the channel paths and control units that were online and had some connection to a device or set of devices of the LCU at the time the report was requested appear in the report.

In a z/VM or alternate VM guest system environment, the report for an z/OS system that is authorized via the VM RMCHINFO directory option, shows static configuration data. Measurement data is not available.

Contents of the report

RMF - IOQUEUE I/O Queuing Activity										Line 92 of 684			
Command ===>										Scroll ===> CSR			
CPU= 3/ 3 UIC=2540 PR= 0										System= S5C Total			
03:20:32	I= 11%	DCM Group		Cont	Del	Q	AVG	CHPID	%DP	%CU	AVG	AVG	
Path	DCM	CTL Units	MN MX DEF	Rate	Length		CSS	Taken	Busy	Busy	CUB	CMR	
D6	5F00			0048				0.39	0.0	0.0	0.0	0.2	
				0048	0.0	0.00	0.4	0.85	0.0	0.0	0.0	0.2	
B0 PF	8000			0069				84.24	0.0	0.0	0.0	0.2	
B1 PF	8000			0069				82.79	0.0	0.0	0.0	0.2	
B2 NP	8000			0069				0.00	0.0	0.0	---	---	
B3 NP	8000			0069				0.00	0.0	0.0	---	---	
95 PF	8000			0069				81.48	0.0	0.0	0.0	0.2	
				0069	0.0	0.00	0.4	248.52	0.0	0.0	0.0	0.2	
B0 NP	8100			006A				0.00	0.0	0.0	---	---	
B1 NP	8100			006A				0.00	0.0	0.0	---	---	
B2 PF	8100			006A				122.94	0.0	0.0	0.0	0.2	
B3 PF	8100			006A				122.42	0.0	0.0	0.0	0.2	
95 NP	8100			006A				0.00	0.0	0.0	---	---	

Figure 169. IOQUEUE Report

Field descriptions

Table 124. 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. If applicable, the path attribute is indicated with the CHPID: PF preferred path NP non-preferred path NS path attribute not specified For devices residing in control units that do not support path attributes, only the CHPID is displayed.
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.
LCU	The hexadecimal identifier of 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.

Table 124. Fields in the IOQUEUE Report (continued)

Field Heading	Meaning
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. $\text{Cont Rate} = \frac{\# \text{ Enqueued Requests}}{\text{Monitor I Interval Time}}$
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. $\text{Del Q Lngth} = \frac{\text{Accumulated Queue Length} - \# \text{ Enqueued Requests}}{\# \text{ Enqueued Requests}}$
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. $\text{AVG CSS} = \frac{\text{Channel Subsystem Time}}{\# \text{ I/O Operations Accepted}}$
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. $\text{CHPID Taken} = \frac{\# \text{ I/O Operations on that Path}}{\text{Monitor I Interval Time}}$
% 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. $\% \text{ DP Busy} = \frac{\text{DPB}}{\text{DPB} + \text{CUB} + \text{SUC}} * 100$ 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
% 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. $\% \text{ CU Busy} = \frac{\text{CUB}}{\text{DPB} + \text{CUB} + \text{SUC}} * 100$ 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 124. 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. $\text{AVG CUB} = \frac{\text{Control Unit Busy Time}}{\# \text{ I/O Operations Accepted on that Path}}$
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). $\text{AVG CMR} = \frac{\text{Initial Command Response Time}}{\# \text{ I/O Operations Accepted on that Path}}$

Report options

```

RMF Monitor II - I/O Queuing Activity Options

Command ==>

Specify one of the options below. For LCU number, a single number,
a list of numbers, and a range of numbers is valid. To exit press END.

Device Class ==> DASD_   Specify one of the following classes:
                        DASD, TAPE, COMM, CHRDR, UNITR OR GRAPH

LCU Number   ==> _____ Ex: D:F,4E,55

```

Figure 170. IOQUEUE Report Options Panel

You can specify either a class or a device number.

Device Class

Allows you to specify the device class. If you leave this field empty, RMF uses DASD.

LCU Number

Allows you to request specific logical control unit numbers. The numbers must be in three-digit hexadecimal format. You can specify any combination of a single number, a list of numbers, or a range of numbers. Your entry must not exceed 32 characters, including commas and colons.

LLI - Library List report

The information shown in the LLI report provides the status of the key system libraries that are defined in the following lists:

- Load module link list
- Pageable link pack area list
- List of authorized libraries (APF list)

This information can help you to check whether the status of these libraries is correct for your current environment.

How to request this report

- In ISPF, specify **L** on the Monitor II Primary Menu. This leads you to the Library List and OPT Settings Selection Menu. Here you can select:

1 Link list

LNKLSTxx - Link Library List

2 LPA list

LPALSTxx - LPA Library List

3 APF list

IEAAPFxx - Authorized Program Library List

- In the command interface of an ISPF or TSO/E display session, specify:

LLI

for the Link Library List

LLI LPA

for the LPA Library List

LLI APF

for the Authorized Program List

In addition, you can specify optional parameter A to create a report with more details. [Table 125 on page 287](#) lists the values that are available with this parameter only. Due to performance reasons, the default for each command is just to collect basic information.

Contents of the report - Link Library List

Command ===>		RMF - LLI Program Library Information				Line 1 of 38	
						Scroll ==> HALF	
		CPU= 37/ 35 UIC=2540 PR= 0				System= RMF8 Total	
		----- Link Library List <IPL				> -----	
DevNum	DevType	Volser	APF	Ext	Data Set Name		
0020F	33903	630D14	Y	3	SYS1.LINKLIB		
0020F	33903	630D14	Y	2	SYS1.MIGLIB		
0020F	33903	630D14	Y	3	SYS1.CSSLIB		
00975	33903	SMSRMF	Y	1	RMF530.GRSREP.LINKLIB		
0020F	33903	630D14	Y	1	SYS1.SHASLINK		
0020F	33903	630D14	Y	1	SYS1.SHASMIG		
0020E	33903	MVSTGT	Y	1	SYS1.RMF.V630.D05.LINKLIB		
0031A	9345-2	RMFUSR	Y	1	DRIVER.SYS1.LINKLIB		

Figure 171. 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

```

Command ==>          RMF - LLI Program Library Information          Line 1 of 7
                                                                Scroll ==> HALF

                      CPU= 37/ 35 UIC=2540 PR=   0          System= RMF8 Total

----- LPA Library List -----
DevNum  DevType  Volser  Data Set Name
0020F   33903   630D14  SYS1.LPALIB
00975   33903   SMSRMF  RMF530.GRSREP.LPALIB
0020E   33903   MVSTGT  SYS1.RMF.V630.D05.LPALIB
0020F   33903   630D14  SYS1.ISAMLPA
0020F   33903   630D14  SDSF.ISFLPA
0020E   33903   MVSTGT  SYS1.TCP.SEZALPA
0020E   33903   MVSTGT  SYS1.REXX.V130.SEAGLPA

```

Figure 172. LLI Report - LPA Library List

The pageable link pack area list is determined by Parmlib members LPALSTxx. It is a fixed-size list that 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

```

Command ==>          RMF - LLI Program Library Information          Line 1 of 47
                                                                Scroll ==> HALF

                      CPU= 37/ 35 UIC=2540 PR=   0          System= RMF8 Total

----- APF List - Format STATIC -----
DevNum  DevType  Volser  SMS RACF  Data Set Name
0020F   33903   630D14  N  N      SYS1.LINKLIB
0020F   33903   630D14  N  N      SYS1.SVCLIB
0020F   33903   630D14  N  N      SYS1.LINKLIB
0020F   33903   630D14  N  N      SYS1.ISAMLPA
0020F   33903   630D14  N  N      SYS1.VTAMLIB
0020F   33903   630D14  N  Sec=?  SYS1.JES3LIB
0020F   33903   630D14  N  N      SYS1.NFSLIB
0020F   33903   630D14  N  N      SYS1.SCEERUN
0020F   33903   630D14  N  N      SYS1.SEPWMOD1
0020F   33903   630D14  N  N      SYS1.SEPWMOD2

```

Figure 173. 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 125. 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.

Table 125. Fields in the LLI Report (continued)	
Field Heading	Meaning
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.

OPT - OPT Settings report

The OPT Settings report displays information about the currently active OPT member in the PARMLIB and the current settings of OPT parameters.

How to request this report

- In ISPF, specify **L** on the Monitor II Primary Menu. This leads you to the Library List and OPT Settings Selection Menu. Here you can select:

4

IEAOPTxx - OPT Settings

- In the command interface of an ISPF or TSO/E display session, type the command OPT.

Contents of the report

```

RMF - OPT Settings
Line 1 of 39
Command ==> Scroll ==> CSR

CPU= 3/ 2 UIC=1027 PR= 0 System= T2 Total

OPT: 00 Time: N/A
-- Parameter -- - Default - -- Value -- Unit ----- Description -----
ABNORMALTERM Yes Yes Y/N Abnormal terminations in routing
ABSMSUCAPPING No No Y/N Absolute, permanent MSU capping
BLWLINTHD 20 20 sec Time blocked work waits for help
BLWLTRPCT 5 5 0/00 CPU cap. to promote blocked work
CCCAWMT 3200 3200 usec Alternate wait management time
CCCSIGUR 45 22 msec Min. mean-time-to-wait threshold
CNTCLIST No No Y/N Clist commands count individually
CPENABLE 10,30|0,0 10,30 % Threshold for TPI (low,high)
DVIO Yes Yes Y/N Directed VIO is active
ERV 500 500/CB SU Enqueue residency CPU Service/DP
FULLPRESYSTEM No No Y/N System AS can preempt other work
HIPERDISPATCH Yes Yes Y/N Hiperdispatch is desired/active
IFAHONORPRIORITY Yes Yes Y/N Allows CPs to help zAAPs
IIPHONORPRIORITY Yes Yes Y/N Allows CPs to help zIIPs
INITIMP 0 0/FE # INITIMP value/DP for initiators
IRA405I 70,50,50 70,50,50 % Fixed storage of <16M,16M-2G,tot
MANAGENONENCLAVE No No Y/N Manage non-enclave work
MAXPROMOTETIME 6 6 *10s Holder allowed to run promoted
MCCAFCTH 400,800 2848,5696 # Threshold for storage (low,ok)
MCCFXEPR 92 92 % Fixed storage threshold < 16 MB
MCCFXTPR 80 80 % Fixed online storage threshold
MT_CP_MODE 1 1 # MT CP mode
MT_ZIIP_MODE 1 1 # MT zIIP mode
PROJECTCPU No Yes Y/N CPU projection for zAAPs, zIIPs
RCCFXET 82,88 82,88 %/A Fixed<16 MPL threshold (low,high)
RCCFXTT 66,72 66,72 %/A Fixed MPL threshold (low,high)
RMPTTOM 3000 3000 msec SRM invocation interval
RTPIFACTOR 100 100 % PI affects server routing weights
STORAGENSWDP Yes Yes Y/N Sets non-swap. ASID non-dispatch.
STORAGESEVERMGMT No No Y/N Storage I/O priority management
STORAGEWTOR Yes YNA WTOR to cancel AS in shortage
SUPPSAFINFOMSG No No Y/N Suppress SAF informational msgs
TIMESLICES 1 1 # Time slices for discretionary wrk
VARYCPU No No Y/N VARYCPU is enabled
VARYCPUMIN 1 1 # VARYCPUMIN value
WASROUTINGLEVEL 0 0 # WebSphere routing level
WLMIRDSTRUC 4DIGITS NONE type WLM IRD structure type
ZAAPAWMT 3200 3200 usec AWM time value for zAAPs
ZIIPAWMT 3200 3200 usec AWM time value for zIIPs

```

Figure 174. OPT Settings

Field descriptions

Table 126. Fields in the OPT Settings report	
Field Heading	Meaning
OPT	Suffix xx in the name of the active option member IEAOPTxx. The option member contains parameters that affect system resource manager (SRM) decisions.
Time	Timestamp when the IEAOPTxx member was activated. If the system programmer did not change the active IEAOPTxx member of SYS1.PARMLIB since the last IPL, then 'N/A' is shown.
Parameter	Name of the WLM OPT parameter.
Default	Default value(s) of the parameter. If more than one default exists, the values are separated by a vertical bar ().

Table 126. Fields in the OPT Settings report (continued)	
Field Heading	Meaning
Value	Current value(s) of the parameter. This value may differ from the value originally specified. With two values displayed, separated by '/', the second value is provided by SRM. Also, parameters that are not set in the IEAOPTxx member are shown with the default value, if not changed otherwise. For information on how SRM handles the settings of OPT parameters, refer to the z/OS MVS Initialization and Tuning Reference . When RMF cannot obtain any data for a parameter, 'No Data' is shown.
Unit	Unit in which the parameter value is measured.
Description	Basic description of the purpose of the parameter. For detailed information, refer to the z/OS MVS Initialization and Tuning Reference .

PGSP - Page Data Set Activity report

The PGSP report provides information on page data set activity. The reporting interval is the period between any two consecutive Monitor II requests.

The PGSP report can help you to determine whether the optimum size has been allocated for each page data set.

How to request this report

- In ISPF, specify **3** on the Resource Report Selection menu.
- In TSO/E, use **PF7** to select the PGSP report.
- Command interface:

Display session

```
PGSP
```

Background session

```
PGSP
```

Special considerations of report output

A Monitor I session measuring page data set activity must be active when you request the report.

Dashes (---) in the data fields indicate that RMF could not provide a value because the page device has been varied online during the reporting interval, or a Monitor I interval ended.

Contents of the report

RMF - PGSP Page Data Set Activity								Line 1 of 4
Command ==>								Scroll ==> PAGE
CPU= 5/ 1 UIC= 65K PR= 0								System= TRX1 Total
S	VOLUME	DEV	DEV	%SLOTS	PAGE	I/O REQ	AVG	PAGES
T	SERIAL	NUM	TYPE	IN USE	TRAN	TIME	RATE	PER I/O V
								10:22:26
P	TRX1PP	0445D	33903	21.21	0.000	0.000	0.000	PAGE.VTRX1PP.PLPA
C	TRX1PP	0445D	33903	0.10	0.000	0.000	0.000	PAGE.VTRX1PP.COMMON
L	TRX1P1	0455D	33909	0.00	0.000	0.000	0.000	PAGE.VTRX1P1.LOCAL1
S	N/A	N/A	N/A	0.19	0.000	0.000	0.000	N/A

Figure 175. PGSP Report

Field descriptions

Table 127. 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 S SCM (Storage Class Memory)	
VOLUME SERIAL	The volume serial number of the volume on which the data set resides. N/A for SCM.	
DEV NUM	The device number. N/A for SCM.	
DEV TYPE	The device type. N/A for SCM.	
% SLOTS IN USE	The percentage of the slots in the page data set that are in use. When you request the report, RMF calculates the percentage from the Monitor I sampling values. $\% \text{ SLOTS IN USE} = \frac{\# \text{ Slots in the Data Set} - \# \text{ Available Slots}}{\# \text{ Slots in the Data Set}} * 100$ In case of SCM, this value is the percentage of 4K SCM blocks in-use by ASM in relation to the number of total blocks available to ASM.	
PAGE TRAN TIME	The page transfer time in seconds. When you request the report, RMF calculates the value from the current Monitor I sampling values. $\text{PAGE TRAN TIME} = ((\text{USE} * \text{INT}) / \text{N}) / \text{XFER}$ 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.	
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	

Table 127. Fields in the PGSP Report (continued)	
Field Heading	Meaning
DATA SET NAME	The name of the page data set being monitored. N/A for SCM. If a page data set name is longer than 23 characters, it is truncated in the report to 22 characters followed by an asterisk (*). If a data set has bad slots, the data set name is preceded by an asterisk (*). When the operating system detects errors in a data set that prevent its further use, RMF can no longer monitor the data set. RMF indicates that monitoring is terminated by <i>preceding</i> the data set name with two asterisks (**).

SDS - Sysplex Data Server report

The report provides statistics about the usage of the SMF Data Buffer, and the usage of the Sysplex Data Services.

The numbers presented in the first section aid in finding the optimal size of the SMF buffer to hold as many SMF records as an installation might want to keep for immediate sysplex reporting. The second section contains statistics about the exploitation of the callable services. You can use these numbers to optimize the usage of the callable services within other applications than RMF.

The SDS report requires that the RMF address space has been started. Otherwise, no statistics can be provided, neither from the reporting system, nor from any remote system in the sysplex.

How to request this report

- In ISPF, specify **8** on the Resource Report Selection menu.
- Command interface:

Display session

```
SDS
```

Contents of the report

```

RMF - SDS RMF Sysplex Data Server                               Line 1 of 13
Command ==>                                                    Scroll ==> HALF

CPU= 37/ 35 UIC=2540 PR=  0                                System= SYS1 Total

RMF Sysplex Data Server Statistics
Report Start 09/23/2021 16:49:24, End 09/30/2021 12:28:21, Duration  6d 19h

SMF Buffer Statistics, Start 09/23/2021 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 176. SDS Report

Field descriptions

Table 128. Fields in the SDS Report	
Field Heading	Meaning
RMF Sysplex Data Server Statistics	

Table 128. Fields in the SDS Report (continued)	
Field Heading	Meaning
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

The report is a "snapshot" report that describes the contention or ownership at the time RMF processes the request for the report. The SENQ report tracks contention for or ownership of serially-reusable resources. SENQ reports only the contention caused by ENQ and DEQ macro instructions.

Note: For information about the contention caused by the RESERVE macro instruction, use the reserve activity report.

Especially when invoked during a display session, the report can help you to determine, on a real-time basis, which resources and jobs are contributing to any bottlenecks caused by resource contention.

For a detailed description of the different parameters, see *z/OS Resource Measurement Facility User's Guide*.

How to request this report

- In ISPF, specify **1** on the Resource Report Selection menu.

By default, you get the summary report. Other reports can be selected from the Report Options panel.

- In TSO/E, use **PF8** to select the SENQ report.
- Command interface:

Display session

```
SENG {S
      {D
      {A,sysname
      {E,sysname
      {majorname[,minorname]}
```

Background session

```
SENG {(S
      {(D
      {(A,sysname
      {(E,sysname
      {(majorname[,minorname])}
```

Different types of SENQ reports

You can request the SENQ report as:

- Summary report for all resources that had contention
- Detail report for all resources that had contention or for a specific resource identified by name
- Report of resources held by a specific system, whether or not there is a contention

Summary report

The summary report includes all resources that had a contention. It describes the number of tasks that own each resource and the number of tasks waiting for the resource. [Figure 177 on page 295](#) 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 179 on page 295](#) 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 178 on page 295](#) shows a sample resource report for a specific system.

Contents of the report

```
Command ===>          RMF - SENQ System Enqueue Contention          Line 1 of 4
                                                                    Scroll ===> HALF
                        CPU= 37/ 35 UIC=2540 PR=   0          System= SYS1 Total

14:52:05 TSK   TSK   TSK MAJOR NAME
        OWN   WTE   WTS  MINOR NAME

           1     0     1   CLRSHARE
           1     0     1   CLRVSAM.HBB4420.ILOG7          (SYSS)
           1     0     1   SYSDSN
           1     0     1   KEYES.MYSLM.MACLIB          (SYSS)
```

Figure 177. SENQ Summary Report

```
COMMAND ===>          RMF - SENQ System Enqueue Contention          Line 1 of 14
                                                                    Scroll ===> HALF
                        CPU= 37/ 35 UIC=2540 PR=   0          System= SYS1 Total

11:44:26
JOBNAME  SYSTEM  ASID  REQ  MAJOR NAME
                        MINOR NAME

ADMPRINT M303    185  E0    ADMPRNTQ
                        BACKGROUND DRIVER EXECUTING.          (SYSS)
HSM       M303    13   S0    ARCGPA
                        ARCMCDS                                (SYSS)
D10MCW1   M303    116  S0    BLXDASDS
                        D83INFO.BLGPNLS                        (SYSS)
D32VJF1   M303    88   S0    CLR12.UN
                        IUSER                                  (SYSS)
D31BAR1   M303    217  E0    GIMSMP
                        C87JPLX.SMPCSI.CSI                    (SYSS)
D31BAR1   M303    35   E0    SPFDSN
                        D31BAR1.IAREND.P06.PLS                (SYSS)
D75DVB1   M303    41   E0    D75DVB1.IATOSDR.W1A.ASM          (SYSS)
D96JRF1   M303    163  E0    D96JRF1.IEFDB413.PS1.PLS          (SYSS)
```

Figure 178. SENQ GRS Report

```
COMMAND ===>          RMF - SENQ System Enqueue Contention          Line 1 of 4
                                                                    Scroll ===> HALF
                        CPU= 37/ 35 UIC=2540 PR=   0          System= SYS1 Total

11:48:44
JOBNAME  SYSTEM  ASID  REQ  MAJOR NAME
                        MINOR NAME
                        SYSDSN

D75TYT1   AQXI    136  E0    D75TYT1.IATCNNJ.P10.ASM          (SYSS)
D75TYT1A  AQXA    29   EW
GALER     AQXI    62   E0    GALER.PA21100.PTMPRINT          (SYSS)
GALERH    AQXA    30   SW
```

Figure 179. SENQ Detail Report

Field descriptions

Table 129. 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.	
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 <i>qname</i> field in the ENQ and DEQ macro instructions, is one to eight characters in length; it is aligned under the MAJORNAME heading. The minor name, which corresponds to the <i>rname</i> field in the ENQ and DEQ macro instructions, can be from 1 to 255 characters in length. However, only 44 characters can appear in the report. When a minor name exceeds 44 characters, it is truncated in the report, and an asterisk (*) following the scope indicates that the name has been truncated. If the minor name contains unprintable characters, RMF reports in the form 'name', where <i>name</i> appears as up to 44 hexadecimal digits. Each minor name is aligned under the MINORNAME heading. RMF recognizes only 44 characters. Therefore if two minor names (both with the same major name) are longer than 44 characters and differ only beyond the forty-fourth character, RMF cannot distinguish between them. The scope of the resource follows the minor name. A resource with a scope of 'SYSTEMS' is followed by (SYSS). A resource with a scope of 'SYSTEM' is followed by (STEP). When the major and minor names are blank on the report, the last non-blank values (above) are valid.	

Report options

```
RMF Monitor II - System Enqueue Options

Command ==>

Specify one of the options below. To exit press END.

Enqueue contention report:

  Summary      ==> YES      Specify YES for a summary or NO for
                           a detailed report.

Enqueue report by system:

  All          ==> ___      Specify YES for all owned resources or
                           or NO for exclusively owned resources only.
  System ID    ==> _____ Specify the system holding the resources.

Enqueue report by major-/minorname:

  Majorname    ==> _____
  Minorname    ==> _____
```

Figure 180. SENQ Report Options Panel

You can specify either a summary report, a report by system, or a report by major-/minorname.

Summary

Allows you to specify a summary report.

YES

Is the default value. A summary report includes the scope of the resource, the number of tasks waiting for exclusive use of the resource, and the number of tasks waiting for shared use of the resource.

NO

Causes a detailed report to be generated.

All

Allows you to specify a report that includes all resources that a system holds in a global serialization complex.

System ID

The system for which the enqueue report is requested.

To create a report that includes all resources owned by the system, also specify YES for All.

Majorname, minorname

Allows you to specify a detailed report for a specific resource that had contention. The **majorname** is a 1 to 8 character major name of a serially-reusable resource. If you specify only a major name, RMF lists all resources grouped under the major name. The optional **minorname** contains the minor name of the resource. The minor name can be 1 to 30 characters.

If you specify **S** or **D** as majorname, you need to specify a minorname.

A and **E** cannot be used as majornames.

SENQR - System Enqueue Reserve report

The report is a 'snapshot' report that describes the status of all RESERVE requests outstanding at the time when RMF processes the request for the report.

The SENQR report enables you to track RESERVE macro instructions issued to reserve a shared direct access device (shared DASD) for use by a particular system.

Figure 181 on page 298 shows a sample SENQR report requested for all devices.

How to request this report

- In ISPF, specify **2** on the Resource Report Selection menu.
By default, you get all volumes; you can select a specific volume from the Report Options panel.
- In TSO/E, use **PF9** to select the SENQR report.
- Command interface:

Display session

```
SENQR {ALLVSR}  
      {volser }
```

Background session

```
SENQR {(ALLVSR)}  
      {(volser) }
```

Contents of the report

```
RMF - SENQR System Enqueue Reserve                               Line 1 of 4  
Command ===>                                                    Scroll ===> HALF  
  
CPU= 37/ 35 UIC=2540 PR= 0                                     System= SYS1 Total  
  
14:52:57      SYSTEM ENQUEUE RESERVE REPORT  
JOBNAME ASID SYSTEM REQ VOLUME DEV  RSV MAJOR  MINOR  
  
CATALOG  34 AQTS  SO  TS0020 00AF1 CNV SYSIGGV2 CATALOG.VTS0020  
BMORRISP 70 AQTS  EO  RES84Z 00D4E ON SYSVT0C RES84Z  
JES2     31 AQTS  EO  JES2PK 00702 OFF SYSZJES2 JES2PKSYS1.BKUPCKPT  
CATALOG  34 AQTS  SO  CAT212 004B1 OFF SYSZVVD5 CAT212
```

Figure 181. SENQR Report

Field descriptions

Table 130. Fields in the SENQR Report	
Field Heading	Meaning
JOBNAME	The name of the job that issued the RESERVE macro instruction for the device identified under DEV.
SYSTEM	The identifier of the system on which the job that owns or requests the resource is running.
ASID	The address space identifier of the job that issued the RESERVE macro instruction for the device identified under DEV.
REQ	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.

Table 130. Fields in the SENQR Report (continued)	
Field Heading	Meaning
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 <i>qname</i> 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 <i>rname</i> field in the RESERVE macro instruction, can be from 1 to 255 characters in length. However, only 32 characters can appear in the report. When a minor name exceeds 32 characters, it is truncated in the report, and an asterisk (*) indicates the name has been truncated. If the minor name contains unprintable characters, RMF reports it in the form 'name', where <i>name</i> is shown in hexadecimal digits and is only 29 digits in length. Each minor name is aligned under the MINOR heading. RMF recognizes only 44 characters. Therefore if two minor names (both with the same major name) are longer than 44 characters and differ only beyond the forty-fourth character, RMF cannot distinguish between them.

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 182. SENQR Report Options Panel

Enter the reserved volume name. For a report of all reserved volumes enter ALLVSER.

SPAG - Paging Activity report

The Paging Activity report presents overview information on system paging activity.

This report enables you to see the paging activity of your system more clearly. This is due to the fact that most of the report fields reflect rates that show the activity since the last request for the report. Exceptions are the following fields: TIME, AFC (length of the available frame queue), HI UIC (highest UIC), and ESF AVL (number of available expanded storage frames).

Each report consists of one line of data that gives a "snapshot" view of system paging activity at the time the report was requested. When you invoke the report repetitively, you can build a table showing the differences over a period of time. [Figure 183 on page 300](#) shows how you can repeat requests for the report to build a table of system paging activity.

How to request this report

- In ISPF, specify **4** on the Resource Report Selection menu.
- In TSO/E, use **PF10** to select the SPAG report.
- Command interface:

Display session

SPAG

Background session

SPAG

Contents of the report

```
RMF - SPAG Paging Activity                                Line 1 of 9
Command ===>                                           Scroll ===> HALF

CPU= 37/ 35 UIC=2540 PR= 0                               System= SYS1 Total

  TIME    LPA CSA SWP  PGS-SWPD PRIV_IN PRV V&H TAR  HI  ES MIG  ESF  MIG
        IN  IN  OUT   IN   OUT BLK  NBK OUT I+O CWS  AFC  UIC  RTE AGE  AVL  RTE
14:06:45 --- --- ---  ---  ---  ---  ---  ---  ---  ---  ---  ---  ---  ---  ---
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 --- --- ---  ---  ---  ---  ---  ---  ---  ---  ---  ---  ---  ---  ---
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 183. SPAG Report

Field descriptions

Table 131. Fields in the SPAG Report	
Field Heading	Meaning
TIME	The time the report was requested.
LPA IN	The rate of LPA pages paged in. After the initial request, this field represents the rate since the previous report.
CSA IN	The rate of CSA pages paged in. After the initial requests, this field represents the rate since the previous report.
SWP OUT	The rate of successful swap-outs.
PGS-SWPD IN	The rate of pages swapped in. After the initial request, this field represents the rate since the previous report.
PGS-SWPD OUT	The rate of pages swapped out. After the initial request, this field represents the rate since the previous report.
PRIV_IN BLK	The rate of private area (VIO + non-VIO) pages paged in, in blocks, not including the first page. After the initial request, this field represents the rate since the previous report.
PRIV_IN NBK	The rate of private area (VIO + non-VIO) pages paged in. After the initial request, this field represents the rate since the previous report. This rate includes single pages plus the first page of each block.
PRV OUT	The rate of private area (VIO + non-VIO) pages paged out. After the initial request, this field represents the rate since the previous report.
V&H I+O	The rate of hiperspace and VIO pages paged in and paged out. After the initial request, this field represents the rate since the previous report.

Table 131. Fields in the SPAG Report (continued)	
Field Heading	Meaning
TAR CWS	The target working set size for the common area.
AFC	The total number of frames currently available.
HI UIC	The highest unreferenced interval count (UIC). Values greater than 9999 are displayed as nnK to indicate a multiple of 1000. The maximum value is 65K.
ES RTE	The rate of pages sent to expanded storage. After the initial request, this field represents the rate since the previous report. In a system without expanded storage, the field heading appears in the report, but the field is blank.
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

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 184 on page 302](#) is an example of repeated requests.

This report detects abnormal situations when they occur so you can request a more detailed report to further analyze the situation.

How to request this report

- In ISPF, specify **5** on the Resource Report Selection menu.
- In TSO/E, use **PF11** to select the SRCS report.
- Command interface:

Display session

```
SRCS
```

Background session

```
SRCS
```

Contents of the report




Figure 184. SRCS Report

Field descriptions

Field Heading	Meaning	
TIME	The time the report was requested.	
AFC	The average number of available frames.	
HI UIC	The highest unreferenced interval count (UIC). Values greater than 9999 are displayed as nnK to indicate a multiple of 1000. The maximum value is 65K.	
SQA F	The total number of SQA frames, including frames in central storage.	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 and restricted use common service area (RUCSA) frames.	
L+C FF	The total number of fixed LPA, CSA, and RUCSA 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.	

Table 132. Fields in the SRCS Report (continued)

Field Heading	Meaning
CPU UTL	The average processor utilization percentage for all general purpose processors (CPs) currently online. If Monitor I CPU gathering is active, it is the MVS view of CPU utilization which is the percentage of the time that the general purpose processors were busy: $\text{CPU UTL} = \frac{\text{Time Range} - \text{Sum of Wait Times}}{\text{Time Range}} * 100$ The time range is the sum of the times the processors were online. With HiperDispatch mode active, it is the sum of the times the processors were online but not parked. If Monitor I CPU gathering is not active, CPU UTL is the SRM view of CPU utilization (CCVTUTILP). Note: The CPU UTL value is a snapshot of CPU usage over a short period of time and is identical to the first value in the CPU= field in the header. For details, please refer to Table 115 on page 253.
IN Q	The current length of the SRM in queue.
OUT LOG	The current number of address spaces that are logically swapped out.
OUT RQ	The current length of the SRM out ready queue.
OUT WQ	The current length of the SRM out wait queue.

Chapter 4. Real-time reporting with Monitor I

Monitor I produces interval reports that are created at the end of a measurement interval, for example, 30 minutes.

You can obtain Monitor I session interval reports during or at the end of RMF processing, or they can be generated at a later time by the Postprocessor.

[“CHAN - Channel Path Activity report” on page 341](#)

[“CPU - CPU Activity report” on page 347](#)

[“CRYPTO - Crypto Hardware Activity report” on page 369](#)

[“DEVICE - Device Activity report” on page 377](#)

[“ENQ - Enqueue Activity report” on page 395](#)

[“FCD - FICON Director Activity report” on page 407](#)

[“IOQ - I/O Queuing Activity report” on page 417](#)

[“PAGESP - Page Data Set Activity report” on page 431](#)

[“PAGING - Paging Activity report” on page 433](#)

[“TRACE - Trace Activity report” on page 467](#)

[“VSTOR - Virtual Storage Activity report” on page 472](#)

See Chapter 5, [“Long-term overview reporting with the Postprocessor,” on page 307](#) for a description of these reports.

Chapter 5. Long-term overview reporting with the Postprocessor

Postprocessor reports are based on data gathered as SMF records by RMF (Monitor I, Monitor II, and Monitor III), by web servers, and by Lotus Domino servers.

This information unit describes the following report types:

- **Interval and Duration reports**

Reports are available as single-system and sysplex reports.

- **Exception report**

Presents a summary of values that exceeded installation-defined thresholds.

- **Overview report**

Provides an improved version of the Exception and Summary report and offers data for further processing in spreadsheet or other applications.

- **Summary report**

Presents an overview of system activity.

Postprocessor reports are available as either textual reports or XML reports, or both. You can use the XML output format of a report for further processing, for example, with an XML parser.

[How to work with Postprocessor XML reports in z/OS Resource Measurement Facility User's Guide](#) provides or navigates to all required information on how to produce and view XML reports.

Table 133 on page 307 presents an overview of available formats for all Postprocessor reports.

Table 133. Available formats for Postprocessor reports			
Report	Text	XML produced by Postprocessor job	XML available in Spreadsheet Reporter
CACHE	yes	yes	yes
CF	yes	yes	yes
CHAN	yes	yes	yes
CPU	yes	yes	yes
CRYPTO	yes	yes	yes
DEVICE	yes	yes	yes
DOMINO	yes		
EADM		yes	yes
ENQ	yes	yes	yes
ESS	yes	yes	yes
FCD	yes	yes	yes
HFS	yes	yes	yes
HTTP	yes		
IOQ	yes	yes	yes

Table 133. Available formats for Postprocessor reports (continued)

Report	Text	XML produced by Postprocessor job	XML available in Spreadsheet Reporter
OMVS	yes	yes	yes
PAGESP	yes	yes	yes
PAGING	yes	yes	yes
PCIE		yes	yes
SDEVICE	yes	yes	yes
SDELAY		yes	yes
TRACE	yes		
VSTOR	yes	yes	yes
WLMGL	yes	yes	yes
XCF	yes	yes	yes
Exception report	yes		
Overview report	yes	yes	
Summary report	yes		

Interval and duration reports

The Postprocessor can generate interval reports based on data gathered as SMF records by RMF (Monitor I, Monitor II, and Monitor III), by web servers, and by Lotus Domino servers.

Interval reports can be created either as **single-system reports** using the report option:

REPORTS(option)

or as **sysplex reports** with the report option:

SYSRPTS(option)

The Postprocessor can either get its input from data sets with SMF records from all systems in the sysplex, or it can access all current SMF records in the sysplex automatically using the RMF Sysplex Data Server.

If the SMF records are collected by the Data Gatherer at a different release or service level than the RMF Postprocessor, the SMF records are converted for use with RMF.

For details on how to call the Postprocessor with the different options and capabilities, see *z/OS Resource Measurement Facility 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 305](#).

Samples of the interval reports printed during a Monitor II session appear in [Chapter 3, “Snapshot reporting with Monitor II,” on page 251](#).

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.

<i>Table 134. Interval and Duration Reports</i>				
Report Option	Report Name	Gathered by	SMF Record	Gathering Options required
REPORTS(CACHE)	Cache subsystem activity	Monitor I	74.5	CACHE
SYSRPTS(CF)	Coupling facility activity	Monitor III	74.4	SYS(TYPE(...,74(4),...)) in active SMFPRMxx parmlib member
REPORTS(CHAN)	Channel path activity	Monitor I	73	CHAN
REPORTS(CPU)	CPU activity	Monitor I	70.1	CPU
REPORTS(CRYPTO)	Crypto hardware activity	Monitor I	70.2	CRYPTO
REPORTS(DEVICE)	Device activity	Monitor I	74.1	DEVICE(type)
REPORTS(DOMINO)	Lotus Domino server activity	Lotus Domino server	108.1, 108.3	gathererd by Lotus Domino servers
REPORTS(EADM)	EADM Activity	Monitor III	74.10	SYS(TYPE(...,74(10),...)) in active SMFPRMxx parmlib member
REPORTS(ENQ)	Enqueue activity	Monitor I	77	ENQ(SUMMARY DETAIL[,majorname[,minorname]])
REPORTS(ESS)	Enterprise Disk Systems activity	Monitor I	74.5, 74.8	ESS(options)
REPORTS(FCD)	FICON director activity	Monitor I	74.7	FCD
REPORTS(HFS)	HFS statistics	Monitor III	74.6	SYS(TYPE(...,74(6),...)) in active SMFPRMxx parmlib member and optionally Monitor III gatherer option HFSNAME(ADD(hfsname))
REPORTS(HTTP)	HTTP server activity	IBM HTTP Server (IHS) powered by Domino	103.1, 103.2	gathered by IBM HTTP Servers powered by Domino
REPORTS(IOQ)	I/O queuing activity	Monitor I	78.3	IOQ(options)
REPORTS(OMVS)	OMVS kernel activity	Monitor III	74.3	SYS(TYPE(...,74(3),...)) in active SMFPRMxx parmlib member

Table 134. Interval and Duration Reports (continued)

Report Option	Report Name	Gathered by	SMF Record	Gathering Options required
REPORTS(PAGESP)	Page data set activity	Monitor I	75	PAGESP
REPORTS(PAGING)	Paging activity	Monitor I	71	PAGING
REPORTS(PCIE)	PCIE activity	Monitor III	74.9	SYS(TYPE(...,74(9),...)) in active SMFPRMxx parmlib member
REPORTS(SDELAY)	Serialization delay	Monitor III	72.5	SYS(TYPE(...,72(5),...)) in active SMFPRMxx parmlib member
SYSRPTS(SDEVICE)	Shared device activity	Monitor I	74.1	DEVICE(DASD TAPE)
REPORTS(TRACE)	Trace activity	Monitor I	76	TRACE(variable[,options list])
REPORTS(VSTOR)	Virtual storage activity	Monitor I	78.2	VSTOR(S D [,jobname1,jobname2,...] jobname)
SYSRPTS(WLMGL)	Workload activity	Monitor I	72.3	WKLD
REPORTS(XCF)	XCF activity	Monitor III	74.2	SYS(TYPE(...,74(2),...)) in active SMFPRMxx parmlib member
The ENQ, SDELAY, and TRACE report are only available as interval reports.				

In addition, the Postprocessor can create the following interval reports based on data collected during a Monitor II background session. Duration reports are not available. You can find a description of these reports in [Chapter 3, “Snapshot reporting with Monitor II,” on page 251](#).

Table 135. Monitor II Interval Reports

Report Option	Report Name	SMF Record
REPORTS(ARD/ARDJ)	Address space resource data	79.2
REPORTS(ASD/ASDJ)	Address space state data	79.1
REPORTS(ASRM/ASRMJ)	Address space SRM data	79.5
REPORTS(CHANNEL)	Channel path activity	79.12
REPORTS(DEV/DEVV)	Device activity	79.9
REPORTS(IOQUEUE)	I/O queuing activity	79.14
REPORTS(PGSP)	Page data set activity	79.11
REPORTS(SENQ)	System enqueue contention	79.7
REPORTS(SENQR)	System enqueue reserve	79.6
REPORTS(SPAG)	Paging activity	79.4
REPORTS(SRCS)	Central storage/Processor/SRM	79.3

Single-system report header

```

                                R E P O R T   T I T L E
z/OS 3.1      SYSTEM ID  cccc      DATE mm/dd/yyyy      INTERVAL mm.ss.ttt      PAGE nnnn
              report_version    TIME hh.mm.ss          CYCLE s.ttt SECONDS
  
```

Figure 185. 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 3.1	The version of the operating system.
SYSTEM ID cccc	The SMF system ID of this system.
report_version	The version of the RMF Postprocessor. This field also indicates if SMF records were converted due to being collected by the Data Gatherer at another version or service level. The following values are possible: RPT VERSION 3.1 RMF The version of the RMF Postprocessor. No SMF records were converted. UP-CONVERTED TO 3.1 LEVEL nnn SMF records were converted from a lower release or service level to this RMF Postprocessor service level (<i>nnn</i>). DOWN-CONVERTED TO 3.1 LEVEL nnn SMF records were converted from a higher service level to this RMF Postprocessor service level (<i>nnn</i>).
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
z/OS 3.1      SYSPLEX sysplex      DATE mm/dd/yyyy      INTERVAL mmm.ss.ttt      PAGE nnnn
              report_version    TIME hh.mm.ss          CYCLE ss.ttt SECONDS
  
```

Figure 186. Header of a Sysplex report

The difference in this header compared to the single-system report is that the sysplex name is shown instead of the SMF system ID.

Duration report header

REPORT TITLE				
z/OS 3.1	SYSTEM ID cccc	START mm/dd/yyyy-hh.mm.ss	INTERVAL hhh.mm.ss	PAGE nnnn
	report_version	END mm/dd/yyyy-hh.mm.ss	CYCLE s.ttt SECONDS	

Figure 187. Header of a Single-System Duration Report

The following differences exist in the header of a duration report when compared to an interval 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.
- For the *report_version*, if there is a mix of SMF records converted and not converted, the following value appears: **SOME RECORDS CONVERTED TO3.1 RMF**

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 *z/OS Resource Measurement Facility User's Guide*. The overview condition names (also shortly called overview names or overview conditions) that refer to single fields in a report are shown in the rightmost column of the table of spreadsheet range names for the report, so that you can correlate them with a field heading.

For full details of the algorithms used to determine the overview condition, and the SMF record fields they apply to, see *z/OS Resource Measurement Facility User's Guide*.

CACHE - Cache Subsystem Activity report

The Cache Subsystem Activity report provides cache statistics on a subsystem basis as well as on a detailed device-level basis.

With the help of cache control units, access time to data that resides on a DASD can be reduced to the minimum allowed by the speed of the channels, providing significant I/O response time improvements. Important questions for performance analysts are:

- *How many control units do I need?*
- *What is the optimum size of the cache?*
- *How much non-volatile storage (NVS) do I need?*
- *Is the cache performing effectively?*

The Cache Subsystem Activity report provides answers to these questions.

How to request this report

Monitor I gathers data for this report with the default option CACHE as SMF record type 74.5. If you want to suppress gathering, you have to specify option NOCACHE.

To produce this report, specify

```
REPORTS(CACHE(options))
```

This report is also available in XML output format. [How to work with Postprocessor XML reports in z/OS Resource Measurement Facility User's Guide](#) provides all required information on how to produce and view XML reports.

Example URLs for the DDS API:

```
http://ddshost:8803/gpm/rmfpp.xml?reports=CACHE(SUBSYS)
http://ddshost:8803/gpm/rmfpp.xml?reports=CACHE(SUMMARY)
http://ddshost:8803/gpm/rmfpp.xml?reports=CACHE(DEVICE)
```

Different report levels

The contents of the report depends on the reporting options:

Cache Summary reporting - REPORTS (CACHE (SUMMARY)) (see [“Cache Summary reporting” on page 313](#))

This generates a report with three sections:

- Cache Subsystem Summary
- Top-20 Device List by DASD I/O Rate
- Top-20 Device List by total I/O Rate

The Summary report provides an overview on all subsystems with the most relevant data. The two device lists contain the 20 devices which show the highest DASD I/O rates (devices that should be investigated for potential cache hit improvements) or that show the highest I/O rates (hopefully with the best cache hit rates).

Subsystem-level reporting - REPORTS (CACHE (SUBSYS)) (see [“Subsystem-level reporting” on page 314](#))

This generates a report with these sections:

- Cache Subsystem Status
- Cache Subsystem Overview
- Cache Subsystem Device Overview
- RAID Rank Activity, if available

The subsystem-level report gives an overall view of the storage controller, that is the amount of cache storage and non-volatile storage installed. In addition, the performance analyst finds the number of I/O requests sent to the control unit and their resolution in the cache (*hits*). Furthermore, a list of all volumes attached to the subsystem is part of the report, showing their specific utilization of the cache.

The suboptions SSID/EXSSID can be used to select or exclude specific control units.

Device-level reporting - REPORTS (CACHE (DEVICE)) (see [“Device-level reporting” on page 320](#))

This generates, in addition to the report previously described, a report with two sections:

- Cache Device Status
- Cache Device Activity

The device-level report provides detailed information for each single device attached to the selected control unit. The status section shows whether DASD FAST WRITE is active. The report is intended to help analyze cache usage in detail on the basis of the information about the applications that access these volumes.

Note: When comparing I/O rates in the DASD Activity report and in the Cache Subsystem Activity report, you may see differences due to different ways how I/Os are counted:

- In the DASD Activity report, one I/O is counted for one SSCH or RSCH instruction. There can be record chaining, for example for paging I/O, which is not reflected in the SSCH count.
- In the Cache Subsystem Activity report, one I/O is counted for each cache request, and one I/O chain may cause several cache requests.

This different I/O counting can lead to higher or lower I/O rates in the Cache Subsystem Activity report than in the DASD Activity report.

Cache Summary reporting

Cache Subsystem Summary

The report offers you a top-down approach to analyze the storage subsystems in your configuration because you can see at a glance the most important data. Looking at this report, the storage subsystems causing problems can be easily identified and analyzed in a second Postprocessor run requesting more details.

C A C H E S U B S Y S T E M S U M M A R Y																			PAGE 1	
z/OS 3.1				SYSTEM ID SYS1 RPT VERSION 3.1 RMF					DATE 04/25/2023 TIME 03.14.33				INTERVAL 15.00.003							
SSID	CU-ID	TYPE	CACHE	NVS	I/O RATE	--CACHE READ	HIT DFW	RATE-- CFW	----	DASD I/O STAGE	DEL	NVS	OTHER	ASYNC RATE	TOTAL H/R	READ H/R	WRITE H/R	% READ		
2180	2182	2107-951	124G	4096	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		
4000	4042	2107-981	246G	16G	20.0	20.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.999	0.999	1.000	100.0			
4100	41A8	2107-981	246G	16G	786.6	786.5	0.1	0.0	0.0	0.0	0.0	0.0	0.0	1.000	1.000	1.000	100.0			
4300	43AD	2107-981	246G	16G	17.9	17.9	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.999	N/A	N/A	100.0			
4400	44C4	2107-981	246G	16G	0.1	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.898	0.857	1.000	71.7			

Figure 188. 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 3.1				SYSTEM ID SYS1 RPT VERSION 3.1 RMF				DATE 04/25/2023 TIME 03.14.33				INTERVAL 15.00.003					
*** DEVICE LIST BY DASD I/O RATE ***																	
VOLUME SERIAL	DEV NUM	SSID	% I/O	I/O RATE	---CACHE READ	HIT DFW	RATE-- CFW	----	DASD I/O STAGE	DEL	NVS	OTHER	ASYNC RATE	TOTAL H/R	READ H/R	WRITE H/R	% READ
SYSDS2	047C3	4700	23.3	2.3	0.4	1.8	0.0	0.1	0.0	0.0	0.0	0.0	0.2	0.972	0.876	1.000	22.3
SYSASP	047BF	4700	62.2	6.1	1.8	4.3	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.995	0.984	1.000	30.0
MVSQIK	0405B	4000	2.4	0.5	0.5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.958	0.958	N/A	100.0
MVSOIB	043AD	4300	33.7	6.0	6.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.997	0.997	N/A	100.0
SYSFPP	044C4	4400	29.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.649	0.581	1.000	83.8
SYSDSP	047C2	4700	14.3	1.4	0.1	1.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.995	0.952	1.000	9.9
...																	
*** DEVICE LIST BY TOTAL I/O RATE ***																	
VOLUME SERIAL	DEV NUM	SSID	% I/O	I/O RATE	---CACHE READ	HIT DFW	RATE-- CFW	----	DASD I/O STAGE	DEL	NVS	OTHER	ASYNC RATE	TOTAL H/R	READ H/R	WRITE H/R	% READ
MVSTG4	041AD	4100	44.6	351.2	351.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.000	1.000	N/A	100.0
MVSLIB	04100	4100	27.7	217.7	217.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.000	1.000	1.000	100.0
SYSUSR	04A21	4A00	76.6	139.6	139.6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.000	1.000	N/A	100.0
MVSTG3	041AC	4100	16.5	129.8	129.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.000	1.000	N/A	100.0
SYSSM9	04D08	4D00	19.6	127.4	122.5	4.8	0.0	0.0	0.0	0.0	0.0	0.0	0.1	1.000	1.000	1.000	96.2
SYSSMD	04D15	4D00	15.1	98.4	94.6	3.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.000	1.000	1.000	
96.1																	

Figure 189. Cache Subsystem Activity Report - Top-20 Device Lists

Table 136. 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 137 on page 315 and Table 141 on page 320.	
Note: The reported storage capacities for cache and non-volatile storage (NVS) represents only the Cluster Processor Complex in the storage server, that controls the subsystem. Since a typical storage server has two clusters, you must double the reported capacities to get the actual sizes.	
CACHE	Amount of physical storage that is configured in this storage subsystem (in megabytes unless otherwise noted).
NVS	Amount of physical non-volatile storage (NVS) that is configured in this storage subsystem (in megabytes unless otherwise noted).
% I/O	Percentage of I/O requests to this volume, compared to the total number of I/O requests sent to the subsystem it is attached to. This is not the percentage of all I/O requests in the system.

Subsystem-level reporting

Cache subsystem status and overview

The first section shows the configuration of the selected storage subsystem. This includes the amount of cache configured (installed) and available for use, and how much storage, if any, is pinned because of a DASD failure. Finally, the overall subsystem status shows whether NON-VOLATILE STORAGE and CACHE FAST WRITE are active.

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													
z/OS 3.1		SYSTEM ID SYS1				DATE 04/25/2023		INTERVAL 10.00.010		PAGE 1			
		RPT VERSION 3.1 RMF				TIME 08.50.00							
SUBSYSTEM	2107-01	CU-ID	B005	SSID	1800	CDATE	04/25/2023	CTIME	08.50.00	CINT	10.00		
TYPE-MODEL	2107-981	MANUF	IBM	PLANT	75	SERIAL	0000000YT111						

C A C H E S U B S Y S T E M S T A T U S													

SUBSYSTEM STORAGE		NON-VOLATILE STORAGE				STATUS							
CONFIGURED	252128M			CONFIGURED	16384M	NON-VOLATILE STORAGE		- ACTIVE					
AVAILABLE	216498M			PINNED	0.0	CACHE FAST WRITE		- ACTIVE					
PINNED	0.0												
OFFLINE	0.0												

C A C H E S U B S Y S T E M O V E R V I E W													

TOTAL I/O	85821	CACHE I/O	85821										
TOTAL H/R	1.000	CACHE H/R	1.000										
C A C H E I / O R E Q U E S T S		---READ I/O REQUESTS---				---WRITE I/O REQUESTS---							
	COUNT	RATE	HITS	RATE	H/R	COUNT	RATE	FAST	RATE	HITS	RATE	H/R	%
NORMAL	35724	59.5	35724	59.5	1.000	15454	25.8	15454	25.8	15454	25.8	1.000	69.8
SEQUENTIAL	23279	38.8	23279	38.8	1.000	11364	18.9	11364	18.9	11364	18.9	1.000	67.2
CFW DATA	0	0.0	0	0.0	N/A	0	0.0	0	0.0	0	0.0	N/A	N/A
TOTAL	59003	98.3	59003	98.3	1.000	26818	44.7	26818	44.7	26818	44.7	1.000	68.8

R E Q U E S T S		C A C H E M I S S E S				-----MISC-----							
	READ	RATE	WRITE	RATE	TRACKS	RATE	DELATED DUE TO NVS		COUNT	RATE			
NORMAL	0	0.0	0	0.0	0	0.0	DELATED DUE TO CACHE		0	0.0			
SEQUENTIAL	0	0.0	0	0.0	0	0.0	DFW INHIBIT		0	0.0			
CFW DATA	0	0.0	0	0.0			ASYNC (TRKS)		14737	24.6			
TOTAL	0	RATE	0.0										

---CKD STATISTICS---		---RECORD CACHING---		---SYNCH I/O ACTIVITY---		---HOST ADAPTER ACTIVITY---		-----DISK ACTIVITY-----					
WRITE	0	READ MISSES	0	REQ	HITS	BYTES	BYTES	RESP	TIME	BYTES	BYTES		
WRITE HITS	0	WRITE PROM	17037	/SEC	/REQ	/REQ	/SEC		/REQ	/SEC	/SEC		
				READ	7.0	1.000	3.7K 362.4K	READ	0.000	0	0		
				WRITE	0.0	0.000	3.7K 163.8K	WRITE	17.124	16.0K	344.1K		

Figure 190. Cache Subsystem Activity Report - Status and Overview

Table 137. 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 when the cache interval started.
CTIME	Time when the cache interval started.
CINT	Cache interval time. In interval reports, the format is <i>mm.ss</i> , while in duration reports the format is <i>hh.mm.ss</i> .
TYPE-MODEL	Device type and model.

Table 137. Fields in the Cache Subsystem Activity Report - Header (continued)	
Field Heading	Meaning
TYPE MODEL MANUF PLANT SERIAL	The hardware description of the disk system.
VOLSER	Volume serial number of the reported DASD device (only for device-level reporting).
NUM	Device number of the reported DASD device (only for device-level reporting).

Table 138. Fields in the Cache Subsystem Activity Report - Subsystem Status	
Field Heading	Meaning
Note: The reported storage capacities for cache and non-volatile storage (NVS) represents only the Cluster Processor Complex in the storage server, that controls the subsystem. Since a typical storage server has two clusters, you must double the reported capacities to get the actual sizes.	
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.
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.

Table 139. 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 I/O REQUEST - The channel command DEFINE EXTENT specifies the way the cache will be used. There are three categories (NORMAL, SEQUENTIAL, CFW DATA) and a TOTAL value:	
NORMAL	Cache will be managed by <i>least-recently-used (LRU)</i> algorithm for making cache space available.
SEQUENTIAL	Tracks following the track assigned in the current CCW chain are promoted. They will be transferred from DASD to cache in anticipation of a short-term requirement.
CFW DATA	WRITE and READ-AFTER-WRITE requests are processed in cache. The data might not be written to DASD. Because CFW does not use the NVS, the application is responsible for restoring the data after a cache or subsystem failure.
TOTAL	This is either the sum of I/O requests, the total I/O rate, or the average hit ratio for the three categories previously described.
READ I/O REQUESTS - Cache I/O requests that searched or read data from DASD. This is the number of channel operations that had at least one SEARCH or READ command but no WRITE commands. This is counted for cached devices only.	
COUNT	Total number of SEARCH/READ requests.
HITS	Number of SEARCH/READ requests that completed without accessing the DASD.
H/R	Hit Ratio - number of SEARCH/READ hits compared to the total number of SEARCH/READ requests.
WRITE I/O REQUESTS - Cache I/O requests that wrote data to DASD. This is the number of channel commands that had at least one WRITE command. It is counted for cached devices only.	
COUNT	Total number of WRITE requests.
FAST	Total number of DASD/CACHE FAST WRITE requests.
HITS	Number of DASD/CACHE FAST WRITE requests that completed without accessing the DASD (fast write hit).
H/R	Number of DASD/CACHE FAST WRITE hits compared to the sum of all READ and WRITE requests (excluding ICL and BYPASS).
%READ	Percentage of READ requests compared to the sum of all READ and WRITE requests (excluding ICL and BYPASS).
CACHE MISSES -Cache misses are calculated as the difference between total I/O requests and the number of cache hits. They are shown for normal, sequential, and CFW requests. A cache miss occurs if a record that is to be read from or written to a data set is not found in the cache. In either case, the track in which this record is located is searched on the DASD and transferred to cache. Subsequent requests referring to the same record can then be processed in cache only, provided that the data has not yet been moved out because cache space is needed for other data. Other reasons for cache misses are DFW bypass and DFW inhibit. In the first case, a lack of NVS caused the data to be written to DASD immediately, while in the second case, usage of DFW is inhibited (for example, the device does not allow DFW).	
READ	Number of SEARCH/READ requests that needed access to DASD because the data could not be found in the cache.
WRITE	Number of WRITE requests that needed access to DASD because the data could not be found in the cache.
TRACKS	Number of tracks transferred from DASD to cache.
TOTAL	Total number of I/O requests that needed access to DASD because the data could not be found in the cache.

Table 139. Fields in the Cache Subsystem Activity Report - Subsystem Overview (continued)	
Field Heading	Meaning
MISC - Miscellaneous cache activities	
DELAYED DUE TO NVS	DASD Fast Write operations delayed due to non-volatile storage space constraints.
DELAYED DUE TO CACHE	Operations delayed due to cache space constraints.
DFW INHIBIT	If DASD FAST WRITE is active, this is the number of WRITE requests that <i>inhibited</i> DASD FAST WRITE. If DASD FAST WRITE is inactive, this is the number of WRITE requests that directly accessed the DASD, even with DASD FAST WRITE turned on.
ASYNCR	Number of tracks transferred from cache to DASD asynchronously to transfer from the channel to release space in the cache and the NVS. A high number of ASYNCR I/Os is an indicator for an over-committed cache or NVS.
CKD STATISTICS - CKD (Count-Key-Data) is a format used to store data on DASD. The counts shown in this section are contained in the total write count.	
WRITES	Number of write I/O requests in CKD format.
WRITE HITS	Number of write I/O requests in CKD format that could be resolved in the cache.
RECORD CACHING - Record caching is done dynamically upon a decision made by DCME or the microcode. It may improve overall cache performance if caching of whole tracks would waste cache storage. The decision is based on the number of I/Os, the hit ratio, and the locality of reference of a certain entity of data.	
READ MISSES	Number of instances in which a record requested for READ was not found in the cache, and access to DASD was required.
WRITE PROM	Number of instances in which a record requested for WRITE was found in the cache, and access to DASD was not required.
SYNCH I/O ACTIVITY - Synchronous I/O read and write activities	
REQ/SEC	The average number of synchronous I/O read and write requests per second.
HITS/REQ	Ratio of cache hits for synchronous I/O read and write requests.
HOST ADAPTER ACTIVITY ¹⁾ - I/O activity of normal, sequential and CFW read and write requests.	
BYTES/REQ	The average number of transferred bytes per read and write request.
BYTES/SEC	The average number of transferred bytes per second for read and write requests.
DISK ACTIVITY ¹⁾ - Transfer activity from hard disk to cache and vice versa.	
RESP TIME	Response time in milliseconds per read and write request.
BYTES/REQ	The average number of transferred bytes per read and write request.
BYTES/SEC	The average number of transferred bytes per second for read and write requests.
¹⁾ available for the IBM TotalStorage DS family	

Duration reports: An asterisk behind a status field or the field VOLUME SERIAL indicates that the contents of the field has changed during the duration interval.

Cache subsystem device overview and RAID rank activity

The first section lists all the devices in the subsystem. Each line shows the most important statistics for the device it represents. The I/O rate is divided into two groups (cache hits and DASD I/O), showing the different types of I/O activity in each group.

A RAID rank is a set of physical volumes. Several logical volumes as well as parallel access volumes are associated with a single RAID rank. Such a subsystem consists of multiple RAID ranks. If several higher utilized logical volumes are mapped to the same RAID rank, DASD skew is likely to appear. Knowing which logical volumes are associated to a certain RAID rank allows the storage administrator to move logical volumes from one RAID rank to another and thus optimally balance the load on the RAID ranks.

The RAID Rank Activity section in this report only appears for each 2105 subsystem. For 2107 subsystems, RAID rank activity is measured in the ESS Statistics report. The RAID Rank Activity

section provides information about each RAID rank belonging to the subsystem. It shows the physical characteristics of a RAID rank, the details for READ and WRITE requests, and lists the volumes with the highest activity.

C A C H E S U B S Y S T E M A C T I V I T Y																	
z/OS 3.1				SYSTEM ID SYS1				DATE 04/25/2023				INTERVAL 10.00.010				PAGE	2
				RPT VERSION 3.1 RMF				TIME 08.50.00									
SUBSYSTEM		2107-01		CU-ID		B005		SSID 1800		CDATE 04/25/2023		CTIME 08.50.00		CINT 10.00			
TYPE-MODEL		2107-981		MANUF		IBM		PLANT 75		SERIAL 0000000YT111							

C A C H E S U B S Y S T E M D E V I C E O V E R V I E W																	

VOLUME	DEV	XTNT	%	I/O	---CACHE	HIT	RATE--	----	DASD I/O	RATE----	ASYNC	TOTAL	READ	WRITE	%		
SERIAL	NUM	POOL	I/O	RATE	READ	DFW	CFW	STAGE	DEL	NVS	OTHER	RATE	H/R	H/R	H/R	READ	
*ALL			100.0	143.0	98.3	44.7	0.0	0.0	0.0	0.0	0.0	24.6	1.000	1.000	1.000	68.8	
RLB000	B000	0002	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	
RLB001	B001	0000	3.2	4.6	2.7	1.9	0.0	0.0	0.0	0.0	0.0	0.7	1.000	1.000	1.000	58.2	
RLB002	B002	0000	1.9	2.7	2.1	0.6	0.0	0.0	0.0	0.0	0.0	0.2	1.000	1.000	1.000	78.6	
RLB003	B003	0000	0.7	1.0	0.4	0.6	0.0	0.0	0.0	0.0	0.0	0.4	1.000	1.000	1.000	40.5	
RLB004	B004	0000	1.8	2.6	1.6	0.9	0.0	0.0	0.0	0.0	0.0	0.1	1.000	1.000	1.000	63.1	
RLB005	B005	0000	0.6	0.9	0.8	0.1	0.0	0.0	0.0	0.0	0.0	0.1	1.000	1.000	1.000	88.7	
RLB006	B006	0000	0.8	1.2	1.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	1.000	1.000	N/A	100.0	
RLB007	B007	0000	1.5	2.2	1.3	0.9	0.0	0.0	0.0	0.0	0.0	0.7	1.000	1.000	1.000	58.5	
RLB008	B008	0000	1.3	1.8	1.6	0.2	0.0	0.0	0.0	0.0	0.0	0.2	1.000	1.000	1.000	90.8	
RLB009	B009	0000	4.6	6.6	4.8	1.9	0.0	0.0	0.0	0.0	0.0	0.7	1.000	1.000	1.000	71.6	
RLB00A	B00A	0000	1.3	1.9	0.9	1.0	0.0	0.0	0.0	0.0	0.0	0.6	1.000	1.000	1.000	45.8	
RLB00B	B00B	0000	1.0	1.5	1.0	0.4	0.0	0.0	0.0	0.0	0.0	0.3	1.000	1.000	1.000	70.9	
RLB00C	B00C	0000	1.2	1.8	1.5	0.2	0.0	0.0	0.0	0.0	0.0	0.2	1.000	1.000	1.000	86.9	
RLB00D	B00D	0000	1.5	2.2	1.3	0.9	0.0	0.0	0.0	0.0	0.0	0.5	1.000	1.000	1.000	60.7	
RLB00E	B00E	0000	0.9	1.3	1.2	0.1	0.0	0.0	0.0	0.0	0.0	0.1	1.000	1.000	1.000	91.1	

Figure 191. Cache Subsystem Activity Report - Device Overview

Table 140. Fields in the Cache Subsystem Activity Report - Device Overview	
Field Heading	Meaning
VOLUME SERIAL	*ALL All volumes belonging to the reported storage subsystem For all other lines: the volume serial number.
DEV NUM	Device number.
RRID or XTNT POOL	RRID is displayed for 2105 subsystems and is the RAID rank identifier. XTNT POOL is displayed for 2107 subsystems and is the extent pool identifier.
% I/O	Percentage of I/O requests to this volume or category, compared to the total number of I/O requests sent to the subsystem.
I/O RATE	Number of I/O requests per second during the reporting interval.
CACHE HIT RATE - I/O rate of all cache hits.	
READ	Rate of SEARCH/READ requests that completed without accessing the DASD.
DFW	Rate of DFW requests.
CFW	Rate of WRITE and READ-AFTER-WRITE requests that are processed in cache.
DASD I/O RATE - I/O rate of all requests that accessed DASD.	
STAGE	Rate of normal or sequential I/O requests that accessed DASD.
DEL NVS	DASD Fast Write operations delayed due to non-volatile storage constraints.
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 141. 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.

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 139 on page 317](#).

C A C H E D E V I C E A C T I V I T Y													PAGE	5		
z/OS 3.1			SYSTEM ID S52			DATE 04/25/2023			INTERVAL 10.00.010							
			RPT VERSION 3.1 RMF			TIME 08.50.00										
SUBSYSTEM	2107-01	CU-ID	B005	SSID	1800	CDATE	04/25/2023	CTIME	08.50.00	CINT	10.00					
TYPE-MODEL	2107-981	MANUF	IBM	PLANT	75	SERIAL	0000000YT111									
VOLSER	RLB001	NUM	B001	EXTENT	POOL	0000										
----- C A C H E D E V I C E S T A T U S -----																

CACHE STATUS																
DASD FAST WRITE - ACTIVE																
PINNED DATA - NONE																

----- C A C H E D E V I C E A C T I V I T Y -----																

TOTAL I/O	2743	CACHE I/O	2743													
TOTAL H/R	1.000	CACHE H/R	1.000													
CACHE I/O																
REQUESTS	COUNT	READ	I/O	REQUESTS	HITS	RATE	H/R	COUNT	RATE	FAST	I/O	REQUESTS	HITS	RATE	H/R	%
NORMAL	747	1.2	747	1.2	1.000	444	0.7	444	0.7	444	0.7	1.000	62.7			
SEQUENTIAL	849	1.4	849	1.4	1.000	703	1.2	703	1.2	703	1.2	1.000	54.7			
CFW DATA	0	0.0	0	0.0	N/A	0	0.0	0	0.0	0	0.0	N/A	N/A			
TOTAL	1596	2.7	1596	2.7	1.000	1147	1.9	1147	1.9	1147	1.9	1.000	58.2			
REQUESTS	READ	CACHE	MISSSES	RATE	TRACKS	RATE										
NORMAL	0	0.0	0	0.0	0	0.0	DELAYED DUE TO NVS									
SEQUENTIAL	0	0.0	0	0.0	0	0.0	DELAYED DUE TO CACHE									
CFW DATA	0	0.0	0	0.0	0	0.0	DFW INHIBIT									
TOTAL	0	RATE	0.0	ASYNC (TRKS)								443	0.7			

---CKD STATISTICS---			---RECORD CACHING---			---SYNCH I/O ACTIVITY---			---HOST ADAPTER ACTIVITY---			-----DISK ACTIVITY-----				
WRITE	0	READ	MISSSES	0	READ	REQ	HITS	BYTES	BYTES	RESP	BYTES	BYTES	BYTES			
WRITE HITS	0	WRITE	PROM	698	WRITE	/SEC	/REQ	/REQ	/SEC	TIME	/REQ	/SEC	/SEC			
						0.2	1.000	3.5K	9.4K	0.000	0	0	0			
						0.0	0.000	4.1K	7.9K	20.533	23.3K	14.0K				

Figure 192. Cache Subsystem Activity Report - Cache Device Activity (device-level reporting)

Note: In the header of the Cache Device Activity Report, RRID shown for 2105 subsystems is the RAID rank identifier. EXTENT POOL shown for 2107 subsystems is the extent pool identifier.

Table 142. 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.	
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 Pinned data exists for the reported device. UNKNOWN PINNED STATUS The pinned status is undetermined.

The report contains less information for devices which had the cache offline at the end of the interval.

Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*. The following table shows the overview condition names for the Overview report.

Table 143. Overview names in the Cache Subsystem Activity Report		
Field Heading or Meaning	Subsystem Report	Device Report
Subsystem Status / Device Status.		
CACHING	CASSC	CADSC
NON-VOLATILE STORAGE	CASSNVS	
Subsystem Overview / Device Activity		
TOTAL I/O	CASTOT	CADTOT
CACHE I/O	CASCTOT	CADCTOT
CACHE OFFLINE	CASCOFF	
TOTAL H/R	CASHRT	CADHRT
CACHE H/R	CASHR	CADHR
READ I/O REQUESTS RATE NORMAL	CASRN	CADRN
READ I/O REQUESTS RATE SEQUENTIAL	CASRS	CADRS
READ I/O REQUESTS RATE CFWDATA	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

Table 143. Overview names in the Cache Subsystem Activity Report (continued)

Field Heading or Meaning	Subsystem Report	Device Report
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
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

Table 143. Overview names in the Cache Subsystem Activity Report (continued)

Field Heading or Meaning	Subsystem Report	Device Report
CACHE MISSES RATE TOTAL	CASMT	CADMT
MISC (Miscellaneous) DFW operations delayed due to NVS storage constraints rate	CASDFWB	CADDFWB
MISC (Miscellaneous) Operations delayed due to cache space constraints 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
SYNCH I/O ACTIVITY READ REQ/SEC	CASSRR	CADSRR
SYNCH I/O ACTIVITY READ HITS/REQ	CASSRH	CADSRH
SYNCH I/O ACTIVITY WRITE REQ/SEC	CASSWR	CADSWR
SYNCH I/O ACTIVITY WRITE HITS/REQ	CASSWH	CADSWH
HOST ADAPTER ACTIVITY BYTES/REQ READ	CASBRR	CADBRR
HOST ADAPTER ACTIVITY BYTES/SEC READ	CASBRS	CADBRS
HOST ADAPTER ACTIVITY BYTES/REQ WRITE	CASBWR	CADBWR
HOST ADAPTER ACTIVITY BYTES/SEC WRITE	CASBWS	CADBWS
DISK ACTIVITY RESP TIME READ	CASDRRT	CADDRRT
DISK ACTIVITY BYTES/REQ READ	CASDRBR	CADDRBR
DISK ACTIVITY BYTES/SEC READ	CASDRBS	CADDRBS
DISK ACTIVITY RESP TIME WRITE	CASDWRT	CADDWRT
DISK ACTIVITY BYTES/REQ WRITE	CASDWBR	CADDWBR
DISK ACTIVITY BYTES/SEC WRITE	CASDWBS	CADDWBS
Subsystem Device Overview		
I/O RATE (volser)	CADT	
I/O RATE (*ALL)	CASAT	
I/O RATE (*CACHE)	CASCT	
I/O RATE (*CACHE-OFF)	CASOT	
DASD I/O RATE STAGE (volser)	CADSTG	
DASD I/O RATE STAGE (*ALL)	CASASTG	
DASD I/O RATE STAGE (*CACHE)	CASCSTG	
% I/O (*CACHE-OFF)	CASCOIO	
RAID RANK READ REQ RATE	CARRRT	

Table 143. Overview names in the Cache Subsystem Activity Report (continued)

Field Heading or Meaning	Subsystem Report	Device Report
RAID RANK READ REQ AVG MB	CARRMB	
RAID RANK READ REQ MB/S	CARRMBS	
RAID RANK READ REQ RTIME	CARRRTIM	
RAID RANK WRITE REQ RATE	CARWRT	
RAID RANK WRITE REQ AVG MB	CARWMB	
RAID RANK WRITE REQ MB/S	CARWMBS	
RAID RANK WRITE REQ RTIME	CARWRTIM	

CF - Coupling Facility Activity report

A Coupling Facility Activity report is produced for each coupling facility attached to the sysplex. It provides the following information:

- Coupling Facility usage summary
- Coupling Facility structure activity
- Subchannel activity
- CF to CF activity

How to request this report

Monitor III gathers data for this report automatically. If you want to suppress gathering, you have to disable writing SMF record type 74.4.

Note: For the description of the SMF type 74 subtype 4 record, see *z/OS MVS System Management Facilities (SMF)*.

If optimized coupling facility hardware data collection is active, in which case the data gatherer collects coupling facility hardware statistics only on one system in the sysplex, you must include SMF records from all data gathering systems, or at a minimum, the data gathering system that contains the coupling facility information. For more information about optimized collection, see [Controlling sysplex-wide optimized coupling facility hardware data collection \(CFOPT\)](#) in *z/OS Data Gatherer User's Guide*.

To produce this report, specify

```
SYSRPTS(CF)
```

This report is also available in XML output format. [How to work with Postprocessor XML reports in z/OS Resource Measurement Facility User's Guide](#) provides all required information on how to produce and view XML reports.

Example URL for the DDS API:

```
http://ddshost:8803/gpm/rmfpp.xml?reports=CF
```

Contents of the report

A Coupling Facility Activity report is produced for each coupling facility attached to the sysplex. [Figure 193 on page 325](#) gives an example of the overall structure of the Coupling Facility Activity report. It shows the sequencing of the report sections:

- Coupling Facility Usage Summary

- Coupling Facility Structure Activity
- Subchannel Activity
- CF to CF Activity

For a complete Coupling Facility Activity report, it is recommended to combine data from all of the systems in the sysplex. If data from one or more systems is missing, the Structure and Subchannel Activity sections of the report are incomplete. In addition, the PRIM (primary) and SEC (secondary) indicators of synchronously duplexed structures might be missing in the Usage Summary section because this information is gathered only on one member of the sysplex (sysplex master gathering).

SYSPLEX = UTCPLXJ8




Figure 193. Structure of the Coupling Facility Activity report

Coupling Facility Usage Summary section

This section of the Coupling Facility Activity report gives a snapshot of coupling facility storage and specific structure data at the end of the reporting interval. It lists all the structures occupying space either in the coupling facility real storage or in storage class memory (SCM), including those with no currently active connections. For asynchronously duplexed secondary structures, it provides a summary on consumed processing times and performance counts related to asynchronous duplexing. Data shown

for asynchronous duplex sync up requests are accumulated over all reported systems that have a connector to the structure.

All structure summary data is grouped in the report by structure type (LIST, LOCK, CACHE and UNKN in that order). Within the structure type, the report lines are in alphanumeric order by structure name. This ordering puts the structure data in the same relative position across interval reports.

COUPLING FACILITY ACTIVITY														PAGE 1
z/OS 3.1		SYSPLEX UTCPLXJ8 RPT VERSION 3.1 RMF			DATE 04/25/2023 TIME 12.00.00			INTERVAL 030.00.000 CYCLE 01.000 SECONDS						

COUPLING FACILITY NAME = X7CFP87														
TOTAL SAMPLES(AVG) = 1800 (MAX) = 1800 (MIN) = 1799														

COUPLING FACILITY USAGE SUMMARY														

GENERAL STRUCTURE SUMMARY														

TYPE	STRUCTURE NAME	STATUS CHG	ENC	ALLOC SIZE	% OF CF STOR	# REQ	% OF ALL REQ	% OF CF UTIL	AVG REQ/ SEC	LST/DIR ENTRIES TOT/CUR	DATA ELEMENTS TOT/CUR	LOCK ENTRIES TOT/CUR	DIR REC/ DIR REC XI'S	
LIST	DBSVPLX7_SCA	ACTIVE	YES	5M	0.0	10754	1.5	1.8	5.97	3484	6743	N/A	N/A	
	DFHXQLS_G2POOL1	PRIM								182	622	N/A	N/A	
		ACTIVE	NO	63M	0.4	0	0.0	0.0	0.00	96K	96K	N/A	N/A	
LIST	ISTGENERIC	ACTIVE	YES	11M	0.0	20850	0.0	0.0	11.58	1822	1859	N/A	N/A	
	THRLSTSCMKP1_1	PRIM								25K	499	4	N/A	
		ACTIVE	NO	4G	4.1	20114K	4.4	16.5	11174	47	2	0	N/A	
LOCK	DBTPLX5_ALOCK	ACTIVE	N/A	44M	0.3	68212	9.4	9.2	37.90	1428K	8569K	1024	N/A	
CACHE	IRRXCF00_B002	SEC A								1262K	7060K	0	N/A	
		ACTIVE	NO	2M	0.0	0	0.0	0.0	0.00	11K	0	17M	N/A	
										0	0	0	N/A	
										184	181	N/A	0	
										4	1	N/A	0	

STRUCTURE TOTALS				15G	57.0	29430K	100	90.5	16350					

SCM STRUCTURE SUMMARY														

TYPE	STRUCTURE NAME	SCM SPACE ALG %USED	AUGMENTED EST. MAX/ %USED	LST ENTRY EST. MAX/ CUR	LST ELEM EST. MAX/ CUR	--- SCM READ --- CNT/BYTE AVG ST/ X'FERRED STD_DEV	--- SCM WRITE --- CNT/BYTE AVG ST/ X'FERRED STD_DEV	SCM AUX ENABLED CMD/%ALL	DELAYED FAULTS CNT/%ALL					
LIST	THRLSTSCMKP1_1	KP1 16384M 0.0	754M 0.3	9585K 0	57508K 0	4790 5023M 1353.8 832.0	3944 4136M 1762.3 736.6	0 0.0	58749 0.0					

ASYNCHRONOUS CF DUPLEXING SUMMARY														

TYPE	STRUCTURE NAME	TOTAL	ASYNC DUPLEX CF --TRANSMIT TIME-- AVG	OPERATIONS --SERVICE TIME-- STD_DEV	--- ASYNC DUPLEX SYNC_UP REQUESTS --- TOTAL #SUSPEND	--SUSPEND TIME-- AVG	STD_DEV							
LOCK	DBTPLX5_ALOCK	19432K	1.1	22.6	1.2	23.8	43197	1	744.0	0.0	0.0			
	DBTSTPLX_LCK	6700K	1.1	22.2	1.1	23.8	7420	0	0.0	0.0	0.0			

STORAGE SUMMARY														

				ALLOC SIZE	% OF CF STORAGE	----- DUMP SPACE ----- % IN USE	MAX % REQUESTED							
TOTAL CF STORAGE USED BY STRUCTURES				6202M	6.2									
TOTAL CF DUMP STORAGE				1024M	1.0	0.0	0.0							
TOTAL CF AUGMENTED SPACE				2M	0.0									
TOTAL CF STORAGE AVAILABLE				93138M	92.8									
TOTAL CF STORAGE SIZE				100366M										
				ALLOC SIZE	% ALLOCATED									
TOTAL CONTROL STORAGE DEFINED				100366M	5.7									
TOTAL DATA STORAGE DEFINED				0K	0.0									
				ASSIGNED	% IN USE	SUM MAX SCM								
TOTAL CF STORAGE CLASS MEMORY				16384M	0.0	16384M								

PROCESSOR SUMMARY														

COUPLING FACILITY		2827	MODEL H66	CFLEVEL 19	DYNDISP OFF									
AVERAGE CF UTILIZATION (% BUSY)		53.9	LOGICAL PROCESSORS:		DEFINED SHARED	2 0	EFFECTIVE AVG WEIGHT	2.0 0.0						

Figure 194. 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 144. Fields in the Coupling Facility Activity Report - Usage Summary	
Field Heading	Meaning
TOTAL SAMPLES (AVG), (MAX), (MIN)	Average (AVG), the maximum (MAX) and minimum (MIN) sample count for all systems connected to this coupling facility.
GENERAL STRUCTURE SUMMARY	
TYPE	Indicates whether the structure is a list, lock, or cache structure. The structures being reported are grouped by structure type. UNKN indicates a structure for which there was no activity during the interval but that is still allocated in the coupling facility. There are no structure activity details for UNKN structures.
STRUCTURE NAME	The name given to the structure by the coupling facility policy specification in the Function Couple Data Set. It is up to 16 characters and is unique within a sysplex.
STATUS	Indicates status of the structure at the end of the interval: 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. An appended "A" indicates that the structure is asynchronously duplexed. ACTIVE SEC The structure is the rebuilt-new (secondary) structure in a duplexing rebuild process. An appended "A" indicates that the structure is asynchronously duplexed. 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. The PRIM and SEC indicators of synchronously duplexed structures might not appear if data from one or more systems in the sysplex is missing.
CHG	X indicates that the status of this structure changed during the reporting interval.
ENC	YES indicates that the CF structure is encrypted. NO indicates that the CF structure is not encrypted or cannot be encrypted. A CF structure cannot be encrypted if the structure is only used by a down-level system in the Parallel Sysplex, that does not support encryption of CF structures. N/A indicates that encryption of lock (LOCK) structures is not supported.
ALLOC SIZE	The number of bytes set aside in the coupling facility for this structure by the coupling facility policy in the Function Couple Data Set. Storage is allocated in increments of 4K bytes. This storage consists of both control and data storage.
% OF CF STOR	The percentage of the total coupling facility storage allocated to this structure.
# 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.

Table 144. Fields in the Coupling Facility Activity Report - Usage Summary (continued)

Field Heading	Meaning
% OF CF UTIL	The percentage of CF processor time used by the structure. The structure execution time is related to the total CF-wide processor busy time. The sum of the values in this column is less than 100%, because not all CF processor time is attributable to structures. N/A is shown in this field if the CF level is lower than 15.
AVG REQ/SEC	The average number of requests per second for this structure.
LST/DIR ENTRIES	TOT Maximum number of list or directory entries that can reside in coupling facility real storage for the structure. CUR Number of structure list or directory entries which are currently in use and reside in coupling facility real storage. N/A in this and the following fields indicates that the information is not applicable.
DATA ELEMENTS	TOT Maximum number of list elements that can reside in coupling facility real storage. CUR Number of structure list elements which are currently in use and reside in coupling facility real storage.
LOCK ENTRIES	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 Structure Activity section). A high value is an indicator for a performance problem in this structure.
SCM STRUCTURE SUMMARY(See note 2)	
ALG	Type of algorithm that is used by the coupling facility to control the movement of structure objects between coupling facility real storage and storage class memory: KP1 KeyPriority1 UNK Unknown
SCM SPACE	MAX Maximum amount of storage class memory that this structure can use (in bytes). %USED Percentage of maximum amount of storage class memory that is in use by this structure.
AUGMENTED	EST.MAX Estimated maximum amount of CF space that may be assigned as augmented space for this structure (in bytes). %USED Percentage of maximum augmented space that is in use by this structure.

Table 144. Fields in the Coupling Facility Activity Report - Usage Summary (continued)

Field Heading	Meaning
LST ENTRY	EST.MAX Estimated maximum number of list entries that may reside in storage class memory for this structure. CUR Number of existing structure list entries that reside in storage class memory.
LST ELEM	EST.MAX Estimated maximum number of list elements that may reside in storage class memory for this structure. CUR Number of existing structure list elements that reside in storage class memory.
SCM READ	CNT The number of read operations against storage class memory that were either initiated • by a reference to list structure objects residing in storage class memory, or • as a prefetch operation in order to retrieve list structure objects in storage class memory that are expected to be referenced. BYTE X'FERRED SCM read bytes transferred. This is the number of bytes transferred from storage class memory to CF. AVG ST Average service time per SCM read operation to storage class memory in microseconds. STD_DEV Standard deviation of the service time for SCM read operations to storage class memory in microseconds.
SCM WRITE	CNT The number of list write operations performed to storage class memory. BYTE X'FERRED SCM write bytes transferred. This is the number of bytes transferred from CF storage to storage class memory. AVG ST Average service time per SCM write operation to storage class memory in microseconds. STD_DEV Standard deviation of the service time for SCM write operations to storage class memory in microseconds.
SCM AUX ENABLED	CMD SCM auxiliary enabled command count. This is the number of commands that required the use of CF auxiliary frames. %ALL Percentage of the SCM auxiliary enabled command count in relation to all requests for this structure.
DELAYED FAULTS	CNT Number of commands and for multiple list-entry commands, the number of list item references that were delayed due to a fault condition resulting in a required access to storage class memory. %ALL Percentage of delayed faults in relation to all requests for the structure.
ASYNCHRONOUS CF DUPLEXING SUMMARY(See note 3)	

Table 144. Fields in the Coupling Facility Activity Report - Usage Summary (continued)

Field Heading	Meaning
ASYNCR DUPLEX CF OPERATIONS	TOTAL Number of asynchronous duplex operations transmitted from the primary to the secondary structure that completed in the secondary structure. TRANSMIT TIME AVG Average asynchronous duplex operation transmission time for operations sent from the primary to the secondary structure, in microseconds. TRANSMIT TIME STD_DEV Standard deviation of the average asynchronous duplex operation transmission time. SERVICE TIME AVG Average service time to transfer the asynchronous duplex operation to the secondary structure and complete the operation in the secondary structure, in microseconds. SERVICE TIME STD_DEV Standard deviation of the average service time to transfer and complete the asynchronous duplex operation in the secondary.
ASYNCR DUPLEX SYNC_UP REQUESTS	TOTAL Total number of user requests that requested sync up with the primary structure. User requests are summed up over all systems connected to the structure. #SUSPEND Number of user requests that were suspended waiting for asynchronous duplex operations to complete in the secondary structure. Suspend counts are summed up over all systems connected to the structure. SUSPEND TIME AVG Average accumulated suspend time for suspended requests waiting for asynchronous duplex operations to complete in the secondary structure, in microseconds. SUSPEND TIME STD_DEV Standard deviation of the average accumulated suspend time.
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. These totals do not necessarily represent 100% of the facility activity for the interval. There is some amount of storage and request activity overhead that is not attributable to individual structures. For example, the total for # REQ will usually be less than the sum of the TOTAL REQ from the Subchannel Activity Report because the subchannel numbers include facility management command counts whereas the structure numbers do not.
TOTAL CF DUMP STORAGE	Amount and percentage 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. 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 AUGMENTED SPACE	ALLOC SIZE Total amount of CF storage used by all structures as augmented space (in bytes). % OF CF STORAGE Percentage of CF storage used by all structures as augmented space.
TOTAL CF STORAGE AVAILABLE	The amount and percentage of coupling facility space that is not allocated to any structure, not allocated as dump space, and not allocated as augmented space.

Table 144. Fields in the Coupling Facility Activity Report - Usage Summary (continued)

Field Heading	Meaning
TOTAL CF STORAGE SIZE	The total amount of storage in the coupling facility, including both allocated and available space. This value does not include the storage required by the coupling facility code itself, so that it differs from the storage assigned to the coupling facility on the HMC.
TOTAL CONTROL STORAGE DEFINED, TOTAL DATA STORAGE DEFINED, % ALLOCATED	The amount of coupling facility storage that is allowed to be occupied by control information (CONTROL STORAGE) or data (DATA STORAGE). For each structure, plus the dump area, a certain amount of control and data storage is allocated. The coupling facility defines an area called control storage; structure control information is restricted to that area. The remaining storage is called data storage and is used for structure data. If the data storage area becomes full, structure data can then be allocated from the control storage area. If TOTAL DATA STORAGE DEFINED is zero, it means control information can reside anywhere on the coupling facility and there are no allocation restrictions. If the % ALLOCATED field for control storage shows a percentage approaching 100, it means the control storage is close to being completely allocated even though the CF SPACE AVAILABLE field may still show an amount of total free space. Possible customer actions include: • 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.
TOTAL CF STORAGE CLASS MEMORY	ASSIGNED Total CF storage class memory. This is the amount of storage class memory that may be concurrently used as structure extensions. Storage is assigned in increments of 4K bytes. % IN USE Percentage of storage class memory that is in use by all structures of the coupling facility. SUM MAX SCM Sum of the storage class memory maxima defined for all structures of the coupling facility.
PROCESSOR SUMMARY	
COUPLING FACILITY	Coupling facility processor type.
MODEL	Coupling facility processor model.
CFLEVEL	Coupling facility architected function level.
DYNDISP(See note 1)	The dynamic CF dispatching status (ON, OFF, or THIN). THIN indicates that coupling thin interrupts are enabled for the coupling facility (only for CFLEVEL 19 or higher).
AVERAGE 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.
LOGICAL PROCESSORS DEFINED(See note 1)	Number of logical processors defined for the coupling facility.

Table 144. Fields in the Coupling Facility Activity Report - Usage Summary (continued)

Field Heading	Meaning
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 time waiting for work. The reported value shows the ratio of the LPAR dispatch time (CFCC execute and wait time) to the RMF interval length. 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.
LOGICAL PROCESSORS SHARED	The number of shared processors defined for the coupling facility.
LOGICAL PROCESSORS AVG WEIGHT ^(See note 1)	The average weight of shared processors, which is the sum of shared processor weights related to the number of shared processors.

Notes:

1. For CFLEVEL lower than 15, this field is not displayed.
2. SCM statistics are included in the SCM Structure Summary only for those structures that can make use of the SCM storage extension and have set a non-zero maximum SCM size. If none of the structures is configured to exploit SCM, the SCM Structure Summary displays message: "NO STORAGE CLASS MEMORY DATA AVAILABLE".
3. If no structure is configured to exploit asynchronous duplexing, the Asynchronous CF Duplexing Summary includes an informational message "NO ASYNCHRONOUS CF DUPLEXING DATA AVAILABLE".

Coupling Facility Structure Activity section

This section of the Coupling Facility Activity report has detail for each active structure in the coupling facility, including activity data for each system connected to the structure during the reporting interval.

COUPLING FACILITY ACTIVITY

PAGE 1

z/OS 3.1

SYSPLEX SYSPLEX1
RPT VERSION 3.1 RMFDATE 04/25/23
TIME 00.00.00INTERVAL 030.00.000
CYCLE 01.000 SECONDS

COUPLING FACILITY NAME = CF2

COUPLING FACILITY STRUCTURE ACTIVITY

STRUCTURE NAME = DSND801_GBP11			TYPE = CACHE		STATUS = ACTIVE		PRIMARY		ENCRYPTED = NO			
SYSTEM NAME	# REQ TOTAL AVG/SEC	# REQ	REQUESTS		-SERV AVG	TIME(MIC)- STD_DEV	REASON	# REQ	% OF REQ	DELAYED REQUESTS		
			% OF ALL							--- AVG	TIME(MIC) /DEL	--- /ALL
SYS1	80646K 44803	SYNC	157K	0.1	222.9	249.7	NO SCH	78M	96.9	1674	858.4	1622
		ASYN	80M	29.2	597.5	310.7	PR WT	0	0.0	0.0	0.0	0.0
		CHNGD	78M	28.3	INCLUDED	IN ASYNC	PR CMP	0	0.0	0.0	0.0	0.0
		SUPPR	0	0.0			DUMP	0	0.0	0.0	0.0	0.0
							MONOP	75K	0.1	1658	538.6	1.5
SYS2	78721K 43734	SYNC	372K	0.1	197.5	214.8	NO SCH	75M	95.1	997.3	617.7	948.5
		ASYN	78M	28.4	584.8	315.5	PR WT	0	0.0	0.0	0.0	0.0
		CHNGD	75M	27.2	INCLUDED	IN ASYNC	PR CMP	0	0.0	0.0	0.0	0.0
		SUPPR	0	0.0			DUMP	0	0.0	0.0	0.0	0.0
							MONOP	0	0.0	0.0	0.0	0.0
SYS3	68465K 38036	SYNC	1968K	0.7	506.2	316.1	NO SCH	34K	0.0	1151	873.1	0.6
		ASYN	66M	24.1	685.2	320.2	PR WT	0	0.0	0.0	0.0	0.0
		CHNGD	4081K	1.5	INCLUDED	IN ASYNC	PR CMP	0	0.0	0.0	0.0	0.0
		SUPPR	0	0.0			DUMP	0	0.0	0.0	0.0	0.0
							MONOP	4048K	5.9	1261	716.6	74.5
SYS4	48145K 26747	SYNC	96K	0.0	126.2	274.5	NO SCH	47M	98.0	8528	4087	8360
		ASYN	48M	17.4	612.6	425.1	PR WT	0	0.0	0.0	0.0	0.0
		CHNGD	48M	17.3	INCLUDED	IN ASYNC	PR CMP	0	0.0	0.0	0.0	0.0
		SUPPR	0	0.0			DUMP	0	0.0	0.0	0.0	0.0
							MONOP	489K	1.0	7654	3203	77.7
TOTAL	276M 153.3K	SYNC	2593K	0.9	430.8	327.4	NO SCH	200M	72.6	3036	3709	2203
		ASYN	273M	99.1	617.9	339.4	PR WT	0	0.0	0.0	0.0	0.0
		CHNGD	205M	74.3			PR CMP	0	0.0	0.0	0.0	0.0
		SUPPR	0	0.0			DUMP	0	0.0	0.0	0.0	0.0
							MONOP	4612K	1.7	1945	2326	32.5
											-- DATA ACCESS --	
											READS	5550362
											WRITES	98230
											CASTOUTS	22786
											XI'S	1657682

STRUCTURE NAME =		DSND800_LOCK1		TYPE = LOCK		STATUS = ACTIVE		ENCRYPTED = N/A						
SYSTEM NAME	# REQ TOTAL AVG/SEC	# REQ	# REQ	REQUESTS		REASON	# REQ	% OF REQ	DELAYED REQUESTS		EXTERNAL REQUEST CONTENTIONS			
				% OF ALL	-SERV TIME(MIC)- STD_DEV				--- AVG	TIME(MIC) - STD_DEV		/DEL	/ALL	
SYS1	5428K 3015	SYNC	4263K	18.9	7.1	10.7	NO SCH	4972K	91.6	25.5	32.8	23.4	REQ TOTAL	5538K
		ASYN	1165K	5.2	27.2	45.6	PR WT	0	0.0	0.0	0.0	0.0	REQ DEFERRED	172K
		CHNGD	1164K	5.2	INCLUDED	IN ASYNC	PR CMP	0	0.0	0.0	0.0	0.0	-CONT	8759
		SUPPR	0	0.0									-FALSE CONT	4875
								MONOP	0	0.0	0.0	0.0	0.0	
SYS2	6454K 3586	SYNC	5954K	26.3	7.9	13.1	NO SCH	5647K	87.5	19.2	28.1	16.8	REQ TOTAL	6563K
		ASYN	500K	2.2	56.5	89.3	PR WT	0	0.0	0.0	0.0	0.0	REQ DEFERRED	84K
		CHNGD	496K	2.2	INCLUDED	IN ASYNC	PR CMP	0	0.0	0.0	0.0	0.0	-CONT	9354
		SUPPR	0	0.0									-FALSE CONT	4982
								MONOP	0	0.0	0.0	0.0	0.0	
SYS3	5531K 3073	SYNC	5328K	23.6	20.3	10.9	NO SCH	200	0.0	246.9	470.2	0.0	REQ TOTAL	5625K
		ASYN	203K	0.9	124.2	141.9	PR WT	0	0.0	0.0	0.0	0.0	REQ DEFERRED	34K
		CHNGD	200	0.0	INCLUDED	IN ASYNC	PR CMP	0	0.0	0.0	0.0	0.0	-CONT	8073
		SUPPR	0	0.0									-FALSE CONT	4212
								MONOP	0	0.0	0.0	0.0	0.0	
SYS4	5196K 2887	SYNC	5095K	22.5	5.9	20.7	NO SCH	129K	2.5	22.1	30.4	0.5	REQ TOTAL	5262K
		ASYN	101K	0.4	201.0	131.4	PR WT	0	0.0	0.0	0.0	0.0	REQ DEFERRED	23K
		CHNGD	82K	0.4	INCLUDED	IN ASYNC	PR CMP	0	0.0	0.0	0.0	0.0	-CONT	8404
		SUPPR	0	0.0									-FALSE CONT	3960
								MONOP	0	0.0	0.0	0.0	0.0	
TOTAL		22608K 12560	SYNC	21M 1968K 1743K 0	91.3 8.7 7.7 0.0	10.4 53.5	15.6 90.9	NO SCH	11M 47.5 22.2 30.6 10.5	0.0 0.0 0.0 0.0 0.0	REQ TOTAL	23M 312K 35K 15K		

STRUCTURE NAME = IRRXCF00_P001		TYPE = CACHE			STATUS = ACTIVE			ENCRYPTED = NO					
SYSTEM NAME	# REQ TOTAL AVG/SEC	SYNCH	# REQ	REQUESTS		REASON	# REQ	% OF REQ	REQUESTS				
				% OF ALL	-SERV TIME(MIC)- AVG STD_DEV				--- AVG	TIME(MIC) STD_DEV	---/ALL		
SYS1	2004K 1113	SYNC	1976K	19.0	15.5	17.6	NO SCH	3	0.0	451.7	541.8	0.0	
		ASYN	28K	0.3	179.6	393.8	PR WT	0	0.0	0.0	0.0	0.0	
		CHNGD	3	0.0	INCLUDED	IN ASYNC	PR CMP	0	0.0	0.0	0.0	0.0	
		SUPPR	0	0.0			DUMP	0	0.0	0.0	0.0	0.0	
							MONOP	0	0.0	0.0	0.0	0.0	
SYS2	2402K 1334	SYNC	2010K	19.3	29.5	16.4	NO SCH	2585	0.1	15.1	14.2	0.0	
		ASYN	393K	3.8	81.9	95.8	PR WT	0	0.0	0.0	0.0	0.0	
		CHNGD	2597	0.0	INCLUDED	IN ASYNC	PR CMP	0	0.0	0.0	0.0	0.0	
		SUPPR	0	0.0			DUMP	0	0.0	0.0	0.0	0.0	
							MONOP	0	0.0	0.0	0.0	0.0	
SYS3	1573K 873.7	SYNC	1246K	12.0	30.7	17.4	NO SCH	70	0.0	4.8	4.5	0.0	
		ASYN	327K	3.1	78.1	87.7	PR WT	0	0.0	0.0	0.0	0.0	
		CHNGD	70	0.0	INCLUDED	IN ASYNC	PR CMP	0	0.0	0.0	0.0	0.0	
		SUPPR	0	0.0			DUMP	0	0.0	0.0	0.0	0.0	
							MONOP	0	0.0	0.0	0.0	0.0	
SYS4	4416K 2455	SYNC	3579K	34.4	15.0	15.9	NO SCH	789K	34.9	78.3	133.7	27.3	
		ASYN	835K	8.1	128.6	143.5	PR WT	0	0.0	0.0	0.0	0.0	
		CHNGD	793K	7.6	INCLUDED	IN ASYNC	PR CMP	0	0.0	0.0	0.0	0.0	
		SUPPR	0	0.0			DUMP	0	0.0	0.0	0.0	0.0	
							MONOP	0	0.0	0.0	0.0	0.0	
TOTAL		10397K 5776	SYNCH	8813K 1584K 796K 0	84.8 15.2 7.7 0.0	21.0 56.5	19.2 94.0	NO SCH PR WT PR CMP DUMP MONOP	792K 0 0 0 0	7.6 0.0 0.0 0.0 0.0	78.0 0.0 0.0 0.0 0.0	5.9 0.0 0.0 0.0 0.0	-- DATA ACCESS -- READS 2784547 WRITES 2077781 CASTOUTS 0 XI'S 2763783

Figure 195. Coupling Facility Activity Report - Structure Activity

The following table explains the field headings in the Structure Activity section.

Table 145. 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 144 on page 327 .
ENCRYPTED	YES indicates that the CF structure is encrypted. NO indicates that the CF structure is not encrypted or cannot be encrypted. A CF structure cannot be encrypted if the structure is only used by a down-level system in the Parallel Sysplex, that does not support encryption of CF structures. N/A indicates that encryption of lock (LOCK) structures is not supported.
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. Systems with no connection to the structure are not reported in the coupling facility structure activity section.
# REQ TOTAL # REQ AVG/SEC	The sum of all requests (internal and external) that utilize the subchannel. Specifically: - External requests to send/receive data on behalf of a structure. The sum of synchronous and asynchronous requests completed against any structure within this coupling facility per second. This includes requests that changed from synchronous to asynchronous. - Internal requests that utilize the subchannels (but are not aggregated by the structure). 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 four categories described hereafter: SYNC, ASYNC, CHNGD, and SUPPR.
SYNC	Total number of hardware operations that started and completed synchronously to the coupling facility on behalf of connectors to the structure.
ASYNC	Total number of hardware operations that started and completed asynchronously to the coupling facility on behalf of connectors to the structure. The service time is the time for all ASYNC requests (ASYNC and CHNGD).
CHNGD	Total number of hardware operations that changed from synchronous to asynchronous because the operation could not be serviced as synchronous operation. This field reports only those operations which were changed due to a subchannel busy condition and can be used as an indicator of a shortage of subchannel resources. Conversions caused by heuristic sync/async algorithms used to optimize the coupling efficiency of workloads using the CF are not included.
SUPPR	Number of requests whose execution was suppressed by the coupling facility in order to avoid a potential serialization deadlock condition across a duplexed pair of structures. This field does not apply to asynchronously duplexed structures.

Table 145. Fields in the Coupling Facility Activity Report - Structure Activity (continued)

Field Heading	Meaning
# REQ % OF ALL (valid for SYNC, ASYNC, CHNGD, and SUPPR)	The number of requests for this structure, and the percentage this represents of all requests for this structure from any system.
SERVICE TIME - AVG	The average time in microseconds required to satisfy a coupling facility request for this structure.
SERVICE TIME - STD_DEV	The standard deviation of service time for this structure. Even though the average time looks acceptable, the standard deviation could be high, indicating that there is a wide fluctuation in service times for requests. In this case, analyze the coupling facility configuration for possible path or coupling facility bottlenecks in the Subchannel Activity section.
DELAYED REQUESTS	These columns list possible contention reasons for requests sent to the coupling facility.
REASON	The reason for a delayed request can be either a subchannel contention (NO SCH), dump serialization (DUMP) or CF monopolization avoidance (MONOP). For synchronous duplexed requests, also peer subchannel wait time (PR WT) and waiting-for-peer-completion time (PR CMP) is reported. A duplexed request requires two subchannels. PR WT is the time (in microseconds) between the moment when the request was sent to the other duplexed structure instance and when it is sent to this one. PR CMP is the time (in microseconds) between the moment when this structure responded to z/OS and when the other structure instance responded. Both subchannels are busy until the responses from both structure instances are processed by z/OS. If the coupling facility tasks receive excessive requests for the same structure, the coupling facility indicates this situation to the operating system so that cross-system extended services (XES) stops sending requests to the coupling facility for the specific structure. This feature is called <i>CF monopolization avoidance</i> . If a system does not have valid CF monopolization avoidance data, MONOP values are displayed in the report as N/A, and total MONOP values are unavailable for the structure.
# 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 DEFERRED 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.

Table 145. Fields in the Coupling Facility Activity Report - Structure Activity (continued)	
Field Heading	Meaning
-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.
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.

COUPLING FACILITY ACTIVITY															
z/OS 3.1		SYSPLEX UTCPLX38				DATE 04/25/2023				INTERVAL 030.00.000					
		RPT VERSION 3.1 RMF				TIME 13.00.00				CYCLE 01.000 SECONDS					

COUPLING FACILITY NAME = CX7CFP87															

SUBCHANNEL ACTIVITY															

SYSTEM NAME	# REQ TOTAL	CF TYPE	LINKS GEN	-- USE	PTH BUSY	REQUESTS			DELAYED REQUESTS						
	AVG/SEC					# REQ	-SERVICE AVG	TIME(MIC) STD_DEV	# REQ	% OF REQ	/DEL	AVG TIME(MIC) STD_DEV	/ALL		
R7D	3599K	CIB	4	4	0	SYNC	3301K	23.5	9.7	LIST/CACHE	0	0.0	0.0	0.0	0.0
	1999.5	SUBCH	28	28		ASNC	236454	84.1	139.7	LOCK	0	0.0	0.0	0.0	0.0
						CHANGED UNSUCC	0	INCLUDED IN ASYNC		TOTAL	0	0.0			
							0	0.0	0.0						
R70	436212K	ICP	4	4	3482	SYNC	435637K	5.2	3.6	LIST/CACHE	18K	0.0	285.0	215.6	0.0
	242340	SUBCH	28	28		ASNC	141411	70.9	152.5	LOCK	0	0.0	0.0	0.0	0.0
						CHANGED UNSUCC	17622	INCLUDED IN ASYNC		TOTAL	18K	0.0			
							0	0.0	0.0						

CHANNEL PATH DETAILS															

SYSTEM NAME	ID	TYPE	OPERATION MODE			DEGRADED	DISTANCE	PCHID	AID	PORT	IOP IDS				
R7D	C4	CIB	1X	IFB	HCA3-0 LR	N	<1	704	000D	01	06				
	C5	CIB	1X	IFB	HCA3-0 LR	N	<1	705	000D	01	06				
	C6	CIB	1X	IFB	HCA2-0 LR	N	<1	706	000C	02	05				
	C7	CIB	1X	IFB	HCA2-0 LR	N	<1	707	000C	02	05				

Figure 196. Coupling Facility Activity Report - Subchannel Activity

Table 146. 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.
# REQ TOTAL # REQ 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 synchronous requests from this system to the coupling facility.
# REQ ASYNC	Number of asynchronous requests from this system to the coupling facility. This number includes requests that might have started out as synchronous requests but were converted to asynchronous requests due to lack of subchannel or due to the heuristic setting.
# REQ CHANGED	Number of requests from this system that were changed from synchronous to asynchronous due to lack of subchannel. This value is a subset of # REQ ASYNC value.
# REQ UNSUCC	Number of requests which could not be completed due to hardware problems. This number should normally be zero. If it is non-zero, there is a hardware problem that needs to be investigated. The reason it is reported here is to judge to what impact extent hardware problem(s) impact coupling facility performance.
SERVICE TIME - AVG SERVICE TIME - STD_DEV	The average service time in microseconds and the standard deviation of the service time spent for requests to the coupling facility. The average service time in conjunction with its standard deviation can be used to determine potential impacts to the end user. Even though the average service time is low the standard deviation can be high indicating a wide fluctuation. This category is for the request types SYNC, ASYNC, and UNSUCC, the fields are not applicable for column CHANGED.
DELAYED REQUESTS - These columns lists possible contention reasons for requests sent to the coupling facility.	
# REQ LIST/CACHE	Number of delayed requests across all LIST and CACHE structures.
# REQ LOCK	Number of delayed requests across all LOCK structures.

Table 146. Fields in the Coupling Facility Activity Report - Subchannel Activity (continued)	
Field Heading	Meaning
# 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.

Table 147. Fields in the Coupling Facility Activity Report - Subchannel Activity - Channel Path Details	
Field Heading	Meaning
Note: If the hardware cannot provide values for a measurement, the field remains blank.	
SYSTEM NAME	The name of the system attached to the coupling facility (from IEASYSxx Parmlib member, SYSNAME parameter).
ID	The hexadecimal identifier of a channel path (CHPID) that is connected to the coupling facility.
TYPE	Channel path type.
OPERATION MODE	Channel path operation mode. It describes the data rate, bandwidth, protocol, and adapter type of the channel path. A data rate of, for example, 1GBIT denotes a rate of 1.0625 gigabit per second. A bandwidth of, for example, 12X denotes a twelve-fold bandwidth. Protocols: • IFB – InFiniBand • IFB3 – InFiniBand 3 • CEE – Converged Enhanced Ethernet • GEN3 – PCIe third generation protocol Adapter types: • HCA2-O – Host Channel Adapter2-optical • HCA2-O LR – Host Channel Adapter2-optical long reach • HCA3-O – Host Channel Adapter3-optical • HCA3-O LR – Host Channel Adapter3-optical long reach • PCIE-O SR – Peripheral Component Interconnect Express short reach • ROCE LR – RDMA over Converged Ethernet long reach Unknown operation mode: • UNKNOWN
DEGRADED	Character Y in this column indicates that the channel path is operating at reduced capacity (degraded) or not operating at all.
DISTANCE	Estimated distance in kilometers. The value is calculated as follows: $\frac{\text{Average round-trip path time in microseconds}}{10 \text{ microseconds / kilometer}}$ A value of zero means that the time was not measured.
CHID	Physical channel identifier.
AID	The hexadecimal coupling adapter identifier associated with the channel path.
PORT	The hexadecimal port associated with the channel path.
IOP IDS	The hexadecimal identifiers of I/O processors (System Assist Processors) to which the channel path is accessible.

CF to CF Activity section

COUPLING FACILITY ACTIVITY														PAGE	6
z/OS 3.1		SYSPLEX UTCPLX38				DATE 04/25/2023		INTERVAL 030.00.000							
		RPT VERSION 3.1 RMF				TIME 13.39.00		CYCLE 1.000 SECONDS							

COUPLING FACILITY NAME = X7CFP87															

CF TO CF ACTIVITY															

REQUESTS															

PEER CF	-RECEIVER- TYPE	USE	--SENDER-- TYPE	USE		# REQ	AVG/ SEC	SERVICE AVG	TIME(MIC)- STD_DEV	# REQ	% OF REQ	DELAYED REQUESTS /DEL	AVG TIME(MIC) STD_DEV	----- /ALL	
X7CFH89	CS5	3	CS5	3	SYNC	0	0.0	0.0	0.0	0	0.0	0.0	0.0	0.0	
X7CFP87	ICP	8	ICP	8	SYNC	0	0.0	0.0	0.0	0	0.0	0.0	0.0	0.0	
X7CFP89	CIB	2	CIB	2	SYNC	0	0.0	0.0	0.0	0	0.0	0.0	0.0	0.0	
0.0															

CHANNEL PATH DETAILS															

PEER CF	ID	TYPE	R/S	OPERATION MODE			DEGRADED		DISTANCE						
X7CFH89	24	CS5	R	8X	GEN3	PCIE-0 SR	N	<1							
	24	CS5	S	8X	GEN3	PCIE-0 SR	N	<1							
	25	CS5	R	8X	GEN3	PCIE-0 SR	N	<1							
	25	CS5	S	8X	GEN3	PCIE-0 SR	N	<1							
	26	CS5	R	8X	GEN3	PCIE-0 SR	N	<1							
	26	CS5	S	8X	GEN3	PCIE-0 SR	N	<1							
X7CFP89	B9	CIB	R	1X	IFB	HCA3-0 LR	N	<1							
	B9	CIB	S	1X	IFB	HCA3-0 LR	N	<1							
	BA	CIB	R	1X	IFB	HCA3-0 LR	N	<1							
	BA	CIB	S	1X	IFB	HCA3-0 LR	N	<1							

Figure 197. Coupling Facility Activity Report - CF to CF Activity

Table 148. Fields in the CF to CF Activity Section	
Field Heading	Meaning
PEER CF	Name of the remote coupling facility.
RECEIVER SENDER	TYPE CF link type of receiver/sender channel paths. USE The number of receiver/sender paths of named type used for coupling facility communication. Detail data on sender channel paths could be blank if you are using data from an old RMF gatherer or have preallocated data from a previous release of RMF.
REQUESTS	The requests are synchronous (SYNC).
# REQ	The sum of the following signals that have been sent from the subject CF to the remote CF: - Number of halt execution signals. - Number of ready to complete signals. - Number of ready to execute signals. - Number of request suppression signals. - Number of request for suppression accepted signals.
AVG/SEC	Average number of signals/messages per second.
SERVICE TIME - AVG	The average service time in microseconds for all kind of signals that have been sent from the subject CF to the remote CF, including redrives, excluding any delay time.
SERVICE TIME - STD_DEV	The standard deviation of the average service time.
DELAYED REQUESTS	The delayed requests are synchronous (SYNC).
# REQ	The number of signals of all types which have experienced a delay in being sent from the subject CF to this remote CF.
% OF REQ	The percentage of requests delayed.
AVG TIME - /DEL	The average delay time in microseconds over all delayed requests.
AVG TIME - STD_DEV	The standard deviation to the average delay time.
AVG TIME - /ALL	The average delay time in microseconds over all requests, whether delayed or not.

Note: If the hardware cannot provide values for a measurement, the field remains blank.

Table 149. Fields in the Coupling Facility Activity Report - CF to CF Activity - Channel Path Details	
Field Heading	Meaning
PEER CF	Name of the remote coupling facility.
ID	The hexadecimal identifier of a channel path (CHPID) that is connecting both coupling facilities with each other.
TYPE	Channel path type.
R/S	R Receiver channel path. S Sender channel path.
OPERATION MODE	Channel path operation mode. It describes the data rate, bandwidth, protocol, and adapter type of the channel path. For more information about displayed values, refer to Table 147 on page 338 .
DEGRADED	Character Y in this column indicates that the channel path is operating at reduced capacity (degraded) or not operating at all.
DISTANCE	Estimated distance in kilometers. For more information, refer to Table 147 on page 338 .

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 *z/OS Resource Measurement Facility User's Guide*.

Table 150. Overview names in the Coupling Facility Activity Report	
Field Heading or Meaning	Overview Name
Average service time of SYNC operations	SYNCST
SYNC operation rate	SYNCR
Average service time of ASYNC operations	ASYNCS
Ended ASYNC operation rate	ASYNCR
Percentage of changed operations	CHNGDP
Changed operation rate	CHNGDRT
Path busy rate	PBSY
Percent requests delayed due to subchannel contention	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	CWRITER
Cache castout rate	CCOUTRT
Cache cross invalidation rate	CXIRT

Table 150. Overview names in the Coupling Facility Activity Report (continued)

Field Heading or Meaning	Overview Name
Total requests to lock structure or serialized list structure	LCKREQ
Contention on lock structure	LCKCONT
False contention on lock structure	LCKFCONT
Percentage of CF utilization	STUTILP
Percentage of subchannel busy	SUBCHBP
Percentage of storage class memory in use	SCMIUP
Percentage of augmented space in use	AUGMIUP
SCM list entry current to total ratio	SCMLCTR
SCM list element current to total ratio	SCMLECTR
Average service time per SCM read operation	SCMRST
Average service time per SCM write operation	SCMWST
SCM auxiliary enabled commands to total request ratio	SCMAUXR
SCM delayed faults to total request ratio	SCMDFR

CHAN - Channel Path Activity report

The Channel Path Activity report provides information about channel path use.

The report identifies each channel path by identifier and channel path type, and reports both the total channel utilization by the central processing complex (CPC) and the channel utilization of the individual system image (partition).

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

The report includes data for each valid online channel path. Data, however, does not appear for any channel path that was offline at the end of the interval or that was brought online during the interval. Instead, one of the following messages appears in the data field:

NOW ONLINE

Brought online during the interval and still online at the end of the interval

NOW OFFLINE

Taken offline during the interval and still offline at the end of the interval

OFFLINE

Offline for the entire interval

DELETED

Deleted during the interval

MODIFIED

Modified during the interval

INSTALLED

Installed during the interval

For all channels that are managed by dynamic channel path management (DCM), additional information is available. DCM allows an installation to identify channels that they wish to be managed dynamically.

These channels are not assigned permanently to a specific control unit, but belong to a pool of channels. Based on workload requirements in the system, these channels will be assigned dynamically by DCM. On top of the report, there is a consolidated data section for managed channels displaying the total number of channel paths for each type and the average activity data. The character **M** as suffix of the acronym for the channel path type is an indicator that the channel is managed by DCM.

Duration report

Any channel that moved online or offline during the duration interval is indicated by an asterisk following the channel identifier.

In this report, the mode of the central processing complex (CPC) can be the following:

BASIC

The report shows all channels configured in the system. Only data for total utilization is reported. The partition utilization column is blank.

NOW BASIC

The report shows the last active mode. If you combine the SMF records from before and after a power-on-reset (POR) and changed the mode, two modes appear in the SMF records. By combining the intervals of the SMF records into one duration report, RMF displays the last active mode in the mode field. The partition utilization column is blank.

LPAR

The report shows the individual PR/SM logical partition's utilization and the total utilization of the shared ESCON channels, and the partition's and total utilization of the unshared channels.

NOW LPAR

The report shows the last active mode. If you combine the SMF records from before and after a POR and changed the mode, two modes appear in the SMF records. By combining the intervals of the SMF records into one duration report, RMF displays the last active mode in the mode field. The partition utilization column is blank.

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

How to request this report

Monitor I gathers data for this report automatically. If you want to suppress gathering, you need to specify NOCHAN.

To produce this report, specify

```
REPORTS(CHAN)
```

This report is also available in XML output format. [How to work with Postprocessor XML reports in z/OS Resource Measurement Facility User's Guide](#) provides all required information on how to produce and view XML reports.

Example URL for the DDS API:

```
http://ddshost:8803/gpm/rmfpp.xml?reports=CHAN
```

Contents of the report

Notes:

1. On a machine running in LPAR mode, but with only one LPAR defined, the *PART* columns for the *READ*, *WRITE* and *UTILIZATION* fields display a zero value for channels of type FC (FICON).

2. When Channel Path Measurement Facility (CPMF) is not available, for example, on z/OS systems running as z/VM or alternate VM guests, RMF uses sampled data from SRM so that the reported channel utilization is only an approximate value. With increasing channel speed, the channel utilization value becomes more and more inaccurate. Therefore, in such cases, RMF does not provide accurate values of FICON channel utilization.

Beginning with z990 processors, the channel data from SRM is no longer available. As a result, the channel utilization data on a z/OS system running as z/VM or alternate VM guest, is reported as

CHANNEL PATH ACTIVITY																	PAGE	1
z/OS 3.1			SYSTEM ID CB88			DATE 04/25/2023			INTERVAL 14.59.999									
			RPT VERSION 3.1 RMF			TIME 08.00.00			CYCLE 1.000 SECONDS									
IODF = 8E		CR-DATE: 04/25/2023			CR-TIME: 16.47.03			ACT: POR		MODE: LPAR			CPMF: EXTENDED MODE			CSSID: 2		

DETAILS FOR ALL CHANNELS																		

CHANNEL PATH				UTILIZATION(%)			READ(MB/SEC)		WRITE(MB/SEC)		FICON OPERATIONS			ZHPF OPERATIONS				
ID	TYPE	G	SPEED	SHR	PART	TOTAL	BUS	PART	TOTAL	PART	TOTAL	RATE	ACTIVE	DEFER	RATE	ACTIVE	DEFER	
00	OSD		10G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00							
02	OSE		1G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00							
04	OSD		10G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00							
05	OSD		10G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00							

74	CIB			Y	-----	-----												
75	CIB			Y	-----	-----												
92	ICP			Y	-----	-----												
93	ICP			Y	-----	-----												
C0	FC_S	13	8G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.0	0.0	0.0	0.0	0.0	0.0	
CHANNEL PATH				UTILIZATION(%)			READ(MB/SEC)		WRITE(MB/SEC)		PHYSICAL NETWORK IDS							
ID	TYPE	G	SPEED	SHR	PART	TOTAL	BUS	PART	TOTAL	PART	TOTAL	PORT 1			PORT 2			
07	OSD		1G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	NETWORK1			NETWORK1			
08	OSD		1G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	NETWORK1						
0C	OSD		1G	Y	0.00	0.00	0.09	0.00	0.00	0.00	0.00	NETWORK7C5			NETWORK7C5			
13	OSD		1G	Y	0.00	0.00	0.00	0.00	0.00	0.00	0.00	NETWORK7C2			NETWORK7C2			
15	OSD		10G	Y	0.00	0.00	0.16	0.00	0.00	0.00	0.00	NETWORK1						
...																		
CHANNEL PATH				WRITE(B/SEC)			MESSAGE RATE		MESSAGE SIZE		SEND FAIL		RECEIVE FAIL		PHYSICAL			
ID	TYPE	G	SHR	PART	TOTAL	PART	TOTAL	PART	TOTAL	PART	TOTAL	PART	PART	TOTAL			NETWORK ID	
E5	IQD			0	0	0	0	-----	-----	0	0	0	0	0			NETWORK7C5	
F1	IQD			0	0	0	0	-----	-----	0	0	0	0	0			NETWORK7C4	
F2	IQD			0	0	0	0	-----	-----	0	0	0	0	0			NETWORK7C2	

Figure 198. Channel Path Activity report

Table 151. 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.
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 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 in the z/OS Internet library (www.ibm.com/servers/resourcelink/svc00100.nsf/pages/zosInternetLibrary).
CSSID	This field is shown only for z990 processors or follow-on processors and denotes the ID of the monitored logical channel subsystem.

Table 151. Fields in the Channel Path Activity report (continued)

Field Heading	Meaning
CHANNEL GROUP G NO	For each channel type which is managed by DCM, a summary line is shown with the average values for all channels in this group. G indicates the generation and is used to differentiate between channels of the same channel type, when one has significant differences from the other. Newer generations with significant differences (for example, the channel throughput) are indicated by a number (1, 2, ...). For example, for a FICON channel, a number 1 indicates that the channel has an auto-negotiated throughput of 1 Gbit/sec, or a number 2 indicates a throughput of 2 Gbit/sec. The number of channels of the group is given in column NO.
CHANNEL PATH ID	The hexadecimal channel path identifier (CHPID).
CHANNEL PATH TYPE	Type of channel path. You may issue the console command <code>D M=CHP(xx)</code> to see an explanation of the channel path type. If RMF encounters an error while processing the TYPE data, this field is blank. RMF continues to measure channel path activity. Check the operator console for messages.
CHANNEL PATH G	This column indicates the generation and is used to differentiate between channels of the same channel type, when one has significant differences from the other. Newer generations with significant differences are indicated by a number (1, 2, ...). For example, for z/OS, a number 2 indicates that a FICON channel has auto negotiated to a link speed of 2 GB/sec.
CHANNEL PATH SPEED	The channel path speed in bits per second at the end of the interval.
CHANNEL PATH SHR	The indication of whether the channel path is defined as shared between one or more logical partitions. Y indicates that the channel path is shared. A blank indicates it is not shared.
PHYSICAL NETWORK ID(S)	Physical-network identifiers (PNET IDs) of an Ethernet network that is accessible from the ports of the channel path.
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: $\text{PART UTILIZATION (\%)} = \frac{\text{Channel Path Busy Time}}{\text{Channel Path Elapsed Time}} * 100$ For channels like FICON, OSA Express, or OSA Direct Express, which are running in extended CPMF mode, the calculation is as follows: $\text{Part Utilization (\%)} = \frac{\text{LPAR \# of Channel Work Units}}{\text{Max \# of Channel Work Units} * \text{Channel Path Elapsed Time}} * 100$ For some channels like OSAEGbE, FICON EXPRESS/EXPRESS2, this value reflects the microprocessor utilization. For hipersockets, this value is not available.

Table 151. Fields in the Channel Path Activity report (continued)

Field Heading	Meaning
UTILIZATION (%) TOTAL	The channel path utilization percentage for the CPC 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: $\text{Total Utilization (\%)} = \frac{\text{\# SRM Observations of Channel Path Busy}}{\text{\# Samples}} * 100$ For unshared channels in LPAR mode, the value for total utilization is the same as partition utilization. For channels like FICON, OSA Express, or OSA Direct Express, which are running in extended CPMF mode, the calculation is as follows: $\text{Total Utilization (\%)} = \frac{\text{Total \# of Channel Work Units}}{\text{Max \# of Channel Work Units} * \text{Channel Path Elapsed Time}} * 100$ For some channels like OSAEGbE, FICON EXPRESS/EXPRESS2, this value reflects the microprocessor utilization. For hipersockets, this value is not available.
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. For hipersockets, this value is not available.
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 CPC. For hipersockets, this value is not available.
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 CPC.
FICON OPERATIONS	RATE Number of native FICON operations per second. ACTIVE The average number of native FICON operations that are concurrently active during the reporting interval. DEFER Number of deferred native FICON operations per second that could not be initiated by the channel due to the lack of available resources. This field is reported for the CPC.
ZHPF OPERATIONS	RATE Number of ZHPF (High Performance FICON) operations per second. ACTIVE The average number of ZHPF operations that are concurrently active during the reporting interval. DEFER Number of deferred ZHPF operations per second that could not be initiated by the channel due to the lack of available resources. This field is reported for the CPC.

Table 151. Fields in the Channel Path Activity report (continued)	
Field Heading	Meaning
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 CPC. The values are shown in bytes/second. This field is for HiperSockets.
MESSAGE RATE	PART Rate of messages sent by this partition. TOTAL Rate of messages sent by the CPC. This field is for HiperSockets.
MESSAGE SIZE	PART Average size of messages sent by this partition. TOTAL Average size of messages sent by the CPC. This field is for HiperSockets.
SEND FAIL PART	Rate of messages (sent by this partition) that failed. This field is for HiperSockets.
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 CPC) that failed due to unavailable buffers. This field is for HiperSockets.

Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*. The following table shows the overview condition names for the Overview report.

Table 152. Overview names in the Channel Path Activity report	
Field Heading or Meaning	Overview Name
Use the following overview condition if CPMF is not available or for CPMF compatibility mode:	
TOTAL UTILIZATION (%)	CHPB SY, CHGPB SY
Use the following overview conditions for CPMF extended mode:	
UTILIZATION (%) PART	CHLB SY, CHGLB SY
UTILIZATION (%) TOTAL	CHTB SY, CHGTB SY
UTILIZATION (%) BUS	CHBTOT, CHGBTOT
PART READ RATE	CHLREAD, CHGLREAD
TOTAL READ RATE	CHTREAD, CHGTREAD
PART WRITE RATE	CHLWRITE, CHGLWRITE
TOTAL WRITE RATE	CHTWRITE, CHGTWRITE
FICON OPERATIONS RATE	CHFRATE
FICON OPERATIONS ACTIVE	CHFACTV
FICON OPERATIONS DEFER	CHFD FER

Table 152. Overview names in the Channel Path Activity report (continued)

Field Heading or Meaning	Overview Name
ZHPF OPERATIONS RATE	CHFXRATE
ZHPF OPERATIONS ACTIVE	CHFXACTV
ZHPF OPERATIONS DEFER	CHFXDFER
MESSAGE RATE PART	CHLMSGST
MESSAGE RATE TOTAL	CHTMSGST
MESSAGE SIZE PART	CHLMSGSZ
MESSAGE SIZE TOTAL	CHTMSGSZ
SEND FAIL PART	CHLMSGF
RECEIVE FAIL PART	CHLRECF
RECEIVE FAIL TOTAL	CHTRECF

CPU - CPU Activity report

The CPU report has several sections that display information about the active processors, address spaces and running or waiting work units, blocked workloads, all configured partitions, data about each LPAR cluster, capacity limit of each defined capacity group, and power consumption.

The report is divided into the following sections:

CPU Activity

Provides information on the active processors. See [“CPU Activity” on page 348](#).

System Address Space and Work Unit Analysis

Provides overall information about address spaces and running or waiting work units. See [“System Address Space and Work Unit Analysis” on page 350](#).

Blocked Workload Analysis

Provides information about blocked workloads. See [“Blocked Workload Analysis” on page 350](#).

Partition Data Report

If the z/OS system is running in a PR/SM environment in LPAR mode, this section provides data about all configured partitions. If the z/OS system is running as guest under z/VM or alternate VM guest, and the Monitor I data gatherer option VMGUEST has been set, this section provides data about the z/OS guest system. Otherwise, this section is not available. See [“Using the information in the Partition Data Report” on page 357](#).

LPAR Cluster Report

Provides data about each LPAR cluster. See [“Using the information in the LPAR Cluster Report” on page 363](#).

Group Capacity Report

Provides data about the capacity limit of each defined capacity group and about the MSU consumption and actual capping of each partition within these groups. See [“Using the information in the Group Capacity Report” on page 365](#).

Hardware Group Report

The Hardware Group Report displays the settings of the hardware groups and their partitions. See [“Using the information in the Hardware Group Report” on page 366](#).

Power Consumption Report

The Power consumption reports are essential for managing energy efficiency and operational costs. This report provides insights into the power usage patterns of the system over defined intervals, allowing you to monitor and optimize performance. See [“Using the information in the Power Consumption report” on page 367](#).

Remember: The **LPAR Cluster Report** and **Group Capacity Report** sections are not available if the system is running in a z/VM or alternate VM guest environment.

How to request this report

Monitor I gathers data for this report automatically. If you want to suppress gathering, you need to specify NOCPU.

To produce this report, specify

```
REPORTS(CPU)
```

This report is also available in XML output format. [How to work with Postprocessor XML reports in z/OS Resource Measurement Facility User's Guide](#) provides all required information on how to produce and view XML reports.

Example URL for the DDS API:

```
http://ddshost:8803/gpm/rmfpp.xml?reports=CPU
```

Contents of the report

The contents of the CPU Activity Report includes the following parts:

- “CPU Activity” on page 348
- “System Address Space and Work Unit Analysis” on page 350
- “Blocked Workload Analysis” on page 350

CPU Activity

The CPU Activity section reports on logical core and logical processor activity. For each processor, the report provides a set of calculations that are provided at a particular granularity that depends on whether multithreading is disabled (LOADxx PROCVIEW CPU parameter is in effect) or enabled (LOADxx PROCVIEW CORE parameter is in effect).

If multithreading is disabled for a processor type, all calculations are at logical processor granularity.

If multithreading is enabled for a processor type, some calculations are provided at logical core granularity and some are provided at logical processor (thread) granularity. The CPU Activity section displays exactly one report line per thread showing all calculations at logical processor granularity. Those calculations that are provided at core granularity are only shown in the same report line that shows the core id in the CPU NUM field and which is representing the first thread of a core.

The following calculations are on a per logical processor basis when multithreading is disabled and on a per logical core basis when multithreading is enabled:

- Percentage of the interval time the processor was online
- LPAR view of the processor utilization (LPAR Busy time percentage)
- Percentage of a physical processor the logical processor is entitled to use
- Multithreading core productivity (only reported when multithreading is enabled)
- Multithreading core utilization (only reported when multithreading is enabled)

The following calculations are on a per logical processor basis regardless whether multithreading is enabled or disabled:

- MVS view of the processor utilization (MVS Busy time percentage)
- Percentage of the online time the processor was parked (in HiperDispatch mode only)
- I/O interrupts rate (general purpose processors only)
- Percentage of I/O interrupts handled by the I/O supervisor without re-enabling (general purpose processors only)

If RMF is running as a guest under z/VM® or alternate VM and Monitor I Session option NOVMGUEST is active, it only reports the MVS busy time percentage. If you want to measure partition utilization (as well

as the individual CPU utilization of the single guests, namely LPAR busy time percentage), you need to use a z/VM or alternate VM monitor. Performance analysts need both views of CPU utilization. The MVS view is a direct indicator to see a CPU bottleneck, while the LPAR view is important with respect to capacity aspects.

The **LPAR view** of the CPU utilization takes the different states that are possible into account:

- WAIT state
- NON WAIT state being dispatched by PR/SM
- NON WAIT state not being dispatched by PR/SM
- WAIT state being dispatched when the LPAR has dedicated processors

The LPAR Busy time is calculated depending on the status of the logical processor:

Dedicated and LOADxx PROCVIEW CPU is in effect or hardware does not support multithreading

CPU time = Online time - Wait time

Dedicated and LOADxx PROCVIEW CORE is in effect on hardware that supports multithreading

CPU time = MT Core LPAR Busy time

Wait completion = YES (requires multithreading disabled)

CPU time = Dispatch time - Wait time

Wait completion = NO

CPU time = Dispatch time

The LPAR view of CPU utilization is calculated as:

$$\text{LPAR Busy Time(\%)} = \frac{\text{CPU time}}{\text{Online time}} * 100$$

The MVS view of the CPU utilization considers the following states:

- CPU wait state
- CPU busy state (which means NON WAIT state)

In HiperDispatch mode, logical processors can be parked and are not dispatched by z/OS. The MVS BUSY fields in the RMF report reflect the effective used capacity for the logical processors and the entire logical partition. The values are based on the difference between online time and MVS wait time to provide an operating system perspective of busy time. Parked processors in HiperDispatch mode generally reflect unavailable capacity at high physical processor utilizations. The formula for MVS Busy has been changed with HiperDispatch mode to exclude the parked time to show how busy the logical processor was when not parked.

HiperDispatch = NO

Time range = Online time

HiperDispatch = YES

Time range = Online time - Parked time

Note: In HiperDispatch mode, the Total/Average MVS BUSY TIME % does not consider parked processors. Therefore, do not use Total/Average LPAR BUSY TIME % and Total/Average MVS BUSY TIME % to calculate the MVS to LPAR busy ratio.

The **MVS view** of CPU utilization is:

$$\text{MVS Busy Time(\%)} = \frac{\text{Time range} - \text{Wait time}}{\text{Time range}} * 100$$

If multithreading is enabled for at least one processor type, you can use the multithreading core productivity and multithreading core utilization metrics to determine the effectiveness of the configured logical cores.

When the multithreading core productivity (MT % PROD) equals 100% in multithreading mode, all threads on the core are executing work and all core resources are being used. If MT % PROD is less than 100%, the core resources were dispatched to physical hardware but one or more threads on a core were in a wait because they had no work to run.

If multithreading is enabled, the available core capacity can be calculated using the multithreading core utilization and LOG PROC SHARE %:

Available Core Capacity = LOG PROC SHARE % - MT % UTIL

System Address Space and Work Unit Analysis

The **System Address Space and Work Unit Analysis** section of the CPU activity report provides overall address space and work unit information and also provides the minimum, maximum, and average numbers of running or ready to run work units.

The data in this section analyzes the following types of address spaces:

- In storage and ready to execute
- In storage
- Out of storage and ready to execute
- Out of storage and waiting to execute
- Logically out of storage and ready to execute
- Logically out of storage and waiting to execute

Data is also presented on the number of address spaces used by batch users, started tasks (STC), TSO/E users, APPC/MVS transaction schedulers (ASCH), and z/OS UNIX (OMVS). Examining this data can indicate when a backlog of address spaces are waiting to use the processor.

The work unit statistics (MIN, MAX, AVG) are provided per processor type, that is, per standard CPs, zAAPs, zIIPs and zCBPs. The distribution does not distinguish between the processor types (CPs, zAAPs, zIIPs and zCBPs.).

The graphical and numeric presentation of the In-Ready work unit queue distribution provides a detailed view on how many work units are running or waiting for a processor. The distribution does not distinguish between the processor types (CPs, zAAPs, and zIIPs).

Blocked Workload Analysis

If the CPU utilization of a system is at 100%, workloads with low importance (low dispatch priority) might not get dispatched anymore. This could cause problems if the work holds a resource and by that holds up more important workloads. Therefore, any address space or enclave which has ready-to-run work units (TCBs or SRBs), but does not get CPU service within a certain time interval due to its low dispatch priority, will be temporarily promoted by WLM to a higher dispatch priority. This helps to complete low priority work in a finite time period, without permanently delaying high priority work.

The **Blocked Workload Analysis** section lists the number of dispatchable work units that are considered to be blocked and eligible for priority promotion. This section also displays the OPT parameters which define the workload promotion. It also displays the average exploitation of the defined promotion rate during the measurement interval. This information helps you to adjust these OPT parameters. To assess the amount of workload still being blocked, the average and peak number of address spaces and enclaves found blocked and waiting for promotion is also listed.

Using the information in the CPU Activity report

High LPAR/MVS BUSY TIME PERC values could indicate contention for CPU. To check this, add the N+1, ... N+150 percentages in the DISTRIBUTION OF IN-READY WORK UNIT QUEUE (where N is the number of online processors). This sum is the percentage of time when at least one task could not be dispatched. A value higher than 60% implies contention for CPU.

Low LPAR/MVS BUSY TIME PERC values can indicate that other bottlenecks in the system are preventing work from being processed.

An OUT READY average value of more than 1 could reflect processor storage constraints.

```

                                C P U   A C T I V I T Y
                                -----
                                z/OS 3.1          SYSTEM ID SYSF          DATE 09/30/2021          INTERVAL 14.59.999
                                RPT VERSION 3.1 RMF          TIME 12.14.33          CYCLE 1.000 SECONDS
CPU 3906 CPC CAPACITY 9163 SEQUENCE CODE 000000000000C5327
MODEL 776 HIPERDISPATCH=YES
H/W MODEL M03 CHANGE REASON=NONE BOOST TYPE=ZIIP BOOST CLASS=IPL

---CPU---
NUM TYPE ONLINE LPAR BUSY MVS BUSY PARKED LOG PROC SHARE % RATE % VIA TPI
0 CP 100.00 2.84 2.83 0.00 100.0 HIGH 0.00 0.00
1 CP 100.00 0.80 0.79 0.00 100.0 HIGH 0.00 0.00
2 CP 100.00 1.04 1.03 0.00 100.0 HIGH 59.60 5.82
TOTAL/AVERAGE 1.56 1.55 300.0 59.60 5.82
3 IIP 100.00 1.28 1.28 0.00 48.5 MED 0.00 0.00
4 IIP 100.00 0.07 0.07 0.00 0.0 LOW 0.00 0.00
6 IIP 100.00 0.00 0.00 100.00 0.0 LOW 0.00 0.00
TOTAL/AVERAGE 0.45 0.67 48.5

                                C P U   A C T I V I T Y
                                -----
                                z/OS 3.1          SYSTEM ID SYSF          DATE 04/25/2023          INTERVAL 14.59.999
                                RPT VERSION 3.1 RMF          TIME 12.14.33          CYCLE 1.000 SECONDS
SYSTEM ADDRESS SPACE AND WORK UNIT ANALYSIS
-----NUMBER OF ADDRESS SPACES-----
QUEUE TYPES MIN MAX AVG NUMBER OF WORK UNITS (%)
IN 73 77 73.8 <= N 100.0
IN READY 0 2 0.0 = N + 1 0.0
OUT READY 0 0 0.0 = N + 2 0.0
OUT WAIT 0 0 0.0 = N + 3 0.0
LOGICAL OUT RDY 0 0 0.0 <= N + 5 0.0
LOGICAL OUT WAIT 53 58 56.5 <= N + 10 0.0
ADDRESS SPACE TYPES <= N + 15 0.0
BATCH 0 0 0.0 <= N + 20 0.0
STC 117 120 119.6 <= N + 30 0.0
TSO 0 1 0.7 <= N + 40 0.0
ASCH 0 0 0.0 <= N + 60 0.0
OMVS 10 10 10.0 <= N + 80 0.0
<= N + 100 0.0
<= N + 120 0.0
<= N + 150 0.0
> N + 150 0.0
-----NUMBER OF WORK UNITS-----
CPU TYPES MIN MAX AVG N = NUMBER OF PROCESSORS ONLINE UNPARKED (5.0 ON AVG)
CP 0 4 0.0
IIP 0 1 0.0

BLOCKED WORKLOAD ANALYSIS
OPT PARAMETERS: BLWLTRPCT (%) 0.5 PROMOTE RATE: DEFINED 31 WAITERS FOR PROMOTE: AVG 0.000
BLWLINTHD 20 USED (%) 0 PEAK 0

```

Figure 199. CPU Activity Report

Table 153. Fields in the CPU Activity Report	
Field Heading	Meaning
CPU Activity	If multithreading is disabled for the reporting interval (that is, the LOADxx PROCVIEW CPU parameter is in effect), the term logical core refers to a logical processor.
CPU MODEL H/W MODEL CBP MODEL SEQUENCE CODE	The hardware description of the CPU.
CPC CAPACITY	Effective processor capacity available to the central processor complex (CPC), measured in MSU/h.
CBP CAPACITY	Effective zCBP processor capacity available to the central processor complex (CPC), measured in MSU/h.

Table 153. Fields in the CPU Activity Report (continued)

Field Heading	Meaning
CHANGE REASON	Reason of the capacity change: NONE CPC is running at normal capacity (100 % effective capacity). No capacity change occurred. POWERSAVE CPC is running in power-save mode. Capacity change was initiated by the user. MACHINE CPC is running in cycle-steering mode. Capacity change was initiated by the machine. N/A No capacity change reason is reported by the machine.
HIPERDISPATCH	HiperDispatch mode: YES Active NO Not active N/A Not supported by the hardware If the mode changed during the reporting interval, an '*' is appended (for example: NO* indicates a switch from YES to NO).
BOOST TYPE	The boost type that was active at the end of the interval: NONE Boost is inactive. ZIIP zIIP boost. SPEED Speed boost. ALL zIIP and speed boost are both active.
BOOST CLASS	The boost class (or period) within which boost was active: NONE Boost is inactive. IPL IPL (Startup) boost. SHUTDOWN Shutdown boost. RECOVERY Recovery Process boost.
CPU NUM/TYPE	The logical core identification and the processor type.
TIME % ONLINE	The percentage of time the logical core was online.

Table 153. Fields in the CPU Activity Report (continued)

Field Heading	Meaning
TIME % LPAR BUSY	The percentage of the online time that the logical core was dispatched by LPAR. - For a dedicated partition: When LOADxx PROCVIEW CORE is in effect on hardware that supports multithreading: $\text{LPAR BUSY TIME \%} = \frac{\text{MT Core LPAR Busy Time}}{\text{Online Time}} * 100$ Otherwise: $\text{LPAR BUSY TIME \%} = \frac{\text{Online Time} - \text{Wait Time}}{\text{Online Time}} * 100$ - For a non-dedicated partition when Wait Completion is NO: $\text{LPAR BUSY TIME \%} = \frac{\text{Partition Dispatch Time}}{\text{Online Time}} * 100$ The partition dispatch time is the elapsed time that PR/SM dispatched this logical core during the interval. - For a non-dedicated partition when Wait Completion is YES: $\text{LPAR BUSY TIME \%} = \frac{\text{Partition Dispatch Time} - \text{Wait Time}}{\text{Online Time}} * 100$
TIME % MVS BUSY	The percentage of the online time that the logical processor was busy. $\text{MVS BUSY TIME \%} = \frac{\text{Online Time} - (\text{Wait Time} + \text{Parked Time})}{\text{Online Time} - \text{Parked Time}} * 100$ The MVS view of CPU time is not meaningful if the logical processor is parked during the entire reporting interval. In this case, '----' is shown.
TIME % PARKED	The percentage of time that the logical processor was parked. In HiperDispatch mode, processors with a low amount of physical processor share can be parked. That is, they are not dispatched by z/OS and do not attempt to run work. Without HiperDispatch, processors are not parked and '----' is shown.
MT % PROD	The percentage of the maximum core capacity that was used in the reporting interval while the logical core was dispatched to physical hardware. When MT % PROD equals 100% and the LOADxx PROCVIEW CORE parameter is in effect, all threads on the core are executing work and all core resources are being used. If MT % PROD is less than 100%, the core resources were dispatched to physical hardware but one or more threads on a logical core were in a wait because they had no work to run. If a core was reconfigured offline/online during the reporting interval, no multithreading core productivity is calculated and '-----' is shown. If the LOADxx PROCVIEW CPU parameter is in effect, this field is not displayed.
MT % UTIL	The percentage of the maximum core capacity that was used in the reporting interval. $\text{MT \% UTIL} = \text{MT Core Productivity} * \text{TIME \% LPAR BUSY}$ If a core was reconfigured offline/online during the reporting interval, no multithreading core utilization is calculated and '-----' is shown. If the LOADxx PROCVIEW CPU parameter is in effect, this field is not displayed.

Table 153. Fields in the CPU Activity Report (continued)

Field Heading	Meaning
LOG PROC SHARE %	Percentage of the physical processor that the logical processor is entitled to use. Without HiperDispatch, the processing weight is equally divided between the online logical processors. In HiperDispatch mode, logical processors have a high, medium or low share of the physical processor. The share percentage is the average value for the reporting interval, whereas HIGH, MED or LOW indicates the HiperDispatch priority at the end of the reporting interval. When the priority changed during the interval, an '*' is appended. N/A is displayed if the HiperDispatch priority is not indicated by the hardware at the end of the reporting interval.
I/O INTERRUPTS RATE	The total rate per second that this processor handled I/O interrupts. The rate reflects the processing for the entire interval. This might include periods of time when the SRM enabled or disabled this processor for I/O interrupts. The rate includes interrupts handled by the second level interrupt handler (SLIH), as well as those handled by the Test Pending Interrupt (TPI) instruction. $\text{RATE} = \frac{\text{SLIH} + \text{TPI}}{\text{INT}}$ SLIH Interrupts that the second level interrupt handler handled TPI Interrupts that the Test Pending Interrupt instruction handled INT Interval time (seconds)
I/O INTERRUPTS % 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. $\% \text{ VIA TPI} = \frac{\text{TPI}}{\text{SLIH} + \text{TPI}} * 100$ TPI Interrupts that the Test Pending Interrupt instruction handled SLIH Interrupts that the second level interrupt handler handled
For the following four TOTAL/AVERAGE values, the logical processors that are parked during the entire interval are not considered in the calculation of the average TIME % MVS BUSY.	
TOTAL/AVERAGE (CP)	The average or total value for general purpose processors (standard CPs).
TOTAL/AVERAGE (zAAP)	The average value for zAAPs. Only visible if zAAPs are configured online.
TOTAL/AVERAGE (zCBP)	The average value for zCBP processors. Only visible if zCBP processors are configured online.
TOTAL/AVERAGE (zIIP)	The average value for zIIPs. Only visible if zIIPs are configured online.
Multi-Threading Analysis: This information is only displayed when the LOADxx PROCVIEW CORE parameter is in effect. Multithreading information is only shown for those processor types for which at least one logical core was configured online for the complete interval.	
CPU TYPE	Processor type CP, AAP, CBP or IIP.
MODE	The multithreading mode of a processor type designates the number of active threads for each online logical core of this type. If MT MODE is greater than 1, multithreading becomes effective for this processor type.
MAX CF	Multithreading maximum capacity factor for a processor type. The multithreading maximum capacity factor represents the ratio of the maximum amount of work that can be accomplished using all active threads to the amount of work that would have been accomplished within this reporting interval when multithreading was disabled. '-----' is shown when the multithreading maximum capacity factor cannot be calculated.

Table 153. Fields in the CPU Activity Report (continued)

Field Heading	Meaning
CF	Multithreading capacity factor for a processor type. The multithreading capacity factor represents the ratio of the amount of work that has been accomplished within this reporting interval to the amount of work that would have been accomplished with multithreading disabled. '-----' is shown when the multithreading capacity factor cannot be calculated.
AVG TD	Average thread density for a processor type. This value represents the average number of active threads for those cores that were dispatched to physical hardware. '-----' is shown when the average thread density cannot be calculated.
System Address Space and Work Unit Analysis: contains information about the NUMBER OF ADDRESS SPACES categorized by the QUEUE TYPES, in which they have been waiting, and categorized by the ADDRESS SPACE TYPES. Furthermore, the MIN, MAX, and AVG numbers of work units are categorized by the CPU TYPES for which they have been dispatched (that is, for standard CPs, zAAPs, zCBPs and zIIPs). This section also shows how many work units have been waiting in the IN-READY queue (DISTRIBUTION OF IN-READY WORK UNIT QUEUE).	
QUEUE TYPES	Shows the number of address spaces that are waiting in the different queues. For each queue type, the MIN, MAX and AVG numbers of address spaces are displayed. The following queue types are analyzed: IN READY Address spaces that are in central storage and ready to execute or currently in execution. 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.
ADDRESS SPACE TYPES	Shows the total number of address spaces detected during the report interval, categorized by address space types. For each address space type, the MIN, MAX and AVG numbers of active address spaces are displayed. The following ADDRESS SPACE TYPES are analyzed: 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.

Table 153. Fields in the CPU Activity Report (continued)	
Field Heading	Meaning
DISTRIBUTION OF IN-READY WORK UNIT QUEUE	The percentaged and graphical distribution of SRM samples when the number of work units on the IN-READY queue is within a certain range. The correlation is based on N, which is the number of online logical processors when the sample is taken. In HiperDispatch mode, N is the number of online logical processors that are not parked. For example, NUMBER OF WORK UNITS = N + 10 with a percentage of 4.3 (see Figure 199 on page 351) indicates that in 4.3 % of the samples ten work units were waiting for a logical processor.
NUMBER OF WORK UNITS by CPU type	The minimum, maximum and average numbers of running and waiting work units categorized by CPU type (standard CPs, zAAPs, zCBPs and zIIPs).
Blocked Workload Analysis: provides information about blocked address spaces and enclaves.	
OPT PARAMETERS	Lists the OPT parameters which define the workload promotion: BLWLTRPCT (%) Specifies how much of the CPU capacity is to be used to promote blocked workloads. This parameter does not influence the amount of CPU service that a single blocked address space or enclave is given. Instead, this parameter influences how many different address spaces or enclaves can be promoted at the same point in time. If the value specified with this parameter is not large enough, blocked workloads might need to wait longer than the time interval defined by BLWLINTHD. This value is specified as a number between 0 and 200 where 200 accounts for 20.0%. BLWLINTHD Specifies the threshold time interval in seconds for which a swapped-in address space or enclave must wait before being considered to be blocked and eligible for promotion. If the parameters have been changed during the reporting interval, the values are followed by an '*'.
PROMOTE RATE	DEFINED Number of blocked dispatchable work units which may get promoted in their dispatching priority per second. This value is derived from OPT parameter BLWLTRPCT. USED (%) The utilization of the defined promote rate during the reporting interval.
WAITERS FOR PROMOTE	Number of address spaces and enclaves found blocked according to OPT parameter BLWLINTHD: AVG Average number found blocked during the report interval. PEAK Highest number found blocked during the report interval.

Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*. The following table shows the overview condition names for the Overview report.

Table 154. Overview names in the CPU Activity Report	
Field Heading or Meaning	Overview Name
CPC CAPACITY	NOMCAPAC, EFFCAPAC
CBP CAPACITY	NOMCACBP, EFFCACBP
ONLINE TIME PERC for general purpose processors	CONTPER
LPAR BUSY TIME PERC	CPUBSY (LPAR mode only)
LPAR BUSY TIME PERC for zAAPs	AAPBSY
LPAR BUSY TIME PERC for zCBPs	CBPBSY
LPAR BUSY TIME PERC for zIIPs	IIPBSY

Table 154. Overview names in the CPU Activity Report (continued)

Field Heading or Meaning	Overview Name
MVS BUSY TIME PERC	MVSBSY, CPUBSY
MVS BUSY TIME PERC for zAAPs	AAPMBSY
MVS BUSY TIME PERC for zCBPs	CBPMBSY
MVS BUSY TIME PERC for zIIPs	IIPMBSY
TYPE (IN READY)	AVGIARDY
TYPE (other)	MXBATCH, AVGBATCH, MXSTC, AVGSTC, MXTSO, AVGTSO, MXASCH, AVGASCH, MXOMVS, AVGOMVS, AVGOARDY, AVGUIN, AVGUOWT, AVGULRDY, AVGULWT
Number of general purpose processors online	NUMPROC
Number of zAAPs online	NUMAAP
Number of zCBPs online	NUMCBP
Number of zIIPs online	NUMIIP
Percentage of the report interval during which at least <i>n</i> jobs could not be dispatched (with <i>n</i> =1,2,3,4,5,10,15,20,30,40,60,80)	OCPU1, OCPU2, OCPU3, OCPU4, OCPU5, OCPU10, OCPU15, OCPU20, OCPU30, OCPU40, OCPU60, OCPU80
Number of CPs/zAAPs/zIIPs/zCBPs with high/medium/low HiperDispatch share for the partition	HDCPHIGH, HDAPHIGH, HDCBHIGH, HDIPHIGH, HDCPMED, HDAPMED, HDCBMED, HDIPMED, HDCPLOW, HDCBLOW, HDAPLOW, HDIPLow
Percentage of time that the general purpose processor was parked	CPARKPER
Maximum number of in-ready work units for general purpose processors	MXWUCP
Maximum number of in-ready work units for zAAPs	MXWUAAP
Maximum number of in-ready work units for zCBPs	MXWUCBP
Maximum number of in-ready work units for zIIPs	MXWUIIP
Average number of in-ready work units for general purpose processors	AVGWUCP
Average number of in-ready work units for zAAPs	AVGWUAAP
Average number of in-ready work units for zCBPs	AVGWUCBP
Average number of in-ready work units for zIIPs	AVGWUIIP
Percentage of the report interval during which at least <i>n</i> work units could not be dispatched (with <i>n</i> =1,2,3,4,5,10,15,20,30,40,60,80,100,120,150)	WCPU1, WCPU2, WCPU3, WCPU4, WCPU5, WCPU10, WCPU15, WCPU20, WCPU30, WCPU40, WCPU60, WCPU80, WCPU100, WCPU120, WCPU150
Percent multithreading core productivity for general purpose processors	MTPROD
Percent multithreading core productivity for zIIPs	IIPPROD
Percent multithreading core utilization for general purpose processors	MTUTIL
Percent multithreading core utilization for zIIPs	IIPUTIL

Using the information in the Partition Data Report

When RMF is running in a Processor Resource/Systems Manager (PR/SM) environment in LPAR mode, the **Partition Data Report** section of the **CPU Activity** report provides data about all configured partitions active at the end of the reporting interval, independent of the operating system running in each partition.

When RMF is running on a z/OS guest in a z/VM or alternate VM guest environment, and the Monitor I data gatherer option VMGUEST has been set when the SMF record was collected, then the report section provides data about the z/OS guest system. If you want information about another z/OS guest system, you can run RMF separately on that guest system.

The report contains the following information:

- Header information
- Partition data
- Logical partition processor data
- Average processor utilization percentages

The header information gives an overview of the LPAR mode characteristics:

- MVS partition name
- Image capacity — information related to software pricing
- Number of configured partitions
- Number of physical processors in total and per type
- Wait completion indicator
- Dispatch interval
- If a group of LPARs on the same CEC is managed towards a combined capacity limit, the name of the group and the common capacity limit is displayed
- Capping information

The section PARTITION DATA is grouped by general purpose and special purpose processor types and provides the following information:

- Name
- Status
- Boost type
- Weighting share of resources
- Defined and consumed service units
- Capping information

The section LOGICAL PARTITION PROCESSOR DATA provides the following information about the partition's processors:

- Number and type of processors assigned to this partition
- The partition's effective dispatch time
- The partition's total dispatch time

The section AVERAGE PROCESSOR UTILIZATION PERCENTAGES provides the following information about the partition's processors:

- Logical constraint percentages. If multithreading is enabled, the percentages shown for logical processor resources can be applied to logical core resources.
 - The partition's average effective utilization of the logical processor resource
 - The partition's average total utilization of the logical processor resource
- Physical constraint percentages. If multithreading is enabled, the percentages shown for physical processor resources can be applied to physical core resources.
 - The average LPAR Management utilization of the physical processor resource on behalf of the partition
 - The partition's average effective utilization of the physical processor resource
 - The partition's average total utilization of the physical processor resource

PARTITION DATA REPORT																PAGE	3																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																														
z/OS 3.1				SYSTEM ID SYSF				DATE 04/25/2023				INTERVAL 14.59.999																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																			
				RPT VERSION 3.1 RMF				TIME 12.14.33				CYCLE 1.000 SECONDS																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																			
MVS PARTITION NAME				SYSF				PHYS PROC NUM				105				GROUP NAME		SYSDPLEX		INITIAL CAP		NO																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																									
IMAGE CAPACITY				482				CP				76				LIMIT		15000*		LPAR HW CAP		NO																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																									
NUMBER OF CONFIGURED PARTITIONS				41				AAP				3				AVAILABLE		9087		HW GROUP CAP		NO																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																									
WAIT COMPLETION				NO				IFL				20								ABS MSU CAP		NO																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																									
DISPATCH INTERVAL				DYNAMIC				ICF				2																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																			
				IIP				4																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																							
----- PARTITION DATA -----																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																															
-----MSU-----								-----CAPPING-----								-----LOGICAL PARTITION PROCESSOR DATA-----								-----AVERAGE PROCESSOR UTILIZATION PERCENTAGES-----																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																							
NAME								S BT WGT								DEF								ACT								LOGICAL PROCESSORS								PHYSICAL PROCESSORS																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																							
SYSF								A N 200								500								5								N N N								WLM%								NUM								TYPE								EFFECTIVE								TOTAL								EFFECTIVE								TOTAL								LPAR								MGMT								EFFECTIVE								TOTAL																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																							
COHCF0E								A N 200								0								7								N N N								0.0								1								CP								00.00.54.300								00.00.54.526								6.03								6.06								0.00								0.08								0.08																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																															
COH2								A N 350								0								8								N N N								0.0								4								CP								00.00.57.182								00.00.59.315								1.59								1.65								0.00								0.08								0.09																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																															
COH3								A N 70								0								5								N N N								0.0								4								CP								00.00.39.000								00.00.40.554								1.08								1.13								0.00								0.06								0.06																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																															
IRD1								A N 50								0								4								N N N								0.0								6.0								CP								00.00.30.358								00.00.32.020								0.56								0.59								0.00								0.04								0.05																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																															
IRD1CFB								A N 200								0								2								Y N N								0.0								1								CP								00.00.11.737								00.00.11.900								1.30								1.32								0.00								0.02								0.02																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																															
IRD2								A N 50								0								3								N N N								0.0								8.0								CP								00.00.21.835								00.00.23.289								0.30								0.32								0.00								0.03								0.03																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																															
IRD3								A N 50								0								4								N N N								0.0								8.0								CP								00.00.27.889								00.00.29.510								0.39								0.41								0.00								0.04								0.04																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																															
SYSD								A N 50								0								329								N N N								0.0								3								CP								00.40.54.118								00.40.55.177								90.89								90.93								0.00								3.59								3.59																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																															

VMSSL								A N 20								0								565								N N N								0.0								6								CP								01.10.19.994								01.10.30.346								78.15								78.34								0.02								6.17								6.18																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																															
ZVM								A N 150								0								0								N N N								0.0								8								CP								00.00.00.000								00.00.00.000								0.00								0.00								0.00								0.00								0.00																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																															
PHYSICAL																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																															
TOTAL								3480																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																																							

Figure 200. CPU Activity Report - Partition Data Report

Note: An asterisk (*) next to any value indicates a change to this value during the measurement interval.

Table 155. Fields in the Partition Data Report	
Field Heading	Meaning
Header Information	
MVS PARTITION NAME	The partition running the z/OS system which requested this report. VMSystem is displayed if the report was requested by a z/OS system running in a z/VM or alternate VM guest environment.
IMAGE CAPACITY	CPU capacity available to the MVS image measured in MSUs (millions of service units) per hour. The field is calculated as minimum of the following capacities: - the capacity based on the partition's logical CP configuration (includes online and standby [can be configured online] CPs) - the defined capacity limit of the partition, if available (image softcap) - the capacity limit of the related WLM capacity group, if the partition belongs to a capacity group - the absolute physical hardware capping limit - the capacity based on the hardware group capping limit. For z/OS systems running as a z/VM or alternate VM guest, the field displays the CPU capacity available to the z/VM or alternate VM partition.
NUMBER OF CONFIGURED PARTITIONS	The total number of activated and deactivated configured partitions. This number does not include the partition reported by the name *PHYSICAL*.
PHYS PROC NUM or VM PROC NUM	Number of physical processors in total and per processor type. Starting with IBM System z9 processors, IFLs and zAAPs are reported separately, and no longer as ICFs. If the data is reported for a z/OS system running as a z/VM or alternate VM guest, the field presents the number of processors that are assigned to the z/VM or alternate VM partition.

Table 155. Fields in the Partition Data Report (continued)

Field Heading	Meaning
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	Time (in milliseconds) a processor can be used when dispatched. This value is specified on the Logical Partition Control (LPCTO) frame on the Processor Controller Element (PCE). DYNAMIC appears in this field if a value is not specified and implies that the length of time a processor is assigned to a partition is dynamically allocated. See <i>PR/SM Planning Guide</i> for more information.
GROUP NAME	Name of the capacity group to which the partition belongs, if it is managed towards a common group capacity limit.
LIMIT	Capacity limit (in MSUs) defined for the partition's capacity group. An '*' following the limit value indicates that this partition started to be a member of this capacity group less than four hours ago. This partition will have a different view of unused group capacity and, therefore, may cap differently than existing group members.
AVAILABLE	Long-term average of CPU service in MSUs/h which would be allowed by the limit of the capacity group but is not used by its members. If the value is negative, the group is capped.
INITIAL CAP	Indicates whether the operator has set 'Initial Capping ON' in the logical partition controls of the Hardware Management Console (HMC) for the partition.
LPAR HW CAP	Indicates whether an absolute physical hardware capping limit has been defined in the logical partition controls of the HMC for any processor type of the partition.
HW GROUP CAP	Indicates whether an absolute hardware group capping limit has been defined in the logical partition group controls of the HMC for any processor type of the partition.
ABS MSU CAP	Indicates whether the ABSMSUCAPPING parameter has been set in the active IEAOPTxx parmlib member for the partition.
Partition Data	
NAME	The name that identifies a partition D Note: - The partition identified by the name *PHYSICAL* is not a configured partition. Data reported for *PHYSICAL* is shown only in columns DISPATCH TIME DATA - TOTAL, PHYSICAL PROCESSORS - LPAR MGMT, and PHYSICAL PROCESSORS - TOTAL. - When data about a z/OS system in a z/VM or alternate VM guest environment is reported, the *VMSys* line reports the time used by z/VM or alternate VM itself.
S	The current status of the partition: A Activated D Deactivated. The LPAR is configured but there are currently no logical CPUs online for this partition. If a partition is deactivated, the rest of the report line is blank.

Table 155. Fields in the Partition Data Report (continued)

Field Heading	Meaning
BT	The boost type that was active at some point within the interval. - For processor type CP: S Speed boost was active. N Speed boost was inactive. - For processor type IIP: I zIIP boost was active. N zIIP boost was inactive.
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 processors used in this partition.
MSU	Shows capacity information for a partition in terms of MSUs per hour. This information is shown for general purpose processors only. DEF Defined capacity limit of the partition. For the partition which is gathering the RMF data, this value is equal to the image capacity which is shown in the header of the report. ACT Actual consumption based on the logical processor effective dispatch time.
CAPPING	Shows capping information for a partition. DEF The hardware capping option of the partition. Each DEF value is a three position character string denoting which hardware capping mechanisms have or have not been applied in the logical partition controls of the HMC for the partition. The values in the first, second and third position of the string are either Y (Yes) or N (No) and have the following meaning: The first character (Y or N) indicates whether "Initial Capping ON" has been set. The second character (Y or N) indicates whether an absolute physical hardware capping limit (maximal number of CPUs) has been defined. The third character (Y or N) indicates whether an absolute hardware group capping limit (maximal number of CPUs) has been defined. For example, a DEF value of "Y N N" indicates that "Initial Capping ON" has been applied to this partition, but the other two options have not. The information provided in this field is useful only for logical partitions with shared processors. WLM% Percentage of time when WLM capped the partition. This information is shown only for general purpose processors.
Logical Partition Processor Data	
PROCESSOR NUM TYPE	The number of physical processors assigned to this partition and its processor type.
EFFECTIVE DISPATCH TIME	The sum of all processors' effective dispatch times for this partition during the measurement interval; expressed in the form HH.MM.SS.TTT. Partition effective dispatch time is the time, excluding LPAR management time, that a processor was assigned to this partition during the measurement interval.

Table 155. Fields in the Partition Data Report (continued)

Field Heading	Meaning
TOTAL DISPATCH TIME	The sum of all processors' dispatch times for this partition during the measurement interval, including LPAR management time. It is possible that the total dispatch time is smaller than the effective dispatch time. This situation occurs when partitions get "overruns" in their dispatch intervals caused by machine delays. The most typical form of this is caused by an MVS partition trying to talk to a coupling facility but getting significant delays or time-outs. It is sometimes symptomatic of recovery problems on the machine. For *PHYSICAL*, this value includes the time during which a physical CPU was busy, but the time could not be attributed to a specific logical partition. This time includes the time PR/SM was controlling the physical processor (LPAR management time), as well as any other time the processor was busy for any reason such as managing coupling facility traffic.
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. Note: If the z/OS system is running as guest under z/VM or alternate VM, and the Monitor I data gatherer option VMGUEST is active, the physical processor utilization represents the logical processor utilization of the z/VM or alternate VM LPAR.	
LOGICAL PROCESSORS - EFFECTIVE	The average partition effective dispatch time percentage. $\frac{\text{Effective Dispatch Time}}{\sum \text{Online Times}} * 100$
LOGICAL PROCESSORS - TOTAL	The average partition total dispatch time percentage. $\frac{\text{Total Dispatch Time}}{\sum \text{Online Times}} * 100$
PHYSICAL PROCESSORS - LPAR MGMT	The average LPAR management time on behalf of the partition reported as a percentage of the measurement interval. $\frac{\text{Total Dispatch Time} - \text{Effective Dispatch Time}}{\# \text{ Physical Processors} * \text{Interval Time}} * 100$ If the total dispatch time is smaller than the effective dispatch time, **** is shown in this column. The calculation for the *PHYSICAL* partition is: $\frac{\text{Total Time *PHYSICAL*}}{\# \text{ Physical Processors} * \text{Interval Time}} * 100$ Time *PHYSICAL* is the time that could not be attributed to a specific logical partition, but was used by PR/SM to control the physical processor (LPAR management time). Note: # Physical Processors denotes the number of physical processors of a certain processor group, either general purpose processors or special purpose processors.
PHYSICAL PROCESSORS - EFFECTIVE	The effective utilization of the physical processor resource by the partition. $\frac{\text{Effective Dispatch Time}}{\# \text{ Physical Processors} * \text{Interval Time}} * 100$
PHYSICAL PROCESSORS - TOTAL	The total utilization of the physical processor resource by the partition. $\frac{\text{Total Dispatch Time}}{\# \text{ Physical Processors} * \text{Interval Time}} * 100$

Table 155. Fields in the Partition Data Report (continued)	
Field Heading	Meaning
TOTAL	The sum of the current weightings of the shared processor resources. 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 *z/OS Resource Measurement Facility User's Guide*. The following table shows the overview condition names for the Overview report.

Table 156. Overview names in the Partition Data Report	
Field Heading or Meaning	Overview Name
PARTITION DATA - MSU DEF	LDEFMSU
PARTITION DATA - MSU ACT	LACTMSU
PARTITION DATA - CAPPING DEF	INICAP, LIMCPU
PARTITION DATA - CAPPING WLM%	WCAPPER
Available long-term average of CPU service (in MSUs/h)	GCMSUAV

Using the information in the LPAR Cluster Report

Starting with zSeries 900 (z900) servers, the Workload Manager is extended to work with PR/SM to dynamically expand resources that are available across LPARs.

An *LPAR cluster* is the subset of the systems that are running as LPARs on the same CEC. Based on business goals, WLM can direct PR/SM to enable or disable CP capacity for an LPAR, without human intervention.

LPAR CPU Management

Based on workload resource demand, the Workload Manager is able to dynamically adjust the number of logical processors and the weight of a logical partition. This allows the system to distribute the CPU resource in an LPAR cluster to partitions where the CPU demand is high. An LPAR cluster is defined as the set of logical partitions in a single CEC that belong to the same parallel sysplex.

The dynamic adjustment of processor resources within the partitions is reflected in the LPAR Cluster report, which provides LPAR views as well as aggregated views on LPAR cluster level.

LPAR CLUSTER REPORT														
z/OS 3.1			SYSTEM ID SYSE				DATE 09/30/2021		INTERVAL 08.52.731				PAGE 5	
			RPT VERSION 3.1 RMF				TIME 08.50.40		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		
SYSDPLEX	SYSD	SYSD	50			50			4	3	1.35	0.05	8192	N/A
	SYSE	SYSE	50			50			4	3	25.78	0.90	12288	N/A
	SYSF	SYSF	200			200			4	3	1.87	0.07	20480	N/A
	TOTAL		300						12	29.01	1.01		40960	N/A

Figure 201. CPU Activity Report - LPAR Cluster Report

Table 157. Fields in the LPAR Cluster Report	
Field Heading	Meaning
CLUSTER	This field identifies a sysplex name associated with the partition. All partitions that have the same cluster name are grouped together.
PARTITION	Name of the logical partition.
SYSTEM	z/OS system name.
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 defined minimum weighting, or 10% below the defined maximum weighting.
Processor Statistics	
NUMBER DEFINED / ACTUAL	Defined and average actual number of general purpose processors assigned to this partition. The actual number might be different from the defined number because of WLM goal achievement reasons.
TOTAL % LBUSY	Total dispatch time reported as a percentage of the logical processor online time: $\frac{\text{Partition Total Dispatch Time}}{\sum \text{Logical Processor Online Times}} * 100$ The contents of this field is equal to the LOGICAL PROCESSOR UTILIZATION - TOTAL column in the Partition Data report.
TOTAL % PBUSY	Total utilization of the physical processor resource by the partition: $\frac{\text{Partition Total Dispatch Time}}{\# \text{ Physical Processors} * \text{Interval Time}} * 100$ The contents of this field is equal to the PHYSICAL PROCESSOR UTILIZATION - TOTAL column in the Partition Data report.
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 *z/OS Resource Measurement Facility User's Guide*. The following table shows the overview condition names for the Overview report.

Table 158. Overview names in the LPAR Cluster Report	
Field Heading or Meaning	Overview Name
WEIGHTING - DEFINED INIT (Cluster)	WDEFC
WEIGHTING - DEFINED INIT (general purpose processors)	WDEFL
WEIGHTING - DEFINED MIN	WMINL
WEIGHTING - DEFINED MAX	WMAXL

Table 158. Overview names in the LPAR Cluster Report (continued)	
Field Heading or Meaning	Overview Name
WEIGHTING - ACTUAL AVG	WACTL
WEIGHTING - ACTUAL MIN%	WMIPL
WEIGHTING - ACTUAL MAX%	WMAPL
PROCESSOR - NUMBER DEFINED (Cluster)	NLDEFC
PROCESSOR - NUMBER DEFINED (Partition)	NLDEFL
PROCESSOR - NUMBER ACTUAL	NLACTL
PROCESSOR - TOTAL% LBUSY (Cluster)	LBUSYC
PROCESSOR - TOTAL% LBUSY (Partition)	LBUSYL
PROCESSOR - TOTAL% PBUSY (Cluster)	PBUSYC
PROCESSOR - TOTAL% PBUSY (general purpose processors)	PBUSYL

Using the information in the Group Capacity Report

You can apply a defined capacity limit not only to one logical partition, but to a group of LPARs on the same CEC and manage this group considering the combined defined capacities of all members of the group.

With the group capacity limit, a third restriction to an LPAR is added. Even when an LPAR is not limited by its weight or its defined capacity, it can be limited by the group capacity. The minimum of the following limitations is applied to any partition:

1. Defined capacity
2. LPAR weights
3. Group capacity limit
4. Absolute physical hardware capping limit
5. Hardware group capping limit (refer to [“Using the information in the Hardware Group Report”](#) on page 366 for more information)

The Group Capacity Report monitors the available capacity of each defined capacity group and the MSU consumption and actual capping of these groups and of each partition within such a group. It helps you to exploit the flexibility to use as much CPU as needed for short periods of time until the 4 hour rolling MSU average exceeds the defined capacity limit for the whole group.

GROUP CAPACITY REPORT											
z/OS 3.1			SYSTEM ID TRX2			DATE 04/25/2023			INTERVAL 05.00.000		
			RPT VERSION 3.1 RMF			TIME 21.55.00			CYCLE 1.000 SECONDS		
									PAGE 3		
GROUP-CAPACITY NAME	LIMIT	PARTITION	SYSTEM	DEF	ACT	WGT	DEF	WLM%	ACT%	ENTITLEMENT - MINIMUM	MAXIMUM
RMFGRP	400	TRX1	TRX1	0	5	40	NO	0.0	0.0	29	400
		TRX2	TRX2	0	11	500	NO	0.0	0.0	370	400
TOTAL					16	540					

Figure 202. CPU Activity Report - Group Capacity Report

Table 159. Fields in the Group Capacity Report	
Field Heading	Meaning
GROUP-CAPACITY NAME	Name of the capacity group.
GROUP-CAPACITY LIMIT	MSU limit defined for the capacity group.
PARTITION	Name of the logical partition.
SYSTEM	Name of the z/OS system.
MSU DEF	User defined capacity limit.

Table 159. Fields in the Group Capacity Report (continued)	
Field Heading	Meaning
MSU ACT	Actual MSU consumption of this partition based on the logical processor effective dispatch time.
WGT	The partition's weighting of the shared processor resources which is used for WLM Group Capacity decisions. In case of hard capped partitions (see field CAPPING DEF), dashes (---) are displayed.
CAPPING DEF	The initial capping option of the partition: YES/NO Initially capped partitions (also referred to as hard capped) are excluded from WLM group capacity management.
CAPPING WLM%	Percentage of time when WLM considers to cap the partition. For more information on WLM capping, refer to topic Workload management and Workload License Charges in <i>z/OS MVS Planning: Workload Management</i> .
CAPPING ACT%	Percentage of time when capping actually limited the usage of processor resources for the partition.
MINIMUM ENTITLEMENT	The minimum share of the MSU limit defined for the capacity group that the partition receives, even if all other partitions within the capacity group are running high workload. N/A is displayed for hard capped partitions.
MAXIMUM ENTITLEMENT	The maximum share of the MSU limit defined for the capacity group that a partition can receive if all other partitions within the capacity group are running without workload. N/A is displayed for hard capped partitions.

Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*. The following table shows the overview condition names for the Overview report.

Table 160. Overview names in the Group Capacity Report	
Field Heading or Meaning	Overview Name
MSU - ACT	GCMSUACT
WGT	GCWEIGHT
MINIMUM ENTITLEMENT	MINENT
MAXIMUM ENTITLEMENT	MAXENT

Using the information in the Hardware Group Report

The Hardware Group Report displays the settings of the hardware groups and their partitions.

With the hardware group capping limit, which can be set in the logical partition controls of the Hardware Management Console (HMC), another restriction to an LPAR is added. This limit enforces an absolute capping for each type of processor defined to every partition in the hardware group. Even when an LPAR is not limited by its weight, its defined capacity, or a group capacity limit, it can be limited by this hardware group capping value. The minimum of the following limitations is applied to any partition:

- Defined capacity
- LPAR weights
- Group capacity limit
- Absolute physical hardware capping limit
- Hardware group capping value

H A R D W A R E G R O U P R E P O R T									
z/OS 3.1		SYSTEM ID R74		DATE 09/30/2021		INTERVAL 08.52.731			
		RPT VERSION 3.1 RMF		TIME 08.50.40		CYCLE 1.000 SECONDS			
HW GROUP		PARTITION		SYSTEM		HW GROUP LIMIT			
NAME									
SYSDPLEX		SYSD		SYSE		CP CBP IIP ICF IFL			
		SYSE		SYSE		0.00 2.00 0.00 0.00 0.00			
		SYSE		SYSE					

Figure 203. CPU Activity Report – Hardware Group Report

Table 161. Fields in the Hardware Group Report	
Field Heading	Meaning
HW GROUP NAME	Name of the hardware group.
PARTITION	Name of the logical partition.
SYSTEM	Name of the z/OS system.
HW GROUP LIMIT	Absolute limit on partition usage of all processor types that are members of the same hardware group, in terms of numbers of CPUs. If the hardware group name or the limit changed during the reporting interval, an asterisk (*) is appended.

Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*. [Table 162 on page 367](#) shows the overview condition names for the Overview report.

Table 162. Overview names in the Hardware Group Report	
Field Heading or Meaning	Overview Name
HW GROUP LIMIT - CP	HGCCP
HW GROUP LIMIT - CBP	HGCCBP
HW GROUP LIMIT - IIP	HGCIIP
HW GROUP LIMIT - ICF	HGCICF
HW GROUP LIMIT - IFL	HGCIFL

Using the information in the Power Consumption report

The Power consumption reports are essential for managing energy efficiency and operational costs. This report provides insights into the power usage patterns of the system over defined intervals, allowing you to monitor and optimize performance.

The Power Consumption report includes three rows of values: the AVG POWER values indicate the average power recorded over the time interval and are displayed in watts; the ENERGY per MIN and ENERGY values represent the amount of energy consumed during one minute or the respective interval, shown in watt-seconds.

The report presented below is based on a sample dataset. [Figure 204 on page 367](#) illustrates the Power Consumption report for the first interval, which is generated using **REPORTS (CPU)**.

POWER CONSUMPTION REPORT										PAGE	4	
z/OS 3.1		SYSTEM ID SYS1			DATE 02/05/2025		INTERVAL 05.00.184					
		RPT VERSION 3.1 RMF			TIME 08.55.01		CYCLE 1.000 SECONDS					
POWER CONSUMPTION												
		----- CPC -----			----- LPAR -----							
		--TOTAL-	UNASSIGNED	--INFRA-	--TOTAL-	PROCESSOR	--I/O--	--MEMORY-				
AVG POWER		18964.13	5340.836	7165.109	601.0667	576.0667	17.00000	8.000000				
ENERGY		5692729	1603233	2150852	180430	172925	5103.125	2401.464				
ENERGY/MIN		1137847	320450	429907	36064	34564	1020.000	400.0000				

Figure 204. CPU Activity Report - Power Consumption Report

Table 163. Fields in the Power Consumption Report

Field Heading	Meaning
CPC TOTAL	The total power or energy of the Central Processor Complex (CPC), containing all available resources that can be utilized for workloads.
CPC UNASSIGNED	The portion of the CPC's total power or energy associated with resources that are in a standby or reserved state.
CPC INFRA	The portion of the CPC's total power or energy allocated for infrastructure-related components within the CPC.
LPAR TOTAL	The total power or energy of the specific Logical Partition (LPAR).
LPAR PROCESSOR	The portion of the LPAR's total power or energy allocated for CPU resources.
LPAR I/O	The portion of the LPAR's total power or energy allocated for Input/Output operations within the LPAR.
LPAR MEMORY	The portion of the LPAR's total power or energy allocated for memory resources within the LPAR.

Spreadsheet and Overview reference

You can make the Power Consumption report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*. [Table 164 on page 368](#) shows the overview condition names for the Overview report.

Table 164. Overview names in the Power Consumption Report

Overview Name	Field Heading or Meaning
CPCTOTE	The total energy consumed during the interval by the Central Processor Complex (CPC), containing all available resources that can be utilized for workloads. The values are represented in watt-seconds.
CPCUNAE	The portion of the CPC's total energy associated with resources that are in a standby or reserved state. The values are represented in watt-seconds.
CPCINFE	The amount of energy consumed by the infrastructure related components within the CPC during the interval. The values are represented in watt-seconds.
CPCMINE	The energy consumption of the CPC within a minute. The values are represented in watt-seconds.
CPCAVGP	The average power recorded by the CPC over a specific time. The values are represented in watts.
LPARTOTE	The total energy consumed by a specific Logical Partition (LPAR) during the interval. The values are represented in watt-seconds.
LPARCPUE	The amount of energy consumed by CPU resources allocated to the LPAR, determining its processing capabilities for executing applications and tasks. The values are represented in watt-seconds.
LPARIOE	The amount of energy consumed by Input/Output operations within the LPAR. The values are represented in watt-seconds.
LPARMEME	The amount of energy consumed by memory resources allocated to the LPAR. The values are represented in watt-seconds.
LPARMINE	The energy consumption of each logical partition over a one-minute interval. The values are represented in watt-seconds.
LPARAVGP	The average power recorded for a specific LPAR over a specific time. The values are represented in watts.

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 Resource Measurement Facility User's Guide* for more information about this processing.

2. If the SMF records contain the same system identifiers, but differs in MVS partition name, number of configured partitions or partition name, the records are processed selectively.

The first SMF record encountered defines the type of records to be accumulated. Any subsequent records that do not have the same characteristics are skipped.

3. If the SMF records contain the same system identifier but differs in number of physical processors, status, wait completion or number of logical processors, the records are processed as if they were from the same system. All records are processed. When a new value is encountered, an asterisk (*) appears next to the changed value on the report. The new value is reported.

CRYPTO - Crypto Hardware Activity report

The **Crypto Hardware Activity** report provides information about the activities in the various cryptographic hardware functions. Most cryptographic hardware functions can only be used through Cryptographic Support for z/OS (ICSF). ICSF is a standard component of z/OS. It provides cryptographic services in the z/OS environment. The report provides the following sections:

- **Cryptographic CCA coprocessors**

This section provides measurements about secure cryptographic functions executed on Common Cryptographic Architecture (CCA) coprocessors, use of secure encrypted key values, clear key and secure PKA operations, and special user cryptographic functions (using the user defined extension (UDX) capability of the card). For cryptographic CCA coprocessors, special attention should be given to RSA key-generation operations because these operations require a high amount of cryptographic processing capacity. Therefore, they are reported in addition to the total number of operations.

- **Cryptographic PKCS11 coprocessors**

This section provides measurements about secure public-key operations executed by cryptographic symmetric- and asymmetric-key functions.

- **Cryptographic accelerators**

This section provides measurements about public key operations (RSA cryptography operations) used with Secure Sockets Layer (SSL) or Transport Layer Security (TLS) protocols which are widely used to help secure e-business applications. The data for cryptographic accelerators is showing details for the two available algorithms, modular exponentiation (ME) and Chinese Remainder Theorem (CRT) for available key lengths (1024, 2048, and 4096 bit). This provides information how the usage of these algorithms affects the utilization of the accelerator.

- **ICSF Services**

The Crypto Hardware Activity report provides performance measurements on selected ICSF activities:

- Using the single and triple Data Encryption Standard (DES) and the Advanced Encryption Standard (AES) to encipher and decipher data.

- Generating and verifying message authentication codes (MAC). The MAC is a value calculated from the message according to a secret shared DES key or AES key and sent to the receiver together with the message. The receiver can recalculate the MAC and compare it with the MAC received. If the MAC values are identical, the message has not been altered during transmission.
- Using public hash functions. A hash is calculated from the transmission data according to a public key or function in cases where it is impossible to share a secret key. If the recalculated hash is identical to the one calculated before transmission, data integrity is ensured.
- Translating and verifying PINs.
- Digital signature generation and verification. A digital signature is created using the data to be signed and a private key, using one of the following algorithms:
 - RSA (Ron Rivest, Adi Shamir and Leonard Adleman)
 - ECC (Elliptic Curve Cryptography)
 - QSA (Quantum Safe)

The digitally signed data is sent to the receiver. The receiver can verify that the signature is valid, using the signer's public key.
- Format preserving encryption (FPE) and Feistel-based encryption (FFX) to encipher, decipher, and translate data while preserving the original formatting of the data.

How to request this report

Monitor I gathers data for this report automatically. If you want to suppress gathering, you need to specify NOCRYPTO.

To produce this report, specify

```
REPORTS(CRYPTO)
```

This report is also available in XML output format. Topic [How to work with Postprocessor XML reports](#) in *z/OS Resource Measurement Facility User's Guide* provides all required information on how to produce and view XML reports.

Example URL for the DDS API

```
http://ddshost:8803/gpm/imfpp.xml?reports=CRYPTO
```

Contents of the report

The data shown for cryptographic coprocessors and accelerators always reflects the total activity in your CPC and the partition activity, 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.




Figure 205. Crypto Hardware Activity Report

Table 165. Fields in the CRYPTO Hardware Activity Report	
Field Heading	Meaning
Cryptographic accelerator or coprocessor performance data reflecting the CPC activity is reported in blocks which are labeled with CPC . Cryptographic accelerator or coprocessor performance data reflecting the partition activity is reported in blocks which are labeled with LPAR .	
Cryptographic CCA Coprocessor	
TYPE	Type that defines the cryptographic CCA coprocessor: Type Meaning CEX2C Crypto Express2 Coprocessor. CEX3C Crypto Express3 Coprocessor. CEX4C Crypto Express4S Coprocessor. CEX5C Crypto Express5S Coprocessor. CEX6C Crypto Express6S Coprocessor. CEX7C Crypto Express7S Coprocessor.
ID	Index that specifies the cryptographic CCA coprocessor.
CPC /LPAR	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.

Table 165. Fields in the CRYPTO Hardware Activity Report (continued)

Field Heading	Meaning
CPC /LPAR KEY-GEN RATE	Rate for RSA-key-generation operations.
Cryptographic PKCS11 Coprocessor	
TYPE	Type that defines the cryptographic PKCS11 coprocessor: Type Meaning CEX4P Crypto Express4S PKCS11 Coprocessor. CEX5P Crypto Express5S PKCS11 Coprocessor. CEX6P Crypto Express6S PKCS11 Coprocessor. CEX7P Crypto Express7S PKCS11 Coprocessor.
ID	Index that specifies the cryptographic PKCS11 coprocessor.
CPC /LPAR	Rate, average execution time (in milliseconds) and utilization percentage for all operations executed on this cryptographic PKCS11 coprocessor.
CPC /LPAR FUNCTION	Rate, average execution time (in milliseconds) and utilization percentage for executed operations, categorized by cryptographic function type: Type Meaning ASYM FAST Fast asymmetric-key function. ASYM GEN Asymmetric-key generation function. ASYM SLOW Slow asymmetric-key function. SYMM COMPLETE Symmetric-key function that returns a complete or final result. SYMM PARTIAL Symmetric-key function that returns partial or incremental results.
Cryptographic Accelerator	
TYPE	Type that defines the cryptographic accelerator: Type Meaning CEX2A Crypto Express2 Accelerator. CEX3A Crypto Express3 Accelerator. CEX4A Crypto Express4S Accelerator. CEX5A Crypto Express5S Accelerator. CEX6A Crypto Express6S Accelerator. CEX7A Crypto Express7S Accelerator.
ID	Index that specifies the cryptographic accelerator.
CPC /LPAR	Rate, average execution time (in milliseconds) and utilization for all operations on this cryptographic accelerator.

Table 165. Fields in the CRYPTO Hardware Activity Report (continued)

Field Heading	Meaning
CPC /LPAR FUNCTION	Rate, average execution time (in milliseconds) and utilization percentage for executed operations, categorized by cryptographic function type: Type Meaning RSA ME 1024 RSA ME 2048 RSA ME 4096 ME Format RSA operations with RSA key length 1024, 2048, 4096 RSA CRT 1024 RSA CRT 2048 RSA CRT 4096 CRT Format RSA operations with RSA key length 1024, 2048, 4096
ICSF services	
ENCRYPTION	SDES RATE Rate of encipher service calls using single DES. SDES SIZE Average number of bytes per service call that have been enciphered using single DES. TDES RATE Rate of encipher service calls using double and triple DES. TDES SIZE Average number of bytes per service call that have been enciphered using double and triple DES. AES RATE Rate of encipher service calls using AES. AES SIZE Average number of bytes per service call that have been enciphered using AES.
DECRYPTION	SDES RATE Rate of decipher service calls using single DES. SDES SIZE Average number of bytes per service call that have been deciphered using single DES. TDES RATE Rate of decipher service calls using double and triple DES. TDES SIZE Average number of bytes per service call that have been deciphered using double and triple DES. AES RATE Rate of decipher service calls using AES. AES SIZE Average number of bytes per service call that have been deciphered using AES. Note: For AES, only service calls and bytes sent to a coprocessor are reported.
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.

Table 165. Fields in the CRYPTO Hardware Activity Report (continued)

Field Heading	Meaning
HASH	SHA-1 RATE Rate of requests to hash using the SHA-1 hash algorithm. SHA-1 SIZE Average number of bytes to be hashed per request using the SHA-1 hash algorithm. SHA-256 RATE Rate of requests to hash using the SHA-224 or the SHA-256 hash algorithm. SHA-256 SIZE Average number of bytes to be hashed per request using the SHA-224 or the SHA-256 hash algorithm. SHA-512 RATE Rate of requests to hash using the SHA-384 or the SHA-512 hash algorithm. SHA-512 SIZE Average number of bytes to be hashed per request using the SHA-384 or the SHA-512 hash algorithm.
PIN	TRANSLATE RATE Rate of requests to translate PIN. VERIFY RATE Rate of requests to verify PIN.
AES MAC	GENERATE RATE Rate of requests to generate AES MACs. GENERATE SIZE Average number of bytes per request for which AES MACs have been generated. VERIFY RATE Rate of requests to verify AES MACs. VERIFY SIZE Average number of bytes per request for which AES MACs have been verified.
RSA DSIG	GENERATE RATE Rate of requests to generate RSA digital signatures. VERIFY RATE Rate of requests to verify RSA digital signatures.
ECC DSIG	GENERATE RATE Rate of requests to generate ECC digital signatures. VERIFY RATE Rate of requests to verify ECC digital signatures.
FORMAT PRESERVING ENCRYPTION	ENCIPHER RATE Rate of requests to encipher data using FPE. ENCIPHER SIZE Average number of bytes per request that have been enciphered using FPE. DECIPHER RATE Rate of requests to decipher data using FPE. DECIPHER SIZE Average number of bytes per request that have been deciphered using FPE. TRANSLATE RATE Rate of requests to translate data using FPE. TRANSLATE SIZE Average number of bytes per request that have been translated using FPE.
QSA DSIG	GENERATE RATE Rate of requests to generate QSA digital signatures. VERIFY RATE Rate of requests to verify QSA digital signatures.

Table 165. Fields in the CRYPTO Hardware Activity Report (continued)

Field Heading	Meaning
FEISTEL-BASED ENCRYPTION	ENCIPHER RATE Rate of requests to encipher data using FFX.
	ENCIPHER SIZE Average number of bytes per request that have been enciphered using FFX.
	DECIPHER RATE Rate of requests to decipher data using FFX.
	DECIPHER SIZE Average number of bytes per request that have been deciphered using FFX.
	TRANSLATE RATE Rate of requests to translate data using FFX.
	TRANSLATE SIZE Average number of bytes per request that have been translated using FFX.

Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*. The following table shows the overview condition names for the Overview report.

Table 166. Overview names in the CRYPTO Hardware Activity Report

Field Heading or Meaning	Overview Name
Cryptographic coprocessor TOTAL RATE	CRYCTR
Cryptographic coprocessor TOTAL EXEC TIME	CRYCTE
Cryptographic coprocessor TOTAL UTIL%	CRYCTU
Cryptographic coprocessor KEY-GENERATION RATE	CRYCKR
Cryptographic accelerator ME(1024) RATE	CRYAM1R
Cryptographic accelerator ME(1024) EXEC TIME	CRYAM1E
Cryptographic accelerator ME(1024) UTIL%	CRYAM1U
Cryptographic accelerator ME(2048) RATE	CRYAM2R
Cryptographic accelerator ME(2048) EXEC TIME	CRYAM2E
Cryptographic accelerator ME(2048) UTIL%	CRYAM2U
Cryptographic accelerator ME(4096) RATE	CRYAM3R
Cryptographic accelerator ME(4096) EXEC TIME	CRYAM3E
Cryptographic accelerator ME(4096) UTIL%	CRYAM3U
Cryptographic accelerator CRT(1024) RATE	CRYAC1R
Cryptographic accelerator CRT(1024) EXEC TIME	CRYAC1E
Cryptographic accelerator CRT(1024) UTIL%	CRYAC1U
Cryptographic accelerator CRT(2048) RATE	CRYAC2R
Cryptographic accelerator CRT(2048) EXEC TIME	CRYAC2E
Cryptographic accelerator CRT(2048) UTIL%	CRYAC2U
Cryptographic accelerator CRT(4096) RATE	CRYAC3R
Cryptographic accelerator CRT(4096) EXEC TIME	CRYAC3E
Cryptographic accelerator CRT(4096) UTIL%	CRYAC3U
ENCRYPTION SDDES RATE	CRYISDER
ENCRYPTION SDDES SIZE	CRYISDES

Table 166. Overview names in the CRYPTO Hardware Activity Report (continued)

Field Heading or Meaning	Overview Name
ENCRYPTION TDES RATE	CRYITDER
ENCRYPTION TDES SIZE	CRYITDES
ENCRYPTION AES RATE	CRYIAER
ENCRYPTION AES SIZE	CRYIAES
Average number of coprocessor calls for AES encipher services	CRYIAEO
DECRYPTION SDES RATE	CRYISDDR
DECRYPTION SDES SIZE	CRYISDDS
DECRYPTION TDES RATE	CRYITDDR
DECRYPTION TDES SIZE	CRYITDDS
DECRYPTION AES RATE	CRYIADR
DECRYPTION AES SIZE	CRYIADS
Average number of coprocessor calls for AES decipher services	CRYIADO
MAC GENERATE RATE	CRYIMGR
MAC GENERATE SIZE	CRYIMGS
MAC VERIFY RATE	CRYIMVR
MAC VERIFY SIZE	CRYIMVS
HASH SHA-1 RATE	CRYIHAR
HASH SHA-1 SIZE	CRYIHAS
HASH SHA-256 RATE	CRYIH2R
HASH SHA-256 SIZE	CRYIH2S
PIN TRANSLATE RATE	CRYIPTR
PIN VERIFY RATE	CRYIPVR
AES MAC GENERATE RATE	CRYIAMGR
AES MAC GENERATE SIZE	CRYIAMGS
AES MAC VERIFY RATE	CRYIAMVR
AES MAC VERIFY SIZE	CRYIAMVS
RSA DIGITAL SIGNATURE GENERATE RATE	CRYIDRGR
RSA DIGITAL SIGNATURE VERIFY RATE	CRYIDRVR
ECC DIGITAL SIGNATURE GENERATE RATE	CRYIDEGR
ECC DIGITAL SIGNATURE VERIFY RATE	CRYIDEVVR
FPE ENCIPHER RATE	CRYIFPER
FPE ENCIPHER SIZE	CRYIFPES
FPE DECIPHER RATE	CRYIFPDR
FPE DECIPHER SIZE	CRYIFPDS
FPE TRANSLATION RATE	CRYIFPTR
FPE TRANSLATION SIZE	CRYIFPTS
QSA DIGITAL SIGNATURE GENERATE RATE	CRYIDQGR
QSA DIGITAL SIGNATURE VERIFY RATE	CRYIDQVR
FFX ENCIPHER RATE	CRYIFXER

Table 166. Overview names in the CRYPTO Hardware Activity Report (continued)

Field Heading or Meaning	Overview Name
FFX ENCIPHER SIZE	CRYIFXES
FFX DECIPHER RATE	CRYIFXDR
FFX DECIPHER SIZE	CRYIFXDS
FFX TRANSLATE RATE	CRYIFXTR
FFX TRANSLATE SIZE	CRYIFXTS
Cryptographic PKCS11 coprocessor TOTAL RATE	CRYPTR
Cryptographic PKCS11 coprocessor TOTAL UTIL%	CRYPTU
Cryptographic PKCS11 coprocessor TOTAL EXEC TIME	CRYPTE
Cryptographic PKCS11 coprocessor SLOW ASYM RATE	CRYP SAR
Cryptographic PKCS11 coprocessor SLOW ASYM UTIL%	CRYP SAU
Cryptographic PKCS11 coprocessor SLOW ASYM EXEC TIME	CRYP SAE
Cryptographic PKCS11 coprocessor FAST ASYM RATE	CRYP FAR
Cryptographic PKCS11 coprocessor FAST ASYM UTIL%	CRYP FAU
Cryptographic PKCS11 coprocessor FAST ASYM EXEC TIME	CRYP FAE
Cryptographic PKCS11 coprocessor SYMM PART RATE	CRYP SPR
Cryptographic PKCS11 coprocessor SYMM PART UTIL%	CRYP SPU
Cryptographic PKCS11 coprocessor SYMM PART EXEC TIME	CRYP SPE
Cryptographic PKCS11 coprocessor SYMM COMPL RATE	CRYP SCR
Cryptographic PKCS11 coprocessor SYMM COMPL UTIL%	CRYP SCU
Cryptographic PKCS11 coprocessor SYMM COMPL EXEC TIME	CRYP SCE
Cryptographic PKCS11 coprocessor ASYM GEN RATE	CRYPAGR
Cryptographic PKCS11 coprocessor ASYM GEN UTIL%	CRYPAGU
Cryptographic PKCS11 coprocessor ASYM GEN EXEC TIME	CRYPAGE

DEVICE - Device Activity report

The Device Activity report provides information for all devices in one or more device classes (such as TAPE or DASD) or for those devices you specify on the DEVICE option.

When used with the Channel Path Activity and I/O Queuing Activity reports, this report can help you analyze the I/O activity at your installation and identify bottlenecks caused by a particular device.

How to request this report

Monitor I gathers data for this report automatically with the default option DEVICE (DASD). If you want to suppress gathering, you need to specify NODEVICE.

To produce this report, specify

```
REPORTS(DEVICE(type))
```

This report is also available in XML output format. [How to work with Postprocessor XML reports in z/OS Resource Measurement Facility User's Guide](#) provides all required information on how to produce and view XML reports.

Example URL for the DDS API:

```
http://ddshost:8803/gpm/rmfpp.xml?reports=DEVICE(NMBR(2000,3FFF))&sysid=SYSA
```

Contents of the report

Each Device Activity report begins on a new page, and the class of devices included in the report is indicated by one of the following titles:

Report Title	What you specified
CHARACTER READER DEVICE ACTIVITY	DEVICE (CHRD)
COMMUNICATION EQUIPMENT ACTIVITY	DEVICE (COMM)
DIRECT ACCESS DEVICE ACTIVITY	DEVICE (DASD)
SYNCHRONOUS I/O DEVICE ACTIVITY	DEVICE (DASD)
GRAPHICS DEVICE ACTIVITY	DEVICE (GRAPH)
MAGNETIC TAPE DEVICE ACTIVITY	DEVICE (TAPE)
UNIT RECORD DEVICE ACTIVITY	DEVICE (UNITR)

Whereas the DASD Activity report provides general DASD device characteristics together with detailed device activity information gathered in time periods where the DASD device was accessed in asynchronous mode (via FICON/HPF paths), the Synchronous I/O Device Activity report section adds an overview on synchronous I/O performance statistics if at least one DASD device actively performed synchronous I/O requests (via synchronous I/O links) during the interval.

The devices included in the report are grouped by logical control unit. The logical control unit provides a way to identifying a related set of devices. Moreover, this organization makes it easier to compare the data in the Device Activity report with the data in the I/O Queuing Activity report.

RMF follows the individual device data lines in each group with a summary line that provides a weighted average or total values for the entire logical control unit.

NMBR and SG suboptions

If the NMBR and SG suboptions are specified together, the device report is divided into two parts.

- The first part of the report contains the devices specified by the NMBR suboption. The devices are sorted by LCU and device number.
- The second part of the report contains the devices specified for the SG suboption. The devices are sorted by storage group and by device numbers within the group.

Note: Some devices might be reported twice, since you can specify a device on the NMBR suboption that is part of a storage group specified on the SG suboption.

Byte-multiplexor-channel-attached device

For any device attached to a byte multiplexor channel, the only measurement data available is the start subchannel (SSCH) + resume subchannel (RSCH) instruction count.

Direct Access Device Activity report

For the DASD Activity report, the information can be sorted by LCU, or storage group, or both. When the storage group (SG) option is specified, the DASD Activity report is sorted by device number within each storage group. The storage group name that a volume is assigned to is always reported, even when the SG option was not selected. If a volume does not belong to a storage group, the STORAGE GROUP field for that volume is blank.

RMF follows the individual device data lines in each group with a summary line that provides average or total values for the entire storage group.

Note: When comparing I/O rates in the DASD Activity report and in the Cache Subsystem Activity report, you may see differences due to different ways how I/Os are counted:

- In the DASD Activity report, one I/O is counted for one SSCH or RSCH instruction. There can be record chaining, for example for paging I/O, which is not reflected in the SSCH count.
- In the Cache Subsystem Activity report, one I/O is counted for each cache request, and one I/O chain may cause several cache requests.

These two ways can lead to higher I/O rates in the Cache Subsystem Activity report than in the DASD Activity report.

Figure 206 on page 379 shows a DASD activity sample report.

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	100
z/OS 3.1			SYSTEM ID SYSB			DATE 04/25/2023			INTERVAL 15.00.000													
			RPT VERSION 3.1 RMF			TIME 15.00.00			CYCLE 1.000 SECONDS													
TOTAL SAMPLES = 900			IODF = D4			CR-DATE: 08/15/2022			CR-TIME: 11.09.47			ACT: POR										
STORAGE GROUP	DEV	DEVICE	NUMBER OF CYL	VOLUME SERIAL	PAV	LCU	DEVID	ACTIVITY	AVG RESP TIME	AVG I/OQ	AVG CMR	AVG DLY	AVG INT	AVG PEND	AVG DISC	AVG CONN	% DEV UTIL	% DEV RESV	AVG NMBR ALLOC	% ANY		
DBNONRLS	05F70	3390A	262668	NRLS17	1.2H	0052	796.4695	1.24	.461	.005	.000	.000	.089	.002	.685	45.53	45.63	0.0	50.0	100.0		
SMSOE	05F71	3390A	262668	S50026	1.0H	0052	0.062	.701	.002	.009	.000	.000	.098	.000	.601	0.00	0.00	0.0	38.0	100.0		
SMSOE	05F72	3390A	262668	S50092	1.0H	0052	1.094	.964	.000	.010	.000	.000	.105	.000	.859	0.09	0.09	0.0	20.0	100.0		
SMSOE	05F73	3390A	262668	S50006	1.0H	0052	1.711	1.15	.012	.009	.265	.000	.364	.090	.683	0.12	0.13	0.0	97.0	100.0		
DATABASE	05F74	3390A	262668	DB2105	1.0H	0052	0.024	.360	.005	.006	.000	.000	.099	.000	.256	0.00	0.00	0.0	0.0	100.0		
DATABASE	05F75	3390A	262668	DB2285	1.0H	0052	1.195	.900	.000	.006	.000	.000	.090	.000	.802	0.10	0.10	0.1	12.0	100.0		
DATABASE	05F76	3390A	262668	DB8136	1.0H	0052	10.8735	.966	.000	.004	.000	.000	.089	.017	.861	0.94	0.95	0.0	102	100.0		
DATABASE	05F77	3390A	262668	DB2185	1.0H	0052	6.3655	1.06	.000	.005	.001	.000	.090	.006	.963	0.61	0.62	0.0	85.0	100.0		
LOGGER	05F78	3390A	262668	LOGR62	1.0H	0052	44.023	.511	.000	.005	.000	.000	.086	.000	.424	1.87	1.87	0.0	28.8	100.0		
DBZHPF	05F79	3390A	262668	DB2080	1.0H	0052	0.024	.916	.003	.012	.000	.000	.105	.000	.809	0.00	0.00	0.0	1.0	100.0		
DBZHPF	05F7A	3390A	262668	DB2280	1.0H	0052	0.024	1.15	.007	.023	.000	.000	.105	.000	1.04	0.00	0.00	0.0	0.0	100.0		
SUBSYS	05F7B	3390A	262668	SUBS8A	1.1H	0052	464.7275	.524	.000	.004	.001	.000	.086	.001	.437	18.67	18.70	0.0	23.0	100.0		
SUBSYS	05F7C	3390A	262668	SUBS6C	1.0H	0052	342.2805	.774	.087	.006	.002	.000	.090	.118	.479	15.97	19.91	1.4	23.0	100.0		
ASYNCH	05F7D	3390A	262668	SYN361	1.0H	0052	0.023	.482	.013	.018	.000	.000	.098	.000	.372	0.00	0.00	0.0	0.0	100.0		
ASYNCH	05F7E	3390A	262668	SYN362	1.0H	0052	0.023	.923	.000	.006	.000	.000	.104	.000	.719	0.00	0.00	0.0	0.0	100.0		
ASYNCH	05F7F	3390A	262668	SYN363	1.0H	0052	0.023	.654	.008	.012	.000	.000	.104	.000	.542	0.00	0.00	0.0	0.0	100.0		
...				LCU		0052	2075.27	.876	.192	.006	.001	.000	.090	.021	.573	0.84	0.88	0.0	2659	100.0		
...																						
06200	33909	10017			1.0H	0055	0.7115	.202	.000	.000	.000	.000	.053	.054	.095	0.01	0.01	0.0	0.0	100.0		
06201	33909	10017			1.0H	0055	5.0025	.186	.000	.000	.000	.000	.054	.040	.092	0.05	0.07	0.0	0.0	100.0		
06202	33909	10017			1.0H	0055	0.8435	.181	.000	.000	.000	.000	.052	.058	.096	0.01	0.01	0.0	0.0	100.0		
06203	33909	10017			1.0H	0055	4.5555	.181	.000	.000	.000	.000	.052	.036	.092	0.04	0.06	0.0	0.0	100.0		
06204	33909	10017			1.0H	0055	1.3375	.198	.000	.000	.000	.000	.053	.048	.097	0.01	0.02	0.0	0.0	100.0		
					LCU	0055	12.448	.188	.000	.000	.000	.000	.053	.041	.093	0.02	0.03	0.0	0.0	100.0		

Figure 206. Direct Access Device Activity Report

Synchronous I/O Device Activity report

The Synchronous I/O Device Activity section is only part of the report if at least one DASD device actively performed synchronous I/O requests using IBM zHyperLink technology. Sorting of columns and insertion of an LCU summary line in this report is identical to that of the Direct Access Device Activity report.

A device with synchronous I/O activity may be mapped back to the synchronous I/O link by which it is reached by looking up the serial number and node descriptor information of the device's storage controller in the RMF Cache Subsystem Device Overview report. Serial number and type-model of the storage controller can be looked up in the Synchronous I/O Link Activity section of the RMF Postprocessing PCIE Activity report to identify the appropriate synchronous I/O link.

Figure 208 on page 380 shows a Synchronous I/O Device Activity sample report.

For DASD devices actually used for synchronous I/O the Synchronous I/O Device Activity report shows detailed IBM zHyperLink activity data:

SYNCHRONOUS I/O DEVICE ACTIVITY															PAGE										
z/OS 3.1			SYSTEM ID SYSF			DATE 04/25/2023			INTERVAL 14.59.998																
			RPT VERSION 3.1 RMF			TIME 06.00.00			CYCLE 1.000 SECONDS																
TOTAL SAMPLES = 304			IODF = 00			CR-DATE: 04/25/2023			CR-TIME: 10.31.31			ACT: ACTIVATE													
						DEVID		ACTIVITY		RATE		AVG RESP TIME		AVG SYNCH I/O		%		%		%		%		--REJECTS--	
STORAGE GROUP	DEV	DEVICE	VOLUME	LCU	SYNCH	I/O	--	ASYNCH	--SYNCH I/O--	ASYNCH	TRANSFER	RATE	REQ	SUCCESS	LINK	BUSY	CACHE	MISS	READ	WRITE	READ	WRITE	READ	WRITE	
02180	33909	SYST10	001C	0.345	0.404	1.024	0.002	0.002	0.384	1.078	0.999	99.85	0.11	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	
02181	33909	SYST11	001C	0.702	0.491	0.073	0.001	0.001	0.112	0.500	0.456	99.94	0.00	0.01	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	
			LCU	001C	1.047	0.895	1.097	0.001	0.001	0.379	1.578	1.455	99.84	0.11	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	

Figure 207. Synchronous I/O Device Activity

SYNCHRONOUS I/O DEVICE ACTIVITY																	PAGE 18
z/OS 3.1				SYSTEM ID SYSB RPT VERSION 3.1 RMF				DATE 04/25/2023 TIME 15.00.00				INTERVAL 15.00.000 CYCLE 1.000 SECONDS					
TOTAL SAMPLES = 900				IODF = D4				CR-DATE: 08/15/2022				CR-TIME: 11.09.47				ACT: POR	
				- DEVICE ACTIVITY RATE -				-- AVG RESP TIME --				AVG SYNCH I/O					
STORAGE GROUP				-- SYNCH I/O -- ASYNCH I/O				-- SYNCH I/O -- ASYNCH I/O				TRANSFER RATE					
DEV	DEV	DEV	VOLUME	LCU	READ	WRITE	READ	WRITE	READ	WRITE	READ	WRITE	% REQ	% LINK	% CACHE	% --REJECTS--	
NUM	TYPE	SERIAL											SUCCESS	BUSY	MISS	READ WRITE	
DBNONRLS	05F70	3390A	NRLS17	0052	1250.84	0.000	796.469	0.023	0.000	1.237	1.765	0.000	99.47	0.06	0.32	0.16 0.00	
DATABASE	05F76	3390A	DB8136	0052	85.007	0.000	10.873	0.027	0.000	0.966	0.348	0.000	99.76	0.07	0.01	0.16 0.00	
DATABASE	05F77	3390A	DB2185	0052	49.694	0.000	6.365	0.027	0.000	1.059	0.204	0.000	99.81	0.08	0.00	0.11 0.00	
SUBSYS	05F7B	3390A	SUBS8A	0052	0.000	37.533	464.727	0.000	0.051	0.524	0.000	0.258	89.41	0.08	0.00	0.00 10.51	
SUBSYS	05F7C	3390A	SUBS6C	0052	21.629	0.000	342.280	0.025	0.000	0.774	0.089	0.000	99.02	0.05	0.01	0.93 0.00	
			LCU	0052	1410.84	37.533	1639.74	0.024	0.051	0.938	2.419	0.258	99.03	0.06	0.40	0.20 0.30	
...																	
...																	
	06200	33909		0055	0.000	959.665	0.711	0.000	0.052	0.202	0.000	6.954	99.98	0.00	0.00	0.00 0.02	
	06201	33909		0055	0.000	997.406	5.002	0.000	0.060	0.186	0.000	7.457	99.86	0.01	0.00	0.00 0.13	
	06202	33909		0055	0.000	1372.14	0.843	0.000	0.050	0.206	0.000	9.403	99.98	0.00	0.00	0.00 0.02	
	06203	33909		0055	0.000	1368.43	4.555	0.000	0.055	0.181	0.000	9.363	99.89	0.01	0.00	0.00 0.10	
	06204	33909		0055	0.000	706.263	1.337	0.000	0.055	0.198	0.000	5.135	99.94	0.00	0.00	0.00 0.05	
			LCU	0055	0.000	5403.90	12.448	0.000	0.054	0.188	0.000	38.31	99.93	0.00	0.00	0.00 0.06	
	0637C	33909		0056	0.000	277.060	0.568	0.000	0.051	0.172	0.000	1.697	99.95	0.00	0.00	0.00 0.05	
	06381	33909		0056	0.000	277.431	0.197	0.000	0.047	0.191	0.000	1.701	99.99	0.00	0.00	0.00 0.01	
			LCU	0056	0.000	554.490	0.764	0.000	0.049	0.177	0.000	3.398	99.97	0.00	0.00	0.00 0.03	

Figure 208. Synchronous I/O Device Activity

Communication equipment activity report

The reports for communication equipment, character reader devices, graphic devices, and unit record devices have the same format. The Communication equipment activity report is shown as example in .

The [Figure 209 on page 380](#) is as follows.

COMMUNICATION EQUIPMENT ACTIVITY																	PAGE	1	
z/OS 3.1				SYSTEM ID SYSA RPT VERSION 3.1 RMF				DATE 04/25/23 TIME 00.00.00				INTERVAL 30.00.000 CYCLE 1.000 SECONDS							
TOTAL SAMPLES = 1,800				IODF = 00				CR-DATE: 01/01/2021				CR-TIME: 00.00.00				ACT: ACTIVATE			
				DEVICE ACTIVITY RATE				AVG RESP TIME				AVG DISC CONN							
DEV	DEVICE	VOLUME	LCU					AVG	AVG	AVG	AVG	AVG	%	%	%	AVG	%		
NUM	TYPE	SERIAL						TIME	TIME	CMR	DB	INT	PEND	DISC	CONN	UTIL	RESV		
								DLY	DLY	DLY	DLY	DLY	DLY	DLY	DLY	DLY	ANY		
																	MT		
																	NOT		
																	RDY		
0C470				0012	0.016	.305	.000	.004	.000	.000	.128	.026	.150	0.00	0.00	0.0	100.0	0	
0C471				0012	0.017	.329	.000	.000	.000	.000	.141	.030	.158	0.00	0.00	0.0	100.0	0	
0C472				0012	0.016	.379	.000	.000	.000	.000	.174	.037	.169	0.00	0.00	0.0	100.0	0	
0C473				0012	1.803	.555	.000	.001	.000	.075	.110	.535	19.5	3.51	99.96	0.0	100.0	0	
0C474				0012	0.000	.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	100.0	0	
0C475				0012	0.000	.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	100.0	0	
0C476				0012	0.000	.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	100.0	0	
0C477				0012	1139.30	.373	.000	.000	.000	.000	.144	.030	.198	22.61	26.04	0.0	100.0	0	
			LCU	0012	1141.16	1.25	.000	.002	.000	.000	.144	.075	.229	3.27	15.75	0.0	100.0	0	

Figure 209. Communication Equipment Activity Report

COMMUNICATION EQUIPMENT ACTIVITY																	PAGE	1
z/OS 3.1			SYSTEM ID		SYSF		DATE 04/25/2023				INTERVAL 15.00.000							
			RPT VERSION 3.1		RMF		TIME 06.15.00				CYCLE 1.000 SECONDS							
TOTAL SAMPLES = 900			IODF = 00		CR-DATE: 04/25/2023				CR-TIME: 10.31.31				ACT: ACTIVATE					
			DEVICE		AVG		AVG		AVG		AVG		AVG		%			
			ACTIVITY		RESP		TQSO		CMR		DB		INT		PEND			
			RATE		DLY		DLY		DLY		DLY		DLY		DISC			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME		TIME		CONN			
			TIME		TIME		TIME		TIME		TIME							

Figure 210. Communication Equipment Activity Report

Magnetic tape device activity report

Magnetic tape devices are used for data storage in various systems, such as the DFSMS (Data Facility Storage Management Subsystem). The Magnetic tape device activity report is shown as example in [Figure 211 on page 381](#).

The following figure shows the Magnetic Tape Device Activity 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																
z/OS 3.1			SYSTEM ID SYSA			DATE 04/25/23			INTERVAL 30.00.000			PAGE 1				
			RPT VERSION 3.1 RMF			TIME 00.00.00			CYCLE 1.000 SECONDS							
TOTAL SAMPLES = 1,801 IODF = 00 CR-DATE: 01/01/2021 CR-TIME: 00.00.00 ACT: ACTIVATE																
DEV	DEVICE	VOLUME	LCU	DEVI	AVG	AVG	AVG	AVG	AVG	AVG	AVG	AVG	%	%	%	NUMBER
NUM	TYPE	SERIAL		ACTIVITY	RESP	IOSQ	CMR	DB	INT	PEND	DISC	CONN	DEV	DEV	DEV	OF
				RATE	TIME	TIME	DLY	DLY	DLY	TIME	TIME	TIME	CON	UTIL	RESV	MOUNTS
																TIME
																ALLOC
09000	3490L		00B7	OFFLINE												
09001	3490L		00B7	OFFLINE												
09002	3490L		00B7	OFFLINE												
09003	3490L		00B7	OFFLINE												
09004	3490L	TT9265	00B7	1169.16	.315	.000	.009	.000	.001	.159	.001	.155	18.14	18.26	0.0	0
09005	3490L	TT9278	00B7	1350.71	.313	.000	.008	.000	.001	.158	.001	.154	20.84	20.99	0.0	1
09006	3490L	TT9303	00B7	1168.83	.363	.000	.002	.000	.001	.136	.002	.225	26.32	26.56	0.0	* 6
09007	3490L	TT9345	00B7	1350.14	.365	.000	.002	.000	.001	.135	.006	.224	30.20	31.06	0.0	6
		LCU	00B7	5038.84	.339	.000	.005	.000	.001	.147	.003	.190	23.88	24.21	0.0	13

Figure 211. Magnetic Tape Device Activity Report

Field descriptions for all the device activity reports

Table 167. 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 five-digit hexadecimal device number that identifies a physical I/O device. The first digit represents the ID of the subchannel set to which the I/O device is physically configured.
DEVICE TYPE	The device type on which the data set resides.
NUMBER OF CYL	The DASD volume capacity (in cylinders).
VOLUME SERIAL	The volume serial number (for direct access and magnetic tape reports) of the volume mounted on the device at the end of the reporting interval.
PAV	The number of parallel access volumes (base and alias) which were available at the end of the reporting interval. If the number has changed during the reporting interval, it is followed by an '*'. If the device is a HyperPAV base device, the number is followed by an 'H', for example, 5.4H. The value is the average number of HyperPAV volumes (base and alias) in that interval. Average # of HPAV devices = $\frac{\text{Accumulated \# of HPAV devices}}{\text{Number of Samples}}$
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 z/VM or alternate VM guest system environment.

Table 167. Fields in the Device Activity Reports (continued)

Field Heading	Meaning
DEVICE ACTIVITY RATE	The rate at which start subchannel (SSCH) instructions to the device completed successfully. $\text{ACTV RATE} = \frac{\text{\# Successful SSCH Instructions}}{\text{Interval Time}}$ This formula applies to the activity rate measured during asynchronous I/O processing. For devices using suspended channel programs, resume I/O requests are included in the SSCH counts. Character 'S' appended to the DEVICE ACTIVITY RATE value of a device shown in the Direct Access Device Activity report indicates that the device performed synchronous I/O requests and that detailed synchronous I/O performance measurements for this device are available in the Synchronous I/O Device Activity report section. For easy comparison the Synchronous I/O Device Activity report lists the asynchronous I/O device activity rate calculated (ASYNCH I/O) in adjacent columns showing - the rate of successfully completed SYNCH I/O READ requests and - the rate of SYNCH I/O WRITE requests which completed successfully during the interval. The synchronous I/O activity rate is calculated as $\text{ACTV RATE} = \frac{\text{\# Successful Synch I/O read (respectively write) requests}}{\text{Interval Time}}$ 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 asynchronous 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. $\text{AVG ACT TIME} = \frac{\text{Device Active Time}}{\text{Measurement Event Count}}$ $\text{AVG RESP TIME} = \text{AVG ACT TIME} + \text{AVG IOSQ TIME}$ 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. For easy comparison the Synchronous I/O Device Activity report lists the asynchronous I/O average response time calculated (ASYNCH I/O) in adjacent columns showing - the average processing time (in milliseconds) per successful SYNCH I/O READ requests and - the average processing time (in milliseconds) per successful SYNCH I/O WRITE request.
AVG SYNCH I/O TRANSFER RATE	READ The number of megabytes per second read during synchronous I/O processing on the device. WRITE The number of megabytes per second written during synchronous I/O processing on the device.
% REQ SUCCESS	Percentage of synchronous I/O requests that completed successfully.
% LINK BUSY	Percentage of synchronous I/O requests that hit a link busy condition when trying to use a synchronous I/O link.
% CACHE MISS	Percentage of synchronous I/O read requests that hit a cache miss condition.

Table 167. Fields in the Device Activity Reports (continued)

Field Heading	Meaning
% REJECTS	READ The percentage of synchronous I/O read requests that were rejected for reasons other than a link busy condition or a read cache miss. WRITE The percentage of synchronous I/O write requests that were rejected for reasons other than a link busy condition.
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. $\text{AVG IOSQ TIME} = \frac{\text{Total IOSQ Time}}{\text{Start Subchannel Count}}$
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). $\text{AVG CMR DLY} = \frac{\text{Initial Command Response Time}}{\# \text{ I/O Operations Accepted on that Path}}$
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 $\text{AVG DB DLY} = \frac{\text{Device Busy Delay Time}}{\text{Measurement Event Count}}$
AVG INT DLY	The average interrupt delay time in units of milliseconds encountered for I/O requests to this device. For each I/O request, the time is measured from when the I/O operation is complete to when the operating system begins to process the status. $\text{AVG INT DLY} = \frac{\text{Device Interrupt Delay Time}}{\text{Measurement Event Count}}$
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. $\text{PEND TIME} = \frac{\text{Device Pending Time}}{\text{Measurement Event Count}}$

Table 167. Fields in the Device Activity Reports (continued)

Field Heading	Meaning
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. $\text{AVG DISC TIME} = \frac{\text{Device Disconnect Time}}{\text{Measurement Event Count}}$ 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. $\text{AVG CONN TIME} = \frac{\text{Device Connect Time}}{\text{Measurement Event Count}}$
% DEV CONN	The percentage of time during the interval when the device was connected to a channel path. $\% \text{ DEV CONN} = \frac{\text{Device Connect Time}}{\text{Interval Time}} * 100$
% 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. $\% \text{ DEV UTIL} = \left(\frac{(\text{CON} + \text{DISC}) / \text{PAV}}{\text{INT}} + \frac{\text{RSV}}{\text{N}} \right) * 100$ CON Device connect time DISC Device disconnect time PAV Number of parallels access volumes (base and alias); in case of non-PAV devices, PAV is set to 1 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. $\% \text{ DEV RESV} = \frac{\# \text{ Device-reserved Samples}}{\# \text{ Samples}} * 100$

Table 167. Fields in the Device Activity Reports (continued)

Field Heading	Meaning
AVG NMBR 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. $\% \text{ ANY ALLOC} = \frac{\# \text{ Samples when the Device was Allocated}}{\# \text{ Samples}} * 100$
% 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. $\% \text{ MT PEND} = \frac{\text{Counter for Mount-Pending Condition}}{\# \text{ Samples}} * 100$
%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. $\% \text{ NOT RDY} = \frac{\# \text{ Samples when the Device was not Ready}}{\# \text{ Samples}} * 100$
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 increments 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. $\text{AVG MOUNT TIME} = \frac{\# \text{ Samples Tape Mount was Pending} * \text{Interval}}{\# \text{ Mounts}}$ If the mount count or the sample count is zero, the result is zero. This field is reported only for magnetic tape devices.

Table 167. Fields in the Device Activity Reports (continued)	
Field Heading	Meaning
TIME DEVICE ALLOC	The total time the device was allocated during the interval, expressed in the form of HH:MM:SS. $\text{TIME DEVICE ALLOC} = \frac{\# \text{ Samples Tape Device was Allocated} * \text{Interval}}{\# \text{ Samples}}$ If the sample count is zero, the result is zero. This field is reported only for magnetic tape devices.

Device data incomplete or missing

Device data can be incomplete or missing because:

- Device not available during entire interval
- Device changed or deleted
- Hardware data not available
- Required data not available
- Device in use
- Average cannot be calculated

Device not available during entire interval

Data lines are included for each device that has been online at least once since Monitor I session initialization. However, data is not reported for devices that were offline at the end of the reporting interval, that came online during the interval, or that were affected by dynamic device reconfiguration during the interval. One of the following messages will appear in the data line indicating the reason why data was not presented:

NOW ONLINE

Brought online during this interval and still online at the end of the interval.

NOW OFFLINE

Taken offline during this interval and still offline at the end of the interval.

OFFLINE

Offline for the entire interval.

DEVICE DYNAMICALLY DELETED

Device dynamically deleted during the interval.

DEVICE DYNAMICALLY CHANGED

A device changed from static to dynamic during the interval, or a device deleted and a new device added with the same device number during the interval.

These messages indicate that the device data is incomplete and may present an inaccurate picture of device activity. If the hardware measurement data for the device is not available, the device data might be incomplete, even when a device has been online for the entire interval.

Device changed or deleted

If devices are changed or deleted from a storage group during the interval, RMF replaces the name of the storage group by ****CHGD**** in the STORAGE GROUP name column of the direct access device activity report. RMF does not provide summary lines for a storage group with ****CHGD**** in the STORAGE GROUP name column. Storage group names are still reported when devices are varied on or offline during the interval.

Hardware data not available

When hardware measurement data is not available, RMF can report values only for fields based on sampled data. It cannot report values for the fields based on hardware measurements; these fields are:

DEVICE ACTIVITY RATE
AVG RESP TIME
AVG DB DELAY
AVG IOSQ TIME
AVG PEND TIME
AVG DISC TIME
AVG CONN TIME
%DEV CONN
%DEV UTIL

Required data not available

When it cannot obtain the required data, RMF prints the status message **HARDWARE DATA UNAVAILABLE** in place of the data. Even if the channel measurement facility and the measurement block update facility are active and the device is online for the entire interval, valid hardware data might not be available. If RMF is unable to obtain valid hardware data, it prints the status message **NO H/W DATA**.

Device in use

RMF prints the status message **NO H/W DATA, DEVICE IN USE BY SYSTEM** when it cannot initialize the channel subsystem interface needed in order to gather the measurement data from the channel subsystem. This may occur for CTC devices that are being used by applications using protocols with never ending channel programs.

Average cannot be calculated

When RMF cannot calculate an average because a division by zero or a division overflow has occurred, four asterisks (****) appear in the field in place of the data.

Overflow condition occurred

Depending on the processor model you have, the hardware measurement data might be incomplete because of an overflow in the measurement timer. Any I/O request that exceeds that maximum time limit causes overflow. For example, chain scheduling, which the system uses when accessing page data sets or printing a SYSOUT data set, results in long channel programs and can cause timer overflow.

For shared DASD, pending times in excess of 8.3 seconds can occur due to RESERVE activity on the sharing system. Overflow conditions in pending time, however, are *not detected*.

Overflow in non-Monitor II reports

In Monitor II reports, the overflow counts are not reported. A value affected by overflow, however, is marked by an asterisk (*). In exception reports and summary reports, overflows are neither detected nor identified, because Monitor II device activity reports, as well as some exception reports and some fields in the summary report, are based on data the Monitor I session collects. The same inaccuracies apply to these reports.

Connect/Disconnect time overflow

Overflow conditions in connect time and disconnect time are detected by the hardware, counted by z/OS, and reported by RMF.

When a connect or disconnect time overflow occurs, RMF prints "HARDWARE DATA INCOMPLETE" on the line following the requests that caused the overflow. The data presented for those requests is most likely inaccurate because the values shown for connect time and disconnect time per request represent what remained after all the long-running requests were discarded. Those values, along with percent device connected, percent device utilized, and average response time, represent the lower bounds of what the

actual values might be. Because at least one request was discarded, all values must be larger than reported; how much larger, however, cannot be exactly determined.

On the same line with "HARDWARE DATA INCOMPLETE", RMF records the values of two counters: total requests that had timer overflow (in either or both timers); and total requests that had connect time overflow. The difference between timer overflow and connect time overflow gives the number of requests that had only disconnect time overflow. If the difference is large, it might indicate that difficulty in reconnecting to the channel is causing delays. Total requests with timer overflow is the difference between the start subchannel count and the measurement event count. RMF records connect time overflow separately.

The counts of requests that had overflow are intended to indicate how much data was lost. For example, if 1000 requests occurred in an hour and only one had overflow, the actual values are probably not much larger than the reported values. However, if the 999 requests measured were all short and all occurred within a short span of time, whereas the one long request lasted for 95% of the hour, the reported data is highly inaccurate. The values reported in the two counters do not take into account how many times a single request had an overflow.

The report of overflows for paging devices does not necessarily indicate a problem. Consult the Page Data Set Activity report and the Workload Activity report to determine whether or not paging delays are a problem. If they are, the device data can be used in conjunction with the two reports to analyze the problem.

Pending time overflow

Because all overflows for pending time are lost, RMF does not accurately report certain shared DASD delays. For example, a request delayed for 18 seconds overflows twice; 16.6 seconds are lost. To RMF, the delay appears to be only 1.4 seconds. Therefore, the AVERAGE PENDING TIME and the AVERAGE RESPONSE TIME values are extremely inaccurate.

For requests with extremely long delays, the missing interrupt handler (MIH) halts the request and reschedules it periodically. MIH estimates the amount of pending time, based on the MIH interval, and adds it to the value RMF reports. Therefore, pending time is lost only for requests that take longer than 8.3 seconds and less than 1.5 times the MIH interval. To increase the accuracy of AVERAGE PENDING TIME and AVERAGE RESPONSE TIME, decrease the MIH interval. An interval of four seconds will ensure that no pending time is lost. However, some performance penalty does occur because of the four-second interval.

Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*. The following table shows the overview condition names for the Overview report.

<i>Table 168. Overview names in the DASD Activity Report</i>	
Field Heading or Meaning	Overview Name
NUMBER OF CYL	DVCAP
DEVICE ACTIVITY RATE	DART
AVG RESP TIME	DRTAVG
AVG IOSQ TIME	DQTAVG
AVG CMR DLY	CMRDL
AVG INT DLY	INTDL
AVG DB DLY	DBDL
AVG PEND TIME	DPTAVG
AVG DISC TIME	DDTAVG
AVG CONN TIME	DCTAVG

Table 168. Overview names in the DASD Activity Report (continued)

Field Heading or Meaning	Overview Name
% DEV UTIL	DVUTL
% DEV RESV	DR
% MT PEND	DMTPEND
SYNC I/O DEVICE READ ACTIVITY RATE	DARSR
SYNC I/O DEVICE WRITE ACTIVITY RATE	DARSW
AVG SYNC I/O READ RESP RATE	DRTSR
AVG SYNC I/O WRITE RESP RATE	DRTSW
SYNC I/O READ XFER RATE	DTRSR
SYNC I/O WRITE XFER RATE	DTRSW

Data inaccuracies in duration report

When you request a duration report for I/O device activity, the identifier of each I/O device that had any VARY activity during the duration interval is followed by a single asterisk (*). The data recorded for such a device is partial; that is, no data was collected during one or more of the measurement intervals included in the duration report because the device was varied during a measurement interval. The calculations the Postprocessor performs to generate a duration report make no adjustments for RMF measurement intervals when no data was collected for a device. Thus, the data for a device that moved online or offline might appear to be inconsistent. For example, the percentages reported for the QLENGTH DISTRIBUTION field do not equal 100% when data for the device is partial.

At the beginning of a Monitor I session, all devices that are online are known to RMF; thus, RMF creates an entry in the type 74 SMF record for each online device that the user requested RMF to monitor. In contrast, any device that is offline at the beginning of the session is unknown to RMF, and no entry in the SMF record is built. When an unknown device is brought online, it becomes known to RMF, and an entry in the SMF record is then built for the device.

When a duration report combines data collected during two or more separate Monitor I sessions, the status of a device can change (for example, from offline to online or from unknown to online). Thus, the following conditions can occur:

- A device was known to RMF but offline during the first Monitor I session and online during subsequent Monitor I sessions but did not change during a measurement interval included in the duration report. In this case, the data is partial and the device identifier is followed by an asterisk.
- A device was offline for one or more measurement intervals and unknown to RMF during all other measurement intervals included in the duration report. In this case, the device identifier is followed by an asterisk, OFFLINE appears in the first data field, and no data is formatted for the device.
- A device was unknown to RMF during a Monitor I session and online for subsequent Monitor I sessions but did not change during a measurement interval included in the duration report. In this case, the data is partial and the device identifier is followed by an asterisk.

The following conditions can occur for storage group reporting:

- The STORAGE GROUP field shows **CHGD** for the volume if the storage group name changes in an SMF record for the duration period.
- The STORAGE GROUP field shows **CHGD** and the device identifier is followed by an asterisk if a volume is not reported in all SMF records of the duration and has changed the storage group name at least once.
- STORAGE GROUP DATA NOT AVAILABLE is reported between the TOTAL SAMPLES field and the report headings if the storage management subsystem is not available in one of the reports during the duration period.

- SMS INTERFACE ERROR, NEW STORAGE GROUP INFORMATION CANNOT BE OBTAINED is reported between the TOTAL SAMPLES = field and the report headings if a system-managed storage interface error occurs in one of the reports during the duration period.

DOMINO - Lotus Domino Server report

The Domino Server family is an integrated messaging and Web application software platform. The Domino Server enables Web clients to communicate with Notes servers.

The Lotus Domino Server report provides information about the activities of a server. The information can be used to analyze the activities of the server in case of problems.

How to request this report

The Postprocessor is using SMF records type 108 as input for the Lotus Domino Server report. These records are not gathered by an RMF monitor, but are written by Domino servers. See *z/OS Resource Measurement Facility User's Guide* for details.

To produce this report, specify

```
REPORTS(DOMINO)
```

Contents of the report

The report consists of two parts:

- Lotus Domino Server Summary

The summary contains one line for each server which is part of the report.

- Lotus Domino Server Details

This part consists of the following sections:

- Definition data (provided by record type 108-3)
- Performance data (provided by record type 108-3)
- Load data (provided by record type 108-1)

LOTUS DOMINO SERVER SUMMARY											
z/OS 3.1	SYSTEM ID LN21	DATE 04/25/2023	INTERVAL 05.00.000				PAGE 1				
	RPT VERSION 3.1 RMF	TIME 18.40.00									
SERVER NAME	AVAILABLE HHH.MM.SS	---- USERS CONNECTED	----- ACTIVE	TASKS	TRANSACTION RATE	ASYNC READS	I/O RATE WRITES	MAIL RATE DELIVERED	RATE SENT	SMTP RATE READS	RATE 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 212. Lotus Domino Server Report - Summary

LOTUS DOMINO SERVER DETAILS											PAGE 2
z/OS 3.1		SYSTEM ID LN21		DATE 04/25/2023		INTERVAL 05.00.000					
		RPT VERSION 3.1 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	2034	MAX	76	THRESHOLD	0	MAX	44032		
CURRENT	2034	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 -----			----- TOP-10 BY R/T -----			----- TOP-10 BY R/T -----		
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 213. Lotus Domino Server Report - Details

Table 169. 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.
ASYNCH I/O RATE - READS	Rate of asynchronous reads.
ASYNCH I/O RATE - WRITES	Rate of asynchronous writes.
MAIL RATE - DELIVERED	Rate of Domino mail messages delivered to local users.
MAIL RATE - SENT	Rate of Domino mail messages sent to other servers.
SMTP RATE - READS	Rate of SMTP messages received from other servers.
SMTP RATE - WRITES	Rate of SMTP messages sent to other servers.

Table 170. 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.

Table 170. Fields in the Domino Server Details Report (continued)	
Field Heading	Meaning
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	
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	

Table 170. Fields in the Domino Server Details Report (continued)	
Field Heading	Meaning
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.

EADM - EADM Activity Report

The EADM Activity Report provides statistics and performance measurements on Extended Asynchronous Data Mover (EADM) activity. EADM activity encompasses Storage Class Memory (SCM) and EADM compression activity.

How to request this report

If the currently active SMFPRMxx parameter settings indicate that SMF record type 74 subtype 10 is to be collected, then RMF Monitor III gathers the data required for the EADM Activity Report into this SMF record.

To produce this report, specify

```
REPORTS (EADM)
```

Note: This report was formerly called SCM Activity Report. The SCM keyword is still accepted and has the same meaning as EADM.

This single-system report is only available in XML output format. Therefore, you need to specify the XPRPTS ddname in your Postprocessor job. How to work with Postprocessor XML reports in *z/OS Resource Measurement Facility User's Guide* provides all required information on how to produce and view XML reports.

Example URL for the DDS API

<http://ddshost:8803/gpm/rmfpp.xml?reports=EADM>

Contents of the report

The Extended Asynchronous Data Mover (EADM) Activity Report consists of three segments:

1. Device/subchannel level information. The EADM device summary segment provides the rate of start subchannel (SSCH) instructions for all EADM devices together with response time statistics consisting of pending, IOP queue and initial command response time.
2. Compression activity information. This segment provides request rates, throughput, and ratios of EADM compression and decompression. These values cover asynchronous compression and decompression activity by EADM. Synchronous compression and decompression is not reported.
3. Storage Class Memory activity information. For each Flash Express adapter, the report provides measurements at both the LPAR and CPC level. The total number of requests, the rate at which requests are processed by the adapter, the rate at which data units were read and written, the average response and IOP queue time is displayed.

Note: If the hardware supports Virtual Flash Memory, Flash Express cards are simulated by cache, and SCM activity is reported in one report line.

RMF Postprocessor Interval Report [System SYSE] : Extended Asynchronous Data Mover Activity Report													
RMF Version: z/OS 3.1 SMF Data: z/OS 3.1 Start : 09/30/2021-15.44.34 End : 09/30/2021-15.59.33 Interval : 15:00:000 minutes													
Device/Subchannel Summary													
Total Number of SSCH : 0 SSCH Rate : 0.00 Avg Function Pending Time : 0.000 Avg Initial Cmd Response time : 0.000													
Compression Activity													
Compression Request Rate : 12.30 Compression Throughput : 65321 Compression Ratio : 45.78 Decompression Request Rate : 23.17 Decompression Throughput : 43216 Decompression Ratio : 0.67													
Storage Class Memory Activity													
Card ID	Util% (LPAR)	Util% (Total)	Read B/Sec (LPAR)	Read B/Sec (Total)	Write B/Sec (LPAR)	Write B/Sec (Total)	Request Rate (LPAR)	Request Rate (Total)	Avg Respons e Time (LPAR)	Avg Respons e Time (Total)	Avg IOP Queue Time (Total)	Request s (LPAR)	Request s (Total)
VFM	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.000	0.000	0.000	0	0

Figure 214. Extended Asynchronous Data Mover (EADM) Activity Report

Table 171. Fields in the EADM Activity Report	
Field Heading	Meaning
Device/Subchannel Summary	
This section provides summary information about the extended asynchronous data mover (EADM) devices or subchannels. EADM subchannels are similar to I/O subchannels in a way that I/O instructions can be issued. However, they do not have channel paths or device numbers assigned, and they are not defined in the I/O configuration. They are created automatically during IPL.	
Total Number of SSCH	The total number of SSCH instructions to all EADM devices in the report interval.
SSCH Rate	The number of SSCH instructions to all EADM devices per second.
Avg Function Pending Time	The average function pending time across all EADM devices in milliseconds. This is similar to function pending time for traditional I/O devices, which is the amount of time between when the SSCH is issued and the first command in the channel program is accepted. $\text{AVG} = \frac{\text{Sum(Function Pending Time)}}{\text{Measurement Event Count}}$

Table 171. Fields in the EADM Activity Report (continued)	
Field Heading	Meaning
Avg IOP Queue Time	The average IOP queue time across all EADM devices in milliseconds. This is unique to EADM devices. It represents the amount of time the request is not accepted by the adapter because it would exceed its maximum capacity. For a particular I/O request, this may occur multiple times. $\text{AVG} = \frac{\text{Sum(IOP Queue Time)}}{\text{Measurement Event Count}}$
Avg Initial Cmd Response Time	The average initial command response time across all EADM devices in milliseconds. This is the time from when the first command does not immediately proceed to execute until the successful start of execution at the EADM resource part. $\text{AVG} = \frac{\text{Sum(Initial Command Response Time)}}{\text{Measurement Event Count}}$
Compression Activity	
Compression Request Rate	The number of compression requests per second.
Compression Throughput	The number of megabytes compressed per second.
Compression Ratio	The ratio between input and output bytes compressed within this interval.
Decompression Request Rate	The number of decompression requests per second.
Decompression Throughput	The number of megabytes decompressed per second.
Decompression Ratio	The ratio between input and output bytes decompressed within this interval.
Storage Class Memory Activity	
Card ID	The identifier of the flash adapter card. VFM is reported if the hardware has configured Virtual Flash Memory.
Following fields are displayed at a system-wide level (Total) and for the current LPAR whereby IOP Queue Time is only available at the total level.	
Util(%)	The average utilization of the flash card during the interval as reported by the EADM measurement facility. The average utilization of Virtual Flash Memory is reported as the percentage of the time spent on System Assist Processors (SAP) for EADM processing compared to the total available SAP time in this reporting interval.
Read(B/Sec)	Bytes read per second.
Write(B/Sec)	Bytes written per second.
Request Rate	The requests processed per second.
Requests	The total number of requests.
Avg Response Time	The average response time per request in milliseconds. The response time represents the CHSC execution time and does not include pending, IOP queue and initial command response time.
Avg IOP Queue Time	The average IOP queue time per request in milliseconds.

ENQ - Enqueue Activity report

The Enqueue Activity report provides information about resources that periodically build up queues of one or more requestors waiting to use the resource. Contention is reported for those resources where access is controlled by jobs that issue ENQ and DEQ macro instructions. RMF records related resource contention status changes signalled by GRS. When contention detail data (such as resource owner or numbers of exclusive/shared waiters) is passed with a signalled contention, RMF attributes these data to queue length buckets and reports individually on related contention measurements. To complement the

picture about resource contentions, the portion of contention change events about which RMF does not obtain detail information is visible as percentage of the total number of contention change events.

Using the information given in the report

Because the amount of time that a requestor must spend waiting for a resource can seriously affect system throughput, the information in this report can be very helpful in locating resources that consistently cause bottlenecks.

Once you have defined a critical resource, such as a serially-reusable resource that can be requested on either an exclusive or shared basis, your installation can improve the situation in a variety of ways. You could change the hardware configuration to release device bottlenecks, change data set placement, or reschedule jobs to improve throughput, or re-specify the installation tuning parameter ERV (enqueue residence value) to give more processor time to the holder of the resource.

The information in the detail report can help you to balance your workload to minimize resource contention.

How to request this report

To gather data for this report, specify as a Monitor I gatherer option

```
ENQ(SUMMARY | DETAIL[,majorname,[minorname]])
```

To produce this report, specify

```
REPORTS(ENQ)
```

Note: The ENQ report is only available as an interval report, not as a duration report.

This report is also available in XML output format. [How to work with Postprocessor XML reports in z/OS Resource Measurement Facility User's Guide](#) provides all required information on how to produce and view XML reports.

Example URL for the DDS API

```
http://ddshost:8803/gpm/rmfpp.xml?reports=ENQ
```

Different report levels

The contents of the report depends on the gathering options:

- Summary report - ENQ(SUMMARY)
- Detail report - ENQ(DETAIL) or ENQ(DETAIL,majorname [,minorname])

The **Summary Activity report** includes:

- All resources for which contention has occurred during the reporting interval.
- A description of the contention time for each resource.
- A queue length distribution and average queue length for each resource.
- Information on the type of requests made (either exclusive or shared).
- The total number of enqueue contention events that occurred.

An enqueue contention event is defined as the period from the time when the resource first has contention until the resource no longer has contention.

- The total number of contention status change events.

Contention status change events are events such as an incident where at least one waiter gets queued for a given resource, or an incident where the number of waiters or the contention owner changes, and also the contention-completion event at the end of an enqueue contention event.

- The percentage of status change events that did not provide contention detail data and therefore can not be attributed uniquely to queue length buckets.

A resource for which contention is still occurring at the end of the interval will be indicated by an asterisk following the TOT field, which is under the CONTENTION TIME field.

The **Detail Activity report** shows several lines of data for all resources for which contention occurs.

- The total number of jobs that own the resource and the names of one or two jobs that own the resource
- The total number of jobs that are waiting for the resource and the names of one or two jobs that are waiting for the resource.
- The identifier of the system on which the job is running following each job name
- An E if the request is exclusive or an S if the request is shared

RMF selects the job names shown in the detail report during the period of maximum contention in the interval by determining the longest contention event in the interval. For that event RMF reports the owners and waiters at the point when the event queue is the longest.

When there are several occurrences of the same length queue, the latest queue is reported. RMF reports the job names that were active at maximum contention even though those jobs might have been processed and flushed from the system by the time the contention no longer exists.

You can request data for a specific resource by specifying a *major* name, with or without a *minor* name. Various combinations of the reporting options can give you a complete picture of both critical resources and the jobs that are impacting system throughput by monopolizing a specific resource.

Contents of the report

The data fields for the summary and detail reports are identical, with one exception: the job names causing maximum contention are printed only when the detail level is requested. Therefore, the fields are discussed only once, and the field that is provided only at the detail level is noted. The data fields are preceded by ENQUEUE SUMMARY ACTIVITY for a summary report or ENQUEUE DETAIL ACTIVITY for a detail report.

ENQUEUE ACTIVITY																								
z/OS 3.1				SYSTEM ID SYS1				DATE 04/25/2023				INTERVAL 14.59.946				PAGE 1								
				RPT VERSION 3.1 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 - --- CONTENTION ---									
MAJOR MIN MAX TOT AVG					TOT OWN ----- WAIT ---					1 2 3 4+ LENGH					-EXCL- -SHARE- EVENT --STAT CHNG--									
MINOR					TOT NAME TOT NAME										MIN MAX MIN MAX TOTAL TOTAL %NODET									
					SYSNAME					SYSNAME														
SYSZJES2																								
SJB.2087F970																								
0.000 0.000 0.000 0.000					1					*MASTER*(E) 1 WEID					(S) 100 0.0 0.0 0.0 1.00					0 0 1 1 2 17 3.1				
					RMFG																			

Figure 215. Enqueue Detail Activity Report

Table 172. Fields in the Enqueue Activity Report	
Field Heading	Meaning
GRS MODE	Shows the GRS mode in which the system is running.
NAME (MAJOR MINOR)	The name of a resource that has one or more requestors waiting. The major name is one to eight characters in length; the minor name can be from 1 to 255 characters, but only 44 characters will be printed. When the name exceeds 44 characters, it is truncated in the report. An asterisk (*) following the resource name indicates that it has been truncated. A resource with a scope of "SYSTEMS" will be followed by (SYSTEMS); a resource with a scope of "SYSTEM" will have no indication; and a resource with a scope of "STEP" will not be included in the report.

Table 172. Fields in the Enqueue Activity Report (continued)

Field Heading	Meaning
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: $\text{AVG CONT TIME} = \frac{\text{Contention Time for the Resource}}{\# \text{ Contention Events}}$
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. The reported counts refer to the period of maximum contention for a resource in the RMF reporting interval. 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 status change events during the interval when the number of requestors queued to the resource was one, two, three, four or more. The samples are taken for each contention status change where RMF receives contention detail data relatable uniquely to one of these four queue length buckets. Examples for a contention status change event are the change of the contention owner or the number of waiters. At each such sample, an accumulator for the observed length is updated. At the end of the measurement interval, the percentage for each queue length is computed. The calculation used for each queue length is: $\% \text{ QLEN} = \frac{\text{Accumulator for that Queue Length}}{\text{Sum of Accumulators for all Queue Lengths}} * 100$
AVG Q LENGTH	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. $\text{AVG Q LENGTH} = \frac{\# \text{ Requestors Waiting}}{\text{Sum of Accumulators for all Queue Lengths}}$
-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.
CONTENTION EVENT TOTAL	The total number of resource contention events that occur during the measurement interval. This is the total number of periods each starting from the time when the resource has contention until the resource no longer has contention.
CONTENTION STAT CHNG	TOTAL The total number of contention status change events. This includes status change events relatable uniquely to specific queue length buckets about which RMF received contention information, as well as those status change events about which RMF has no specific contention detail data. %NODET The percentage of contention status change events which did not provide contention detail data.

Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*. The following table shows the overview condition names for the Overview report.

Table 173. Overview names in the Enqueue Activity Report	
Field Heading or Meaning	Overview Name
CONTENTION TIME - TOT	ENQT
CONTENTION TIME - AVG	ENQAVG
CONTENTION TIME - MAX	ENQMAX
TOTAL EVENT	ENQNE
% CONTENTION STATUS CHANGE EVENTS NO DETAIL	ENQPNOD

Messages

During the measurement of enqueue activity, RMF can encounter situations when no reporting can be done. When such a situation occurs, RMF replaces the report with a message describing the reason no report could be formatted. The messages are:

NO CONTENTION OCCURRED

Explanation: During the interval, no contention activity occurred for the resource or resources being measured. Enqueue activity measurement and reporting continue as specified. This message would appear most frequently when you are requesting the enqueue activity report for a specific resource.

TABLE FULL - USE SPECIFIC NAME OR SHORT INTERVAL

Explanation: During the RMF interval, a period of such high contention activity occurred that the internal working table was filled. As a result, no further enqueue reporting could be done for the interval. The interval report contains the data gathered before the internal table was filled and followed by the message. Subsequent interval reports might not include complete data. Enqueue activity measurement and reporting resume at the start of the next interval.

When the message occurs, you could reduce the length of the RMF interval, or, if you want to ensure that the contention activity for a specific resource is reported, you can request enqueue activity reporting for the specific critical resource.

TERMINATE DUE TO DATA EVENT ERROR - TRY RERUN

Explanation: During the interval, the enqueue measurement routines encountered invalid data while processing a contention event. All enqueue measurement activity is terminated for the session; that is, the recovery from the error includes modifying the enqueue activity option to NOENQ. Because the error encountered might not be a permanent error, you can modify the session options to re-specify enqueue measurement and reporting. If the message occurs again and there are no other indicators of a system problem, report the message to the RMF license holder at your installation.

BAD CPU CLOCK OCCURRED - FIX CPU CLOCK AND RERUN

Explanation: During an interval, the enqueue activity measurement routines detected an error in the CPU clock function. All enqueue measurement activity for the session is terminated; that is, the recovery from the error includes modifying the enqueue activity option to NOENQ. Note that this message is probably one of many indicators that there is a problem with the CPU clock. After the clock has been fixed, re-IPL the system and run the session again, specifying the enqueue activity measurements that you require.

ESS - Enterprise Disk Systems report

The Enterprise Disk Systems report provides measurements about the activities of an enterprise disk system. RMF monitors the activity on an enterprise disk system independently from the source of the activity. Activity may be caused by the z/OS system on which RMF is running or from any other system

using the enterprise disk system. You can use the data contained in this report for checking your current disk configuration, for bottleneck analysis and for capacity planning.

If the Monitor I data gatherer is set up to collect all available data, the report may contain the sections described in the following:

- ESS Link Statistics
- ESS Synchronous I/O Link Statistics
- ESS Extent Pool Statistics¹
- ESS Rank Statistics¹

ESS Link Statistics

For each adapter of an ESS, this section contains statistics about the occurred I/O operations. One adapter supports one or more type of I/O (link type). The following link types are reported:

- ECKD read and write¹
- SCSI read and write
- PPRC send and receive

For each link type, this section provides the average number of transferred bytes and the average number of operations per second as well as their average response time. The I/O intensity shows the utilization of the adapter during the report interval. Use this section for analysis of the external link usage and for capacity planning of the peer-to-peer remote copy (PPRC) links.

ESS Synchronous I/O Link Statistics

For each adapter of an ESS, this section provides statistics about the occurred synchronous I/O operations on defined IBM zHyperLinks. On basis of cache read, cache write and NVS write operations, it shows per synchronous I/O link

- the average number of operations per second
- the rate of bytes transferred per synchronous I/O operation
- the average response time per synchronous I/O operation
- the percentage of synchronous I/O operations that performed successfully

The ESS Synchronous I/O Link Statistics section is only part of the ESS report of a storage controller if at least one IBM zHyperLink is defined on it. Use this section for performance analysis of IBM zHyperLink activity.

ESS Extent Pool Statistics

This section provides capacity and performance information about allocated disk space. For each extent pool, it shows the real and virtual capacity, the number of real and virtual extents and the number of conversions from a virtual into a real extent and vice versa. Such a conversion occurs when an application writes to a virtual extent. A conversion from a real into a virtual extent occurs if an extent is freed or migrated.

Use this section to check the available disk capacity and if required, change the capacity. For example, if you use the FlashCopy function of the enterprise disk system to create an instant point-in-time backup copy of your application data or data base, then you can read from the conversion information provided in this section, whether your provisioning strategy is successful. If it is dissatisfactory, you may decide to add real storage to the subsystem.

¹ available for the IBM TotalStorage DS family

ESS Rank Statistics

This section provides activity statistics about read and write operations in each rank of an extent pool. It also shows the number of arrays and the array width of all ranks. These values show the current configuration. The wider the rank, the more performance capability it has. By changing these values in your configuration, you can influence the throughput of your work. Use this section to detect and resolve performance problems and also for disk space capacity planning.

How to request this report

The default option for Monitor I data gathering is NOESS. Therefore, you must specify the ESS Monitor I gatherer option if you want to get data for this report.

To produce this report, specify

```
REPORTS(ESS)
```

This report is also available in XML output format. [How to work with Postprocessor XML reports in z/OS Resource Measurement Facility User's Guide](#) provides all required information on how to produce and view XML reports.

Example URL for the DDS API

```
http://ddshost:8803/gpm/rmfpp.xml?reports=ESS
```

Contents of the report

Depending on your Monitor I gatherer options and available data, this report consists of up to four sections. The following fields are common for all sections:

Table 174. Common Fields in the Enterprise Disk Systems Report	
Field Heading	Meaning
SERIAL NUMBER	Serial number of the primary control unit.
TYPE-MODEL	ESS type and model.
CDATE	Date when the cache interval started.
CTIME	Time when the cache interval started.
CINT	Cache interval time. In interval reports, the format is <i>mm.ss</i> , while in duration reports the format is <i>hh.mm.ss</i> .
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.	

ESS LINK STATISTICS										PAGE	1
z/OS 3.1		SYSTEM ID VSL1		DATE 09/30/2021		INTERVAL 15.00.000					
		RPT VERSION 3.1 RMF		TIME 08.30.00		CYCLE 1.000 SECONDS					
SERIAL NUMBER 0000002471		TYPE-MODEL 002107-922		CDATE 09/30/2021		CTIME 08.29.30		CINT 15.00			
-----ADAPTER-----		--LINK TYPE--		BYTES		OPERATIONS		RESP TIME		I/O	
SAID TYPE				/SEC		/OPERATION		/OPERATION		INTENSITY	
0004	FIBRE 2Gb	ECKD READ	162.1K	13.7K	11.8	0.3	3.9				
		ECKD WRITE	2.4M	26.5K	92.5	0.8	76.2				

							80.1				
0011	FIBRE 1Gb	NO DATA TO REPORT OR ZERO									
0024	FIBRE 2Gb	SCSI READ	156.0K	13.9K	11.2	0.3	3.6				
		SCSI WRITE	2.5M	26.5K	93.2	0.8	76.8				

							80.4				
0088	FIBRE 2Gb	PPRC SEND	8.5M	50.4K	169.2	16.1	2729.9				
		PPRC RECEIVE	0.0	0.0	0.0	0.0	0.0				

							2729.9				

Figure 216. ESS Link Statistics

RMF issues the informational message 'NO DATA TO REPORT OR ZERO' if the counters for all link types (both read or write) return 'zero'. This happens in the following cases: either there was no ESS activity in the report interval or the ESS did not deliver any data.

Table 175. Fields in the ESS Link Statistics	
Field Heading	Meaning
ADAPTER	Specifies the channel adapter: SAID system adapter identifier TYPE adapter type, for example, FIBRE 2Gb; "Undefined", if RMF could not determine the type
LINK TYPE	Type of I/O operation performed by the adapter, which can be one of the following: - ECKD READ or ECKD WRITE: designates extended count key data I/O - SCSI READ or SCSI WRITE: designates small computer system interface I/O - PPRC SEND or PPRC RECEIVE: designates peer-to-peer remote copy traffic
BYTES /SEC	The average number of bytes transferred per second for all operations of the indicated link type during the reporting interval.
BYTES /OPERATION	The average number of bytes transferred per operation for all operations of the indicated link type during the reporting interval.
OPERATIONS /SEC	The average number of operations of the indicated link type per second during the reporting interval.
RESP TIME /OPERATION	The average response time of operations of the indicated link type during the report interval. This is the entire time from sending out a data block until the notice of receipt from the receiver arrives. This value is measured in milliseconds.
I/O INTENSITY	The portion of the reporting interval during which an adapter was active. It is the product of OPERATIONS/SEC times RESP TIME/OPERATIONS. The I/O intensity is provided as a total for each adapter as well as for each link type. It is measured in milliseconds/second. That is, a value of 1000 for a link type indicates that this link was busy all the time during the report interval. On a Fiber Channel, multiple data blocks can be sent concurrently without waiting for the notices of receipt. Therefore, an I/O intensity greater than 1000 ms per second for an adapter may occur if such concurrent operations had been active.

ESS SYNCHRONOUS I/O LINK STATISTICS													PAGE	1
z/OS 3.1		SYSTEM ID SYSX			DATE 09/30/2021			INTERVAL 10.00.004						
		RPT VERSION 3.1 RMF			TIME 04.50.00			CYCLE 1.000 SECONDS						
SERIAL NUMBER	00000BBT62	TYPE-MODEL	002107-986	CDATE	09/30/2021	CTIME	04.50.01	CINT	10.02					
-----CACHE READ OPERATIONS-----														
SIID	-----LINK TYPE-----	OPS	BYTES	RTIME	%SUCC	OPS	BYTES	RTIME	%SUCC	OPS	BYTES	RTIME	%SUCC	
		/SEC	/OP	/OP		/SEC	/OP	/OP		/SEC	/OP	/OP		
0080	Optical PCIe GEN3 8	60.6	3.6K	9.8	87.8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
0081	Optical PCIe GEN3 8	NO DATA TO REPORT OR ZERO				0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
0180	Optical PCIe GEN3 8	41.5	3.4K	9.7	82.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	

Figure 217. ESS Synchronous I/O Link Statistics

Table 176. Fields in the ESS Synchronous I/O Link Statistics Report	
Field Heading	Meaning
SIID	Synchronous I/O link (IBM zHyperLink) identifier.
LINK TYPE	Combined presentation of synchronous I/O link type (for example, Optical PCIe), link speed (for example, GEN3) and link width by number of PCIe lanes (for example, 8).
CACHE READ OPERATIONS	OPS/SEC Number of synchronous I/O cache read operations per second. BYTES/OP Average number of bytes read per synchronous I/O cache read operation. RTIME/OP Average response time (in milliseconds) per synchronous I/O cache read operation. %SUCC Percentage of synchronous I/O cache read operations that performed successfully.
CACHE WRITE OPERATIONS	OPS/SEC Number of synchronous I/O cache write operations per second. BYTES/OP Average number of bytes written per synchronous I/O cache write operation. RTIME/OP Average response time (in milliseconds) per synchronous I/O cache write operation. %SUCC Percentage of synchronous I/O cache write operations that performed successfully.
NVS WRITE OPERATIONS	OPS/SEC Number of synchronous I/O write operations to the NVS node per second. BYTES/OP Average number of bytes written per NVS write operation. RTIME/OP Average response time (in milliseconds) per NVS write operation. %SUCC Percentage of NVS write operations that performed successfully.

ESS EXTENT POOL STATISTICS												
z/OS 3.1		SYSTEM ID MES9			START 08/06/2024-10.00.00		INTERVAL 001.00.00					
		RPT VERSION 3.1 RMF			END 08/06/2024-11.00.00		CYCLE 1.000 SECONDS					
		TYPE-MODEL 002107-A01			CDATE 08/06/2024		CTIME 10.30.00		CINT 00.30.00			
--EXTENT POOL--												
ID	TYPE	CAPACITY	EXTENTS	CONVERSIONS	CAPACITY	EXTENTS	CONVERSIONS	PHYS TOT	PHYS USED	LOG TOT	LOG USED	RATIO
0000	CKD 16Mb	101000	6003799	1	11091	587823	0	1221703	76585	5596800	9911	0.13
0001	CKD 16Mb	101000	6003799	0	9403	498359	0	1221703	76161	5596800	1802	0.02
0002	CKD 16Mb	8665	515115	0	5760	305280	0	473820	901	473820	901	1.00
0003	CKD 16Mb	8665	515115	0	3291	174423	0	473820	636	473820	636	1.00

Figure 218. ESS Extent Pool Statistics

The **ESS Extent Pool Statistics** section presents overview information on the defined disk capacity of extent pools.

Table 177. Fields in the ESS Extent Pool Statistics	
Field Heading	Meaning
EXTENT POOL	Pool of allocation units for logical volumes. ID extent pool identifier TYPE extent pool type, for example, FIBRE 1Gb or CKD 1Gb
REAL CAPACITY	The capacity of physical storage in gigabytes for real extents in an extent pool. This is available capacity for the operating system.
REAL EXTENTS	Number of real extents in an extent pool. A discrete number of extents can be used to create volumes.
REAL CONVERSIONS	Number of real extent conversions. A virtual extent is converted to a real extent, if an application writes to that extent.
VIRTUAL CAPACITY	Virtual extent pool capacity. The capacity of physical storage in gigabytes for virtual extents in an extent pool.
VIRTUAL EXTENTS	Number of virtual extents in an extent pool. A virtual extent is an extent that has not yet been converted to a real extent.
VIRTUAL CONVERSIONS	Number of virtual extent conversions. A real extent gets converted via migration or by freeing an extent.
COMPRESSION PHYS TOT	Total physical capacity.
COMPRESSION PHYS USED	Used physical capacity.
COMPRESSION LOG TOT	Total logical capacity.
COMPRESSION LOG USED	Used logical capacity.
COMPRESSION RATIO	Compression Ratio: Used logical capacity / Used physical capacity

ESS RANK STATISTICS																			
z/OS 3.1		SYSTEM ID SYSE		DATE 09/30/2021		INTERVAL 30.00.026		PAGE 8											
		RPT VERSION 3.1 RMF		TIME 13.30.00		CYCLE 1.000 SECONDS													
SERIAL NUMBER	00000DKA61	TYPE-MODEL	002107-961	CDATE	09/30/2021	CTIME	13.30.00	CINT	30.00										
--EXTENT POOL--		ADAPT		READ OPERATIONS		WRITE OPERATIONS		ARRAY		MIN	RANK	RAID							
ID	TYPE	RRID	ID	OPS	BYTES	BYTES	RTIME	OPS	BYTES	BYTES	RTIME	SSD	NUM	WDTH	RPM	CAP	TYPE		
				/SEC	/OP	/SEC	/OP	/SEC	/OP	/SEC	/OP								
0000	CKD 1Gb	0000	0000	0.0	0.0	0.0	16.0	0.0	0.0	0.0	96.0		1	6	15	1800G	RAID 5		
		0004	0000	0.0	65.5K	72.8	0.0	0.0	1.3M	2.8K	100.0		1	7	15	2100G	RAID 5		
		0010	000A	190.0	57.2K	10.9M	2.2	8.0	1.1M	8.9M	9.3	Y	1	6	N/A	2400G	RAID 5		
		0012	000A	180.6	57.3K	10.3M	2.3	8.3	1.1M	9.0M	9.5	Y	1	6	N/A	2400G	RAID 5		
		POOL		370.6	57.2K	21.2M	2.2	16.3	1.1M	17.9M	9.4	Y	4	25	0	8700G	RAID 5		
0001	CKD 1Gb	0001	0000	0.0	0.0	0.0	0.0	0.0	1.4M	22.9K	22.9		1	6	15	1800G	RAID 5		
		0005	0000	0.0	0.0	0.0	0.0	0.0	1.6M	22.9K	39.4		1	7	15	2100G	RAID 5		
		000F	000A	82.9	57.3K	4.7M	2.4	4.3	1.0M	4.5M	7.8	Y	1	6	N/A	2400G	RAID 5		
		0011	000A	82.9	57.3K	4.7M	2.4	5.2	1.1M	5.6M	7.5	Y	1	6	N/A	2400G	RAID 5		
		POOL		165.8	57.3K	9.5M	2.4	9.5	1.1M	10.1M	7.7	Y	4	25	0	8700G	RAID 5		
0002	FIBRE 1Gb	0002	0000	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0		1	6	15	1800G	RAID 5		
0003	FIBRE 1Gb	0003	0000	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0		1	6	15	1800G	RAID 5		
0004	CKD 1Gb	0006	0002	0.8	25.2K	20.5K	8.0	23.2	1.3M	30.6M	16.9		1	6	10	7200G	RAID 5		
		0008	0002	0.7	25.2K	17.9K	12.7	19.8	1.3M	26.1M	34.6		1	6	10	7200G	RAID 5		
		000A	0002	0.5	18.3K	8.3K	11.5	17.3	1.5M	26.3M	23.9		1	7	10	8400G	RAID 5		
		POOL		2.0	23.7K	46.7K	10.5	60.4	1.4M	82.9M	24.7		3	19	10	22800G	RAID 5		
0005	CKD 1Gb	0007	0002	0.1	32.8K	3.8K	10.2	0.0	0.0	0.0	0.0		1	6	10	7200G	RAID 5		
		0009	0002	0.1	32.3K	3.6K	8.9	0.0	0.0	0.0	0.0		1	6	10	7200G	RAID 5		
		000B	0002	0.2	32.6K	5.0K	11.7	0.0	0.0	0.0	0.0		1	7	10	8400G	RAID 5		
		POOL		0.4	32.6K	12.5K	10.4	0.0	0.0	0.0	0.0		3	19	10	22800G	RAID 5		
0006	CKD 1Gb	000E	0003	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0		1	5	7	20000G	RAID 6		

Figure 219. ESS Rank Statistics

Table 178. Fields in the ESS Rank Statistics	
Field Heading	Meaning
EXTENT POOL	Pool of allocation units for logical volumes. ID extent pool identifier TYPE extent type, for example, FIBRE 1Gb or CKD 1Gb

Table 178. Fields in the ESS Rank Statistics (continued)	
Field Heading	Meaning
RRID	RAID rank identifiers in the extent pool. Note: The line where RRID = POOL contains the average for all rank values of the entire extent pool.
ADAPT ID	Adapter Pair ID.
READ OPERATIONS OPS/SEC	Number of read operations per second.
READ OPERATIONS BYTES/OP	Average number of bytes per read operation.
READ OPERATIONS BYTES/SEC	Average bandwidth of a read operation.
READ OPERATIONS RTIME/OP	Average response time of read operations in milliseconds.
WRITE OPERATIONS OPS/SEC	Number of write operations per second.
WRITE OPERATIONS BYTES/OP	Average number of bytes per write operation.
WRITE OPERATIONS BYTES/SEC	Average bandwidth of a write operation.
WRITE OPERATIONS RTIME/OP	Average response time of write operations in milliseconds.
ARRAY	SSD If a Y is displayed, then there is at least one solid state drive in the rank array. NUM Number of arrays on the rank. WDTH Sum of DDMs (disk drive modules) of a rank excluding spares of the rank. For example, if you have a RAID-5 array with 6 data disks and 1 parity disk, ARRAY WDTH is 7, or for a RAID-10 with 3 mirrored disks, ARRAY WDTH is 6.
MIN RPM	The slowest drive of the rank in units of 1000 RPM (rounds per minute).
RANK CAP	The sum of bytes of a rank.
RAID TYPE	RAID type found for the rank, for example, - RAID-5 - RAID-10 In the line displaying the average values for the entire extent pool (where RRID = POOL), 'MIXED' is shown if different RAID types have been encountered for the individual ranks in the extent pool.

Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*. The following tables show the overview condition names for the Overview report, divided according to the sections of the enterprise disk systems report.

Table 179. Overview names in the ESS Link Statistics section	
Field Heading or Meaning	Overview Name
BYTES/SEC for SCSI READ	ESTRSRD
BYTES/SEC for SCSI WRITE	ESTRSWR

Table 179. Overview names in the ESS Link Statistics section (continued)

Field Heading or Meaning	Overview Name
BYTES/OPERATION for SCSI READ	ESPSSRD
BYTES/OPERATION for SCSI WRITE	ESPSSWR
OPERATIONS/SEC for SCSI READ	ESARSRD
OPERATIONS/SEC for SCSI WRITE	ESARSWR
RESP TIME/OPERATION for SCSI READ	ESRTSRD
RESP TIME/OPERATION for SCSI WRITE	ESRTSWR
I/O INTENSITY for SCSI READ	ESIOISRD
I/O INTENSITY for SCSI WRITE	ESIOISWR
I/O INTENSITY for SCSI TOTAL	ESIOIST
BYTES/SEC for ECKD READ	ESTRERD
BYTES/SEC for ECKD WRITE	ESTREWR
BYTES/OPERATION for ECKD READ	ESPSERD
BYTES/OPERATION for ECKD WRITE	ESPSEWR
OPERATIONS/SEC for ECKD READ	ESARERD
OPERATIONS/SEC for ECKD WRITE	ESAREWR
RESP TIME/OPERATION for ECKD READ	ESRTERD
RESP TIME/OPERATION for ECKD WRITE	ESRTEWR
I/O INTENSITY for ECKD READ	ESIOIERD
I/O INTENSITY for ECKD WRITE	ESIOIEWR
I/O INTENSITY for ECKD TOTAL	ESIOIET
BYTES/SEC for PPRC SEND	ESTRPSD
BYTES/SEC for PPRC RECEIVE	ESTRPRV
BYTES/OPERATION for PPRC SEND	ESPSPSD
BYTES/OPERATION for PPRC RECEIVE	ESPSPRV
OPERATIONS/SEC for PPRC SEND	ESARPSD
OPERATIONS/SEC for PPRC RECEIVE	ESARPRV
RESP TIME/OPERATION for PPRC SEND	ESRTPSD
RESP TIME/OPERATION for PPRC RECEIVE	ESRTPRV
I/O INTENSITY for PPRC SEND	ESIOIPSD
I/O INTENSITY for PPRC RECEIVE	ESIOIPRV
I/O INTENSITY for PPRC TOTAL	ESIOIPT

Table 180. Overview names in the ESS Extent Pool Statistics section

Field Heading	Overview Name
REAL CAPACITY	ESXRCAP
REAL EXTENTS	ESXRNSG
COMPRESSION PHYS TOT	ESXCTPC
COMPRESSION PHYS USED	ESXCUPC
COMPRESSION LOG TOT	ESXCTLCT
COMPRESSION LOG USED	ESXCULC

Table 180. Overview names in the ESS Extent Pool Statistics section (continued)	
Field Heading	Overview Name
COMPRESSION RATIO	ESXCRAT

Table 181. Overview names in the ESS Rank Statistics section	
Field Heading	Overview Name
READ OPERATIONS OPS/SEC	ESRROP
READ OPERATIONS BYTES/OP	ESRRBOP
READ OPERATIONS BYTES/SEC	ESRRBD
READ OPERATIONS RTIME/OP	ESRRRT
WRITE OPERATIONS OPS/SEC	ESRWOP
WRITE OPERATIONS BYTES/OP	ESRWBOP
WRITE OPERATIONS BYTES/SEC	ESRWBD
WRITE OPERATIONS RTIME/OP	ESRWRT

Table 182. Overview names in the ESS Synchronous I/O Link Statistics section	
Field Heading	Overview Name
OPS/SEC for SYNCH I/O CACHE READ OPERATIONS	ESSOSCR
OPS/SEC for SYNCH I/O CACHE WRITE OPERATIONS	ESSOSCW
OPS/SEC for SYNCH I/O NVS WRITE OPERATIONS	ESSOSNW
BYTES/OP for SYNCH I/O CACHE READ OPERATIONS	ESSBOCR
BYTES/OP for SYNCH I/O CACHE WRITE OPERATIONS	ESSBOCW
BYTES/OP for SYNCH I/O NVS WRITE OPERATIONS	ESSBONW
RTIME/OP for SYNCH I/O CACHE READ OPERATIONS	ESSRTOCR
RTIME/OP for SYNCH I/O CACHE WRITE OPERATIONS	ESSRTOCW
RTIME/OP for SYNCH I/O NVS WRITE OPERATIONS	ESSRTONW
% SUCC for SYNCH I/O CACHE READ OPERATIONS	ESSPSOCR
% SUCC for SYNCH I/O CACHE WRITE OPERATIONS	ESSPSOCW
% SUCC for SYNCH I/O NVS WRITE OPERATIONS	ESSPSONW

FCD - FICON Director Activity report

With the Fibre Channel architecture and Fibre Channel switches (referred to in the following as FICON directors), link busy conditions are not returned. Instead, the FICON director queues the frames internally and sends them through when the port becomes available. This switch latency can grow as contention for ports increases. Therefore, it is important to report this switch latency (per port), this helps for the following tasks:

- Capacity planning
- Analysis of performance problems and bottlenecks
- Identification of contributors to device pending and disconnect times
- Understanding the contention for reconnection status

How to request this report

The default option for Monitor I data gathering is NOFCD. Therefore, it is required that you specify FCD if you want to get this report.

To produce this report, specify

REPORTS(FCD(option))

This report is also available in XML output format. How to work with Postprocessor XML reports in z/OS Resource Measurement Facility User's Guide provides all required information on how to produce and view XML reports.

Example URL for the DDS API

http://ddshost:8803/gpm/rmfpp.xml?reports=FCD

Contents of the report

The measurements provided for a port in the FCD report do not only comprise the I/O for the system on which the report is taken, but include all I/O that is directed through this port, regardless of which LPAR requests the I/O.

F I C O N D I R E C T O R A C T I V I T Y										PAGE	1
z/OS 3.1		SYSTEM ID SYS1		DATE 09/30/2021		INTERVAL 15.00.000					
		RPT VERSION 3.1 RMF		TIME 07.15.00		CYCLE 1.000 SECONDS					
IODF = 99 CR-DATE: 09/14/2019 CR-TIME: 08.45.00 ACT: POR											
SWITCH DEVICE: 0414 SWITCH ID: 01 TYPE: 005000 MODEL: 001 MAN: MCD PLANT: 01 SERIAL: 00000MK00109											
PORT ADDR	UNIT	CONNECTION ID	SERIAL NUMBER	AVG FRAME PACING	AVG FRAME READ	SIZE WRITE	PORT BANDWIDTH (MB/SEC) -- READ --	-- WRITE --	ERROR COUNT		
05	CHP	FA	000000099802	0	808	285	50.04	10.50	0		
07	CHP	4A	00000009F230	0	149	964	20.55	5.01	0		
08	CHP-H	F4	000000070882	0	568	965.5T	70.56	4.02	1		
09	CHP	FC	00000003C03F	0	558	1424	50.07	10.53	0		
08	CHP	F4	00000004C057	0	872	896	50.00	10.56	0		
12	CHP	D5	00000005C86D	0	73	574	20.51	5.07	0		
13	CHP	C8	00000008C1DF	0	868	1134	70.52	2.08	1		
14	SWITCH	----	0000013124DA	0	962	287	50.03	10.59	0		
15	CU	C800	0000000CF811	0	1188	731	20.54	5.00	0		
	CU	CA00	0000000CD111				70.55	3.01	1		
16	CHP	CB	00000009CE35	0	740	1185	70.50	2.06	1		

Figure 220. FICON Director Activity Report

Table 183. Fields in the FICON Director Activity Report.	
Field Heading	Meaning
IODF = xx	The IODF number where xx is the suffix of the IODF data set name.
CR-DATE: mm/dd/yyyy	The creation date of the IODF.
CR-TIME: hh.mm.ss	The creation time of the IODF.
ACT: text	The configuration state where text indicates how the IODF was activated.
SWITCH DEVICE	The hexadecimal number of the switch device of the FICON director for which measurements are being reported.
SWITCH ID	The hexadecimal switch identifier of the FICON director which is associated with this switch device. In case of cascaded switches, '***' may be shown.
TYPE, MODEL, MAN, PLANT, SERIAL	The hardware description of the switch device.
PORT ADDR	The hexadecimal address of the port.

Table 183. Fields in the FICON Director Activity Report. (continued)	
Field Heading	Meaning
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. SERIAL NUMBER The serial number of the connected 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 *z/OS Resource Measurement Facility User's Guide*. The following table shows the overview condition names for the Overview report.

Table 184. Overview names in the FICON Director Activity Report	
Field Heading or Meaning	Overview Name
AVG FRAME PACING	FDAFPT
PORT BANDWIDTH (READ - MB/SEC)	FDMBREAD
PORT BANDWIDTH (WRITE - MB/SEC)	FDMBWRT
ERROR COUNT	FDNERR

HFS - Hierarchical File System Statistics report

The Hierarchical File System Statistics report provides information about activities and storage usage within your z/OS UNIX environment. This data can be used to analyze whether storage and buffer pool definitions are correct, or whether some adjustments should be performed to improve the performance of I/O activities for HFS files.

How to request this report

Monitor III gathers global data for this report as SMF record type 74.6. If you want to get information about specific hierarchical file systems, you have to activate the Monitor III gatherer option HFSNAME(ADD(hfsname)).

To produce this report, specify

REPORTS (HFS)

This report is also available in XML output format. Topic [How to work with Postprocessor XML reports in z/OS Resource Measurement Facility User's Guide](#) provides all required information on how to produce and view XML reports.

Example URL for the DDS API

<http://ddshost:8803/gpm/rmfpp.xml?reports=HFS>

Contents of the report

The report consists of two parts.

HFS Global Statistics Report

The first part of the HFS report provides overall data about I/O activities of HSF files and gives statistics about the various buffer pools which have been defined.

The report can be used as an entry point for performance investigation and capacity planning.

HFS File System Statistics Report

The second part of the report is based on data gathering for specific file systems. You get data about I/O activities and about the internal structure (index) of the HFS files.

Both parts of the report can help you

- in getting a general understanding of the throughput recognized and achieved by HFS to optimally use your resources.
- in identifying potential problems and bottlenecks within HFS and taking corrective actions.

HFS Global Statistics

HFS GLOBAL STATISTICS									
z/OS 3.1		SYSTEM ID SYS1		DATE 09/30/2021	INTERVAL 15.00.000		PAGE 1		
		RPT VERSION 3.1 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 221. HFS Global Statistics Report

Table 185. 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).

Table 185. Fields in HFS Global Statistics Report (continued)	
Field Heading	Meaning
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.
DATA SPACES	Number of data spaces comprising this buffer pool.
I/O ACTIVITY - TOTAL	Total number of buffers in this buffer pool for which I/Os were issued. This is not necessarily the number of actual I/Os issued since multiple buffers can be written in a single I/O request.
I/O ACTIVITY - FIXED	Number of times a buffer was already fixed prior to an I/O request in this buffer pool.
I/O ACTIVITY - %FIXED	Percentage of fixed I/Os.

HFS File System Statistics

HFS FILE SYSTEM STATISTICS											
z/OS 3.1		SYSTEM ID SYS1		DATE 09/30/2021		INTERVAL 15.00.000		PAGE 2			
		RPT VERSION 3.1 RMF		TIME 12.00.00		CYCLE 1.000 SECONDS					
--- ALLOCATION (MB) ---	SIZE	FILE COUNT	I/O RATE	METADATA COUNT	I/O RATE	INDEX COUNT	I/O RATE	INDEX EVENTS	EVENTS	--- COUNT	
FILE SYSTEM NAME: OMVS.SYS4.ROOT											
MOUNT DATE: 09/14/2021 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	
CACHED	0.000	SEQUENTIAL	20	0.022							
		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: 09/14/2021 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	
CACHED	0.000	SEQUENTIAL	0	0.000							
		RANDOM	0	0.000							

Figure 222. HFS File System Statistics Report

Table 186. 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.

Table 186. Fields in the HFS File System Statistics Report (continued)	
Field Heading	Meaning
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).	
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

The HTTP Server is the Web server for the family of WebSphere® application servers which provide the run-time environment for e-business applications.

The HTTP Server report provides information about the activities of a server. The information can be used to analyze the activities of the server in case of problems.

How to request this report

The Postprocessor requires type 103 subtypes 1 and 2 SMF records as input for the HTTP Server report. These records are not gathered by an RMF monitor, but were written by the IBM HTTP Server (IHS) powered by Domino, which is no longer supported in z/OS V2R2. SMF type 103 records created on a prior release of z/OS can still be used to generate a RMF Postprocessor HTTP report.

The IBM HTTP server powered by Apache does not write type 103 subtypes 1 or 2 SMF records, which means that no RMF Postprocessor HTTP report can be generated for that HTTP server.

To produce this report, specify

REPORTS (HTTP)

Contents of the report

The report consists of two parts:

- HTTP Server Summary

The summary contains one line for each server which is part of the report.

- HTTP Server Details

This part consists of two sections:

- Configuration data (provided by record type 103-1)
- Performance data (provided by record type 103-2)

Configuration data is reported together with performance data. Configuration data is not reported, if there is not at least one corresponding performance data record.

If there is no configuration data available, the line

*** NO CONFIGURATION DATA AVAILABLE WITHIN GIVEN RECORD INTERVAL ***

is shown.

HTTP SERVER SUMMARY												PAGE 1
z/OS 3.1	SYSTEM ID RPT VERSION 3.1	SYS1 RMF	DATE 09/30/2021 TIME 08.06.37	INTERVAL 30.00.000								
SERVER NAME SERVER TOKEN	AVAILABLE HHH.MM.SS	REQUEST RATE	RESPONSE RATE	THROUGHPUT IN	RATE OUT	THREADS MAX USED	CACHE SIZE MAX USED	CACHE MAX	FILES USED	TIMEOUTS		
MVS071 4F37-CA73-0005-678E mvs047.tcp.raleigh.ibm.com 3D29-CA40-0006-997C	000.24.35 000.01.10	0.03 2.14	0.03 0.07	5.46 8.67	30.15 41.37	39 0.00 23 0.00	5120 0.90 5120 1.06	0 0.00 0 0.00	0 0.00 0 0.00	1 0		
s390server17.wsclab.washington.ibm.com 3D29-CA4C-0008-7E86	000.07.37	0.00	0.00	0.00	0.00	150 0.00	5120 0.00	NO 0.00	0 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 223. HTTP Server Report - Summary

H T T P S E R V E R D E T A I L S																					
z/OS 3.1		SYSTEM ID SYS1		DATE 09/30/2021		INTERVAL 30.00.000		PAGE 2													
		RPT VERSION 3.1 RMF		TIME 08.06.37																	

SERVER CHARACTERISTICS																					

SERVER NAME:		s390server17.wslab.washington.ibm.com																			
SERVER ROOT IN HFS:		/webtime/development/pr																			
IP-ADDRESS:		9.67.116.9		STARTUP:		09/14/2021-08.03.01															
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									
IN		893		4.85		DNS LOOKUP		0.000		0.000		0.000		200 (OK)		1					
OUT		730		3.97		SERVICE PLUGIN		0.000		0.170		0.060		302 (MOVED TEMPORARILY)		0					
UNKNOWN		0		0.00		CGI		1.580		1.580		1.580		401 (UNAUTHORIZED)		0					
						SSL HANDSHAKE		0.000		0.000		0.000		403 (FORBIDDEN)		0					
						PROXY		9.580		9.580		9.580		404 (NOT FOUND)		2					
														407 (PROXY UNAUTHORIZED)		0					
														500 (INTERNAL SERVER ERROR)		0					

Figure 224. HTTP Server Report - Details

Table 187. 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 values are displayed in the line below, as shown in Figure 223 on page 413 .
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.
CACHE SIZE	MAX Maximum cache size (KB) of 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 188. 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.
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.

<i>Table 188. Fields in the HTTP Server Details Report (continued)</i>	
Field Heading	Meaning
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.
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).	

Table 188. Fields in the HTTP Server Details Report (continued)	
Field Heading	Meaning
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

The I/O Queuing Activity report provides information on the I/O configuration and activity rate, queue lengths, and percentages when one or more I/O components, grouped by a logical control unit (LCU), were busy.

For all channels that are managed by **Dynamic Channel Path Management (DCM)**, additional information is available. DCM allows an installation to identify channels which they wish to be managed dynamically. These channels are not assigned permanently to a specific control unit, but belong to a pool of channels. Based on workload requirements in the system, these channels will be assigned dynamically by DCM. For each LCU with DCM managed channels, a summary line displays the minimum and maximum number of connected DCM managed channels, the number of defined DCM managed channels and accumulated activity data.

An LCU is the set of devices attached to the same physical control unit (or group of control units that have one or more devices in common). Each device belongs to only one LCU, but the I/O processor (SAP - System Assist Processor), which is part of the channel subsystem, manages and schedules I/O work requests to the various devices within the LCU. If an I/O request is unsuccessful because the control unit is busy, the request is queued on the control unit header (CU-HDR) queue. Once the busy condition is resolved, the CU-HDR is then placed in the initiative queue.

PAV base mode is the mode when alias devices are assigned to one PAV base device. An I/O for a PAV base device is executed using aliases assigned to that PAV base device.

HyperPAV mode is the mode when a pool of alias devices is assigned to one LCU. An I/O for a PAV base device can be executed using any alias device of that pool.

SuperPAV mode is the mode when a pool of alias devices is assigned to one LCU and multiple LCUs are grouped into one Alias Management Group (AMG). An I/O for a PAV base device can be executed using any alias device of these multiple alias pools. The favored way is to use the alias device assigned to the same LCU (home LCU) that the PAV base device is assigned to.

Your installation defines your I/O configuration as input to the input/output configuration program (IOCP). The IOCP uses the information you supply to define the relationship between channel paths, control units, and I/O devices. The IOCP generates and assigns LCU identifiers to these groups of channel paths, control units, and I/O devices. The IOCP then places this configuration definition in a configuration data set (IOCDS). RMF uses the configuration definition as well as measurement data gathered during the interval to generate the I/O Queuing Activity report.

How to request this report

To gather data for this report, specify as a Monitor I gatherer option:

```
IOQ(option)
```

To produce this report, specify

```
REPORTS(IOQ)
```

This report is also available in XML output format. How to work with Postprocessor XML reports in *z/OS Resource Measurement Facility User's Guide* provides all required information on how to produce and view XML reports.

Example URL for the DDS API

```
http://ddshost:8803/gpm/rmfpp.xml?reports=IOQ
```

Using the information given in the report

If the Channel Path Activity and I/O Device Activity reports have shown that a problem exists, you can use the information in the I/O Queuing Activity report to pinpoint the reason for contention delays associated with channel paths, control units, and devices. For example, if the I/O Device Activity report shows an unusually large pending time for one or more devices in an LCU, the I/O Queuing Activity report indicates what proportion of the delay is caused by control unit busy and device busy. This proportion indicates which part of the configuration might need adjustment.

You can also use the I/O Device Activity report and I/O Queuing Activity report to analyze the current I/O configuration. The I/O Device Activity report shows which devices belong to each logical control unit. The I/O Queuing Activity report shows which physical control units are part of each logical control unit and which channel paths are connected to each physical control unit.

Data gathering considerations

The report depends on information in the I/O configuration data set (IOCDS). If RMF cannot read the IOCDS, or if the IOCDS has been updated so that the data might not apply to the present configuration, no report is available. For example, when the operator partitions the system in such a way that RMF cannot read the IOCDS because it appears in another partition of a multi-processing system, RMF terminates the I/O Queuing Activity report and issues a message to the operator, I/O QUEUING ACTIVITY RMF REPORT TERMINATED.

Missing data in report fields

When a **LCU has no activity** during the interval, RMF omits that LCU from the report for that interval. If no activity has occurred during the interval for all selected LCUs, the message NO ACTIVITY FOR SELECTED LCUs appears instead of the data after the headings of the report.

If a **channel path was brought online or taken offline** during the interval, data is formatted and an additional line in the report describes its status. If an installed channel path was offline during the whole interval, the additional line identifies the channel path as OFFLINE. If a channel path was taken offline or brought online during an interval, the additional line identifies the channel path as either NOW OFFLINE or NOW ONLINE.

When RMF **cannot obtain valid hardware data** for CONTENTION RATE and DELAY Q LENGTH, it prints the message NO H/W DATA under those headings.

If the **channel measurement facility is inactive** or has been interrupted during the interval, CHANNEL MEASUREMENT FACILITY NOT ACTIVE OR INTERRUPTED appears after the headings where the data normally appears in the report.

If the **diagnosis interface fails** during the interval, DIAGNOSIS INTERFACE FAILURE appears after the headings in the report.

Messages

During the measurement of I/O Queuing activity, you may see one of the following messages in the data line:

LCU DYNAMICALLY CHANGED

A LCU was dynamically changed during the interval.

LCU DYNAMICALLY ADDED

A LCU was dynamically created during the interval.

LCU CHANGE ATTEMPTED

A configuration change was attempted, but did not complete successfully.

Contents of the report

I/O QUEUING ACTIVITY																	PAGE 1	
z/OS 3.1			SYSTEM ID RSC2			DATE 09/30/2021			INTERVAL 15.00.032									
TOTAL SAMPLES = 900			IODF = 07			RPT VERSION 3.1 RMF			TIME 23.44.33			CYCLE 1.000 SECONDS						
			CR-DATE: 09/14/2019			CR-TIME: 16.05.52			ACT: ACTIVATE									

INPUT/OUTPUT PROCESSORS																		
-INITIATIVE QUEUE-			IOP UTILIZATION					-- % I/O REQUESTS RETRIED --					RETRIES / SSCH					
IOP	ACTIVITY RATE	AVG Q LGNTH	% IOP BUSY	% CMPR BUSY	% SCM BUSY	I/O START RATE	INTERRUPT RATE	ALL	CP BUSY	DP BUSY	CU BUSY	DV BUSY	ALL	CP BUSY	DP BUSY	CU BUSY	DV BUSY	
00	1395.107	0.00	1.20	4.30	4.30	1384.433	3573.426	0.0	0.0	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00	
01	1401.048	0.00	0.78	4.30	4.30	1390.375	1307.718	0.2	0.2	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00	
SYS	2796.15	0.00	1.00	4.30	4.30	20824.67	21871.27	1.4	1.4	0.0	0.0	0.0	0.01	0.01	0.00	0.00	0.00	

ALIAS MANAGEMENT GROUPS																		

AMG	DCM GROUP		CHAN	CHPID	% DP	% CU	AVG	AVG		DELAY	AVG			AVG	DATA			
	MIN	MAX	DEF	TAKEN	BUSY	BUSY	CUB DLY	CMR DLY	CONTENTION RATE	Q LGNTH	CSS DLY	HPAV WAIT	MAX	OPEN EXCH	XFER CONC			
00000016			65	0.537	0.00	0.00	0.0	3.6										
			34	0.538	0.00	0.00	0.0	4.5										
			99	0.532	0.00	0.00	0.0	1.3										
			*	1.607	0.00	0.00	0.0	3.1	0.000	0.00	0.8	0.000	0					
00000017			65	0.361	0.00	0.00	0.0	2.1										
			34	0.356	0.00	0.00	0.0	3.8										
			99	0.357	0.00	0.00	0.0	1.2										
			*	1.073	0.00	0.00	0.0	2.4	0.000	0.00	0.7	0.000	0					

LOGICAL CONTROL UNITS																		

LCU/ AMG	CU	DCM GROUP		CHAN	CHPID	% DP	% CU	AVG	AVG		DELAY	AVG			AVG	DATA		
		MIN	MAX	DEF	TAKEN	BUSY	BUSY	CUB DLY	CMR DLY	CONTENTION RATE	Q LGNTH	CSS DLY	HPAV WAIT	MAX	OPEN EXCH	XFER CONC		
0049	5100			43	0.026	0.00	0.00	0.0	0.0									
				55	0.024	0.00	0.00	0.0	0.0									
				44	0.024	0.00	0.00	0.0	0.0									
				56	0.026	0.00	0.00	0.0	0.0									
				*	0.100	0.00	0.00	0.0	0.0	0.000	0.00	0.1	0.000	0	0.00	0.00		
0107 00000016	B101			65	0.084	0.00	0.00	0.0	3.6									
				34	0.083	0.00	0.00	0.0	4.7									
				99	0.082	0.00	0.00	0.0	1.3									
				*	0.250	0.00	0.00	0.0	3.2	0.000	0.00	0.8	0.000	0				
011A 00000016	CF01			65	0.072	0.00	0.00	0.0	3.8									
				34	0.073	0.00	0.00	0.0	4.9									
				99	0.071	0.00	0.00	0.0	1.2									
				*	0.217	0.00	0.00	0.0	3.4	0.000	0.00	0.7	0.000	0				
0106 00000017	B001			65	0.084	0.00	0.00	0.0	2.1									
				34	0.082	0.00	0.00	0.0	4.1									
				99	0.083	0.00	0.00	0.0	1.2									
				*	0.250	0.00	0.00	0.0	2.5	0.000	0.00	0.8	0.000	0				

Figure 225. I/O Queuing Activity Report

The I/O Queuing Activity Report contains three sections with these titles:

INPUT/OUTPUT PROCESSORS

This section shows the measurements accumulated for I/O processors.

ALIAS MANAGEMENT GROUPS

For each defined Alias Management Group (AMG), this section shows performance measurements for all channel paths connected to the LCUs grouped into the AMG.

LOGICAL CONTROL UNITS

For each LCU having online devices, this section shows performance measurements for all channel paths connected to the LCU.

Table 189. 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.
INPUT/OUTPUT PROCESSORS	

Table 189. Fields in the I/O Queuing Activity Report (continued)

Field Heading	Meaning
IOP	The two-digit hexadecimal identifier of the I/O processor (IOP). The IOP data sections are sorted according ascending IOP numbers. Following the last IOP data line is a line that summarizes the measurement data of the individual IOPs. This summary line starts with the character string 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 LENGTH 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. $\text{AVG Q LENGTH} = \frac{\text{Accumulated Queue Length}}{\text{\# I/O Requests on the IOP Queue}} - 1$

Table 189. Fields in the I/O Queuing Activity Report (continued)

Field Heading	Meaning
IOP UTILIZATION	% IOP BUSY The ratio of the number of times the IOP was found busy to the total number of I/O processor samples. $\% \text{ IOP BUSY} = \frac{\# \text{ Busy samples}}{\# \text{ Busy samples} + \# \text{ Idle samples}} * 100$ A high IOP utilization might be caused by a high level of activity in terms of SSCH, I/O or sysplex operations per second or by contention in the I/O configuration. If contention is caused by CP BUSY or CU BUSY conditions, the request is placed on the IOP queue. This is indicated by an AVG Q LENGTH value greater than zero. If contention is caused by DP BUSY conditions, this is not indicated by the AVG Q LENGTH value, because the requests are kept internally. When the IOP is idle, these requests are processed which is reflected by the %IOP BUSY field. % CMPR BUSY The ratio of the number of times the IOP was found busy with EADM compression work (compress or decompress) to the total number of I/O processor samples. $\% \text{ CMPR BUSY} = \frac{\# \text{ EADM compression busy samples}}{\# \text{ Busy samples} + \# \text{ Idle samples}} * 100$ % SCM BUSY The ratio of the number of times the IOP was found busy with SCM work to the total number of I/O processor samples. $\% \text{ SCM BUSY} = \frac{\# \text{ SCM busy samples}}{\# \text{ Busy samples} + \# \text{ Idle samples}} * 100$ 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. $\text{I/O START RATE} = \frac{\# \text{ I/O functions started}}{\text{Interval}}$ 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. $\text{INTERRUPT RATE} = \frac{\# \text{ Processed I/O interrupts}}{\text{Interval}}$

Table 189. Fields in the I/O Queuing Activity Report (continued)

Field Heading	Meaning
% I/O REQUESTS RETRIED	ALL The ratio of the number of retries to the number of I/O functions initially started plus the total number of retries. $\%ALL = \frac{\# \text{ Retries}}{\# \text{ I/O functions started} + \# \text{ Retries}} * 100$ CP BUSY The ratio of the number of I/O operations retried on the I/O processor because the selected channel path was busy, to the number of I/O functions initially started plus the total number of retries. $\%CP \text{ BUSY} = \frac{\# \text{ Retries due to channel path busy}}{\# \text{ I/O functions started} + \# \text{ Retries}} * 100$ 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 to the number of I/O functions initially started plus the total number of retries. $\%DP \text{ BUSY} = \frac{\# \text{ Retries due to director port busy}}{\# \text{ I/O functions started} + \# \text{ Retries}} * 100$ 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 to the number of I/O functions initially started plus the total number of retries. $\%CU \text{ BUSY} = \frac{\# \text{ Retries due to control unit busy}}{\# \text{ I/O functions started} + \# \text{ Retries}} * 100$ 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 to the number of I/O functions initially started plus the total number of retries. $\%DV \text{ BUSY} = \frac{\# \text{ Retries due to device busy}}{\# \text{ I/O functions started} + \# \text{ Retries}} * 100$

Table 189. Fields in the I/O Queuing Activity Report (continued)

Field Heading	Meaning
RETRIES / SSCH	ALL The ratio of the number of retries on the I/O processor to the number of I/O functions initially started. $\text{ALL} = \frac{\text{\# Retries}}{\text{\# I/O functions started}}$ CP BUSY The ratio of the number of retries on the I/O processor because the selected channel path was busy to the number of I/O functions initially started. $\text{CP BUSY} = \frac{\text{\# Retries due to channel path busy}}{\text{\# I/O functions started}}$ 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 to the number of I/O functions initially started. $\text{DP BUSY} = \frac{\text{\# Retries due to director port busy}}{\text{\# I/O functions started}}$ CU BUSY The ratio of the number of retries on the I/O processor because the control unit of the targeted device was busy to the number of I/O functions initially started. $\text{CU BUSY} = \frac{\text{\# Retries due to control unit busy}}{\text{\# I/O functions started}}$ DV BUSY The ratio of the number of retries on the I/O processor because the targeted device was busy to the number of I/O functions initially started. $\text{DV BUSY} = \frac{\text{\# Retries due to device busy}}{\text{\# I/O functions started}}$
ALIAS MANAGEMENT GROUPS	
AMG	The eight-digit hexadecimal system Alias Management Group assigned by I/O Supervisor.
LOGICAL CONTROL UNITS	
LCU/ AMG	LCU/ The four-digit hexadecimal identifier of the Logical Control Unit (LCU). AMG The eight-digit hexadecimal system Alias Management Group assigned by I/O Supervisor, if the LCU is grouped to an AMG. 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.
CU	The four-digit hexadecimal identifier of each physical control unit contained in the logical control unit.
Note: The following fields apply for both Alias Management Groups and Logical Control Units.	

Table 189. Fields in the I/O Queuing Activity Report (continued)

Field Heading	Meaning
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/AMG (in this interval) as well as the installation-specified definition for this LCU/AMG. The line with these values is available only for LCUs/AMGs with DCM managed channels. It contains in addition the accumulated values of the I/O activity rate, the director port contention, and the control unit contention of all DCM managed channels. These values may include also measurements of managed channels which were partially online.
CHAN PATHS	The two-digit hexadecimal channel path identifiers (CHPIDs) of the channel paths that are attached to the physical control units contained in the LCU/AMG. There can be up to eight channel paths in a logical control unit. The channel paths that are offline or moved online or offline during the interval are indicated as follows: 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/AMG. If the control unit supports channel path attributes, RMF displays them together with the channel path: PF preferred path NP non-preferred path NS path not specified In the following cases, RMF cannot find channel path attributes and therefore only displays the CHPID: • for devices residing in control units that do not support path attributes • for offline channels • for summary lines
CHPID TAKEN	The rate at which I/O requests to devices of this LCU/AMG are satisfied by each CHPID during the interval. By reviewing the rate at which each channel path of the LCU/AMG satisfies I/O requests, you can see how evenly the work requests are distributed among the available paths and how effectively those paths are arranged for the LCU/AMG. # I/O Operations Accepted on that Path CHPID TAKEN = ----- Interval Note: If vary activity has occurred during the interval, this field is blank.

Table 189. Fields in the I/O Queuing Activity Report (continued)

Field Heading	Meaning
% DP BUSY	The ratio of the number of times an I/O request was deferred because the director port was busy to the number of attempts to service I/O requests during the measurement interval. This field indicates director port contention. $\% \text{ DP BUSY} = \frac{\text{DPB}}{\text{DPB} + \text{CUB} + \text{SUC}} * 100$ 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
% CU BUSY	The ratio of the number of requests deferred due to control unit busy to the number of attempts to service I/O requests during the measurement interval. This field indicates control unit contention and is reported for each path within the LCU/AMG. $\% \text{ CU BUSY} = \frac{\text{CUB}}{\text{DPB} + \text{CUB} + \text{SUC}} * 100$ 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 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. $\text{AVG CUB DLY} = \frac{\text{Control Unit Busy Time}}{\# \text{ I/O Operations Accepted on that Path}}$
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). $\text{AVG CMR DLY} = \frac{\text{Initial Command Response Time}}{\# \text{ I/O Operations Accepted on that Path}}$
CONTENTION RATE	The rate at which the I/O processor places delayed I/O requests on the CU-HDR for this LCU/AMG. The IOP places an I/O request on the CU-HDR when all paths to the subchannel are busy and at least one path to the control unit is busy. For devices with only one path or for devices where multiple paths exist and the busy condition is resolved immediately over an alternate path, the IOP does not count the condition. $\text{CONTENTION RATE} = \frac{\# \text{ Enqueued Requests}}{\text{Interval}}$
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. $\text{DELAY Q LNGTH} = \frac{\text{Accumulated Queue Length}}{\# \text{ Enqueued Requests}} - 1$

Table 189. Fields in the I/O Queuing Activity Report (continued)	
Field Heading	Meaning
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/AMG, until the channel subsystem first attempts to initiate the operation. $\text{AVG CSS DLY} = \frac{\text{Channel Subsystem Time}}{\# \text{ I/O Operations Accepted}}$
HPAV WAIT	The ratio of the number of I/O requests that could not start because no HyperPAV aliases were available, to the total number of I/O requests for an LCU/AMG: $\text{HPAV Wait} = \frac{\text{I/Os that could not start}}{\text{Total I/Os}}$
HPAV MAX	The maximum number of concurrently used HyperPAV alias devices (including borrowed aliases) for that LCU/AMG during the interval.
AVG OPEN EXCH	The estimated average number of concurrently active I/O operations is provided in the LCU/AMG summary line if at least one FICON channel is connected to the LCU/AMG. $\text{AVG OPEN EXCH} = \frac{\text{CMR+CONN+DISC}}{\text{RMF interval}}$ CMR initial command response time CONN connect time DISC disconnect time
DATA XFER CONC	The data transfer concurrency is provided in the LCU/AMG summary line if at least one FICON channel is connected to the LCU/AMG. $\text{DATA XFER CONC} = \frac{\text{CONN}}{\text{RMF interval}}$

Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*. The following table shows the overview condition names for the Overview report.

Table 190. Overview names in the I/O Queuing Activity Report	
Field Heading or Meaning	Overview Name
ACTIVITY RATE	IOPAC
AVG Q LENGTH	IOPQL
% IOP BUSY	IOPIPB
% CMPR BUSY	IOPECB
% SCM BUSY	IOPSCB
Percent I/O processor idle	IOPIPI
I/O START RATE	IORIFS
INTERRUPT RATE	IORPII
% I/O REQU RETRIED (ALL)	IOPALB

Table 190. Overview names in the I/O Queuing Activity Report (continued)	
Field Heading or Meaning	Overview Name
% I/O REQU RETRIED (CP BUSY)	IOPCHB
% I/O REQU RETRIED (DP BUSY)	IOPDPB
% I/O REQU RETRIED (CU BUSY)	IOPCUB
% I/O REQU RETRIED (DV BUSY)	IOPDVB
RETRIES / SSCH (ALL)	IONALB
RETRIES / SSCH (CP BUSY)	IONCHB
RETRIES / SSCH (DP BUSY)	IONDPB
RETRIES / SSCH (CU BUSY)	IONCUB
RETRIES / SSCH (DV BUSY)	IONDVB
CHPID TAKEN	IOART
% DP BUSY	IODPB
% CU BUSY	IOCUB
CONTENTION RATE	IOCTR
DELAY Q LNGTH	IODLQ
AVG CUB DLY	IOCBT
AVG CMR DLY	IOCMR
AVG CSS DLY	IOCSS
HPAV WAIT	IOHWAIT
HPAV MAX	IOHMAX

OMVS - OMVS Kernel Activity report

The OMVS Kernel Activity report provides information about:

- OMVS System Call Activity
- OMVS Process Activity
- OMVS Inter-Process Communication
- OMVS Memory Map - Shared Library Regions - Queued Signals

How to request this report

Monitor III gathers data for this report automatically. If you want to suppress gathering, you have to disable writing SMF record type 74.3.

To produce this report, specify

```
REPORTS(OMVS)
```

This report is also available in XML output format. [Topic How to work with Postprocessor XML reports in z/OS Resource Measurement Facility User's Guide](#) provides all required information on how to produce and view XML reports.

Example URL for the DDS API

```
http://ddshost:8803/gpm/rmfpp.xml?reports=OMVS
```

Contents of the report

The OMVS Kernel Activity report has these parts:

- OMVS System Call Activity
- OMVS Process Activity
- OMVS Inter-Process Communication
- OMVS Memory Map - Shared Library Regions - Queued Signals

OMVS KERNEL ACTIVITY												PAGE 1	
z/OS 3.1			SYSTEM ID AQT5		DATE 09/30/2021		INTERVAL 30.00.000						
			RPT VERSION 3.1 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							
1200			50			12							
MAXIMUM (TOT)													

MINIMUM			AVERAGE		MAXIMUM		MINIMUM			AVERAGE		MAXIMUM	

CURRENT (TOT)		99	854	1200	12	13	23	0		1.0*	5.3		
OVERRUNS (N/S)		0	5.5*	333	0	1.5*	4.8						
OMVS INTER-PROCESS COMMUNICATION													

MESSAGE QUEUE IDS				SEMAPHORE IDS				SHARED MEMORY IDS				SHARED MEMORY PAGES	
500				500				500				262144	
MAXIMUM (TOT)													

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													

MEMORY MAP STORAGE PAGES				SHARED STORAGE PAGES				MAX SHARED LIBRARY REGION				MAXIMUM QUEUED SIGNALS	
4096				131K				16M				100K	
MAXIMUM (TOT)													

MINIMUM			AVERAGE		MAXIMUM		MINIMUM			AVERAGE		MAXIMUM	

CURRENT (TOT)		0.000	0.000	0.000	1542	1541	1542	16M	16M	16M	0.000		50K
OVERRUNS (N/S)		0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000		99K
Units: (TOT) = Total Value, (N/S) = Number per Second, (H/S) = Hundredth of seconds per Second													

Units: (TOT) = Total Value, (N/S) = Number per Second, (H/S) = Hundredth of seconds per Second

Figure 226. 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 191. Fields in the OMVS Kernel Activity Report	
Heading	Meaning
OMVS SYSTEM CALL ACTIVITY	
SYSCALLS (N/S)	Number of system calls per second processed by the OMVS kernel address space in this interval.
CPU TIME (H/S)	Time spent to process system calls in hundredths of seconds per second.
OMVS PROCESS ACTIVITY	
MAXIMUM PROCESSES USERS PROCESSES PER USER	Maximum number of processes, users, and processes per user defined by OMVS kernel address space initialization parameters (in Parmlib member BPXPRMxx). If one of these values has changed (due to an OMVS restart), it will be reported as '*****'.
CURRENT PROCESSES	Number of OMVS processes controlled by OMVS during this interval.
CURRENT USERS	Number of OMVS users controlled by OMVS.
OVERRUNS PROCESSES	Rate of processes that could not be created by OMVS because the maximum number of processes would have been exceeded.

Table 191. Fields in the OMVS Kernel Activity Report (continued)

Heading	Meaning
OVERRUNS USERS	Rate of OMVS users that could not be created by OMVS because the maximum number of users would have been exceeded.
OVERRUNS PROCESSES PER USER	Rate of processes per user that could not be created by OMVS because the maximum number of processes per user would have been exceeded.
OMVS INTER-PROCESS COMMUNICATION	
MAXIMUM MESSAGE QUEUE IDS SEMAPHORE IDS SHARED MEMORY IDS SHARED MEMORY PAGES	Maximum number of message queue IDs, semaphore IDs, shared memory IDs, and shared memory pages defined by OMVS kernel address space initialization parameters (in Parmlib member BPXPRMxx).
CURRENT MESSAGE QUEUE IDS	Number of message queue IDs during this interval.
CURRENT SEMAPHORE IDS	Number of semaphore IDs during this interval.
CURRENT SHARED MEMORY IDS	Number of shared memory IDs during this interval.
CURRENT SHARED MEMORY PAGES	Number of shared memory pages during this interval.
OVERRUNS MESSAGE QUEUE IDS	Rate of message queue IDs that could not be created by OMVS because the maximum number of message queue IDs would have been exceeded.
OVERRUNS SEMAPHORE IDS	Rate of semaphore IDs that could not be created by OMVS because the maximum number of semaphore IDs would have been exceeded.
OVERRUNS SHARED MEMORY IDS	Rate of shared memory IDs that could not be created by OMVS because the maximum number of shared memory IDs would have been exceeded.
OVERRUNS SHARED MEMORY PAGES	Rate of shared memory pages that could not be created by OMVS because the maximum number of shared memory pages would have been exceeded.
OMVS MEMORY MAP	
MAXIMUM MEMORY MAP STORAGE PAGES SHARED STORAGE PAGES	Maximum number of memory map storage pages and shared storage pages defined by OMVS kernel address space initialization parameters (in Parmlib member BPXPRMxx).
CURRENT MEMORY MAP STORAGE PAGES	Number of memory map storage pages during this interval.
CURRENT SHARED STORAGE PAGES	Number of shared storage pages during this interval.
OVERRUNS MEMORY MAP STORAGE PAGES	Rate of memory map storage pages that could not be created by OMVS because the maximum number of memory map storage pages would have been exceeded.
OVERRUNS SHARED STORAGE PAGES	Rate of shared storage pages that could not be created by OMVS because the maximum number of shared storage pages would have been exceeded.
SHARED LIBRARY REGION	
MAX SHARED LIBRARY REGION	Maximum amount of storage available for shared library region as specified by Parmlib statement SHRLIBRGNSIZE. The values are provided in units of megabytes.
CURRENT SHARED LIBRARY REGION	The current amount of storage in Megabytes available for shared library region.
OVERRUNS SHARED LIBRARY REGION	Rate of attempts to exceed the maximum storage amount for shared library region.
QUEUED SIGNALS	
MAX QUEUED SIGNALS	Maximum amount of queued signals allowed per process as specified by Parmlib statement MAXQUEUEDSIGS.
OVERRUNS QUEUED SIGNALS	Rate of attempts to exceed the maximum number of queued signals.

PAGESP - Page Data Set Activity report

The Page Data Set Activity report provides information about page data set usage for each individual data set. The information about the number of slots used is reported as minimum, maximum, and average values for the interval. Also, the time is provided when the Auxiliary Storage Manager (ASM) considered the data set to be busy, the number of start I/O requests initiated by ASM for the data set, the average page transfer time for each I/O request, and the number of pages transferred to and from the page data set.

The report contains only page data sets that:

- are in use at the end of that RMF measurement interval
- have been deleted during that RMF measurement interval

How to request this report

Monitor I gathers data for this report automatically. If you want to suppress gathering, you need to specify NOPAGESP.

To produce this report, specify

```
REPORTS(PAGESP)
```

This report is also available in XML output format. [How to work with Postprocessor XML reports in z/OS Resource Measurement Facility User's Guide](#) provides all required information on how to produce and view XML reports.

Example URL for the DDS API

```
http://ddshost:8803/gpm/rmfpp.xml?reports=PAGESP
```

Using the information given in the report

You can use the information in the page data set report, for example, to determine whether the optimum size has been allocated for each data set. If the maximum number of slots used is consistently below the number of slots allocated, you might consider reducing the size of the data set to conserve space on the device. However, use caution when reducing the size of the PLPA and common data sets because overflow cannot occur from these data sets to the local data sets.

The **% IN USE** field shows how busy the data set is. If this is above 30% you might see increases in response time. You might then:

- dedicate volumes to page data sets.
- make the sum of all the page space two to four times the number of slots used.
- limit use of **VIO=YES**

Contents of the report

PAGE DATA SET ACTIVITY													PAGE 1	
z/OS 3.1				SYSTEM ID		SYS1		DATE 09/30/2021		INTERVAL 14.59.946				
				RPT VERSION		3.1 RMF		TIME 16.30.00		CYCLE 1.000 SECONDS				
NUMBER OF SAMPLES = 900				PAGE DATA SET AND SCM USAGE										
PAGE SPACE TYPE	VOLUME SERIAL	DEV NUM	DEVICE TYPE	SLOTS ALLOC	---- MIN	SLOTS USED MAX	--- AVG	BAD SLOTS	% IN USE	PAGE TRANS TIME	NUMBER IO REQ	PAGES XFER'D	V I O	DATA SET NAME
PLPA	PGT1B5	001B5	33903	180	168	168	168	12	0.00	0.000	0	0		SYS1.PGT1B5.PLPA
COMMON	PGT1B5	001B5	33903	135000	7145	7148	7146	0	0.00	0.000	16	11		SYS1.PGT1B5.COMMON
LOCAL	PGT1B5	001B5	33903	540000	40287	41912	41018	0	5.67	0.004	1933	12191	Y	SYS1.PGT1B5.LOCAL
LOCAL	PGT80A	0080A	33903	540000	42338	43947	43012	0	4.67	0.003	2069	13129	Y	SYS1.PGT80A.LOCAL
LOCAL	PGT80C	0080C	33903	540000	41388	43197	42128	0	6.22	0.004	2106	13299	Y	SYS1.PGT80C.LOCAL
LOCAL	PGT857	00857	33903	540000	39245	40736	39858	0	5.33	0.004	1987	12575	Y	SYS1.PGT857.LOCAL
LOCAL	PGT859	00859	33903	540000	39809	41273	40438	0	5.89	0.004	1969	12546	Y	SYS1.PGT859.LOCAL
SCM	N/A	N/A	N/A	131072	58500	58501	58500	0	0.00	0.000	106	106		N/A

Figure 227. PAGESP Report

Data Not Available: When a page data set comes online during a report interval, an asterisk is placed next to its name and the following message appears instead of measurement data: NOW AVAILABLE FOR SYSTEM USE.

When a page data set has been deleted during a report interval, an asterisk is placed next to its name and the following message appears instead of measurement data: DATA SET DELETED.

Duration Report: If you have specified a duration report, certain fields (DEV NUM, VOLUME SERIAL, DEVICE TYPE and SLOTS ALLOC) might be distorted due to a lengthy duration interval. When such a change occurs, it is not reflected in the duration report; these fields are set according to the contents of the first type 75 SMF record encountered.

Table 192. Fields in the Page Data Set Activity report	
Field Heading	Meaning
PAGE SPACE TYPE	Page space type, which can be PLPA, COMMON, LOCAL, or SCM (Storage Class Memory).
VOLUME SERIAL	Volume serial number of the volume on which the data set resides. N/A is displayed for page space type SCM.
DEV NUM	Number of the device on which the data set resides. N/A for page space type SCM.
DEVICE TYPE	Device type on which the data set resides. N/A for page space type SCM.
SLOTS ALLOC	The total number of slots each page data set contains. For page space type SCM, the total number of SCM 4K blocks available to ASM is displayed.
SLOTS USED	The number of slots that were being used for paging. For page space type SCM, the total number of SCM 4K blocks in-use by ASM is displayed.
BAD SLOTS	Number of slots that encountered permanent I/O errors. For page space type SCM, the total number of SCM 4K blocks in error is displayed.
% 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. $\% \text{ IN USE} = \frac{\# \text{ Busy Samples}}{\# \text{ Samples}} * 100$
PAGE TRANS TIME	Average number of seconds required to complete a page transfer. $\text{PAGE TRANS TIME} = \frac{(\text{USE} * \text{INT}) / \text{N}}{\text{XFER}}$ 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 or SCM in units of 4K pages.
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

Table 192. Fields in the Page Data Set Activity report (continued)	
Field Heading	Meaning
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. Note: 1. If a data set was dynamically introduced during the interval, its data set name is preceded by an asterisk (*). 2. When the operating system has detected errors in a data set that prevents its further use, the name of the data set is preceded by two asterisks (**). ASM continues to access the data set in read-only mode, and RMF reports this activity. 3. N/A for page space type SCM.

Overview reference

Table 193. Overview names in the Page Data Set Activity report	
Field Heading or Meaning	Overview Name
SLOTS USED - AVG	PSAVGSL
BAD SLOTS	PSBADS
% IN USE	PSBSY
PAGE TRANS TIME	PSPTT
NUMBER IO REQ	PSART
PAGES XFER'D	PSPT

PAGING - Paging Activity report

The Paging Activity report provides information about the demands made on the system paging facilities and the use of central storage and external page storage during the interval.

How to request this report

Monitor I gathers data for this report automatically. If you want to suppress gathering, you need to specify NOPAGING.

To produce this report, specify

```
REPORTS (PAGING)
```

This report is also available in XML output format. *How to work with Postprocessor XML reports in z/OS Resource Measurement Facility User's Guide* provides all required information on how to produce and view XML reports.

Example URL for the DDS API

```
http://ddshost:8803/gpm/rmfpp.xml?reports=PAGING
```

Using the information given in the report

If the non-swap, non-VIO page fault rate (page-ins) is excessively high, it could be the result of over-commitment of central storage.

Other problems to look for are high pageable system area non-swap page-in rates, which could be caused by a poor pack list or a large number of fixed LPA modules. A period of high VIO slot use could be a sign that a specific job is making excessive use of VIO. Always be alert for bad slots because they can cause executing jobs to end abnormally.

Contents of the report

The **Paging Activity** report is formatted as follows:

- CENTRAL STORAGE PAGING RATES
- CENTRAL STORAGE MOVEMENT AND REQUEST RATES
- FRAME AND SLOT COUNTS
- MEMORY OBJECTS AND HIGH VIRTUAL STORAGE FRAMES
- DEDICATED MEMORY

The headers of the sections include the **OPT** field which shows the name of the active option member IEAOPTxx. The option member contains parameters that affect system resource manager (SRM) decisions.

Note: The FRAME AND SLOT COUNTS section is displayed on one report page together with the CENTRAL STORAGE MOVEMENT AND REQUEST RATES section ([Figure 229 on page 437](#)), and therefore does not include the **OPT** field.

Central Storage Paging Rates

This section of the **Paging Activity** report monitors paging rates in central storage below the 2 GB bar. The paging rates monitored are organized into two major groups:

- Page-in rates
- Page-out rates

The page-in and page-out groups are further divided into:

- Swap
- Non-swap (for the page-in group additionally divided into: Block, Non-Block)
- Total (rate and percentage)

All of the above paging data rates appear for one or more of the following:

- Pageable system areas used for non-VIO data, broken down into LPA and CSA
- Address space pages used, broken down into hyperspace data, VIO data and non-VIO data

The rate of page movement within central storage below the 2 GB bar is shown in the bottom left corner of the page.

PAGING ACTIVITY										PAGE
z/OS 3.1		SYSTEM ID R71 RPT VERSION 3.1 RMF				DATE 09/30/2021 TIME 10.35.00		INTERVAL 05.00.000 CYCLE 1.000 SECONDS		
OPT = IEAOPT00 LFAREA SIZE = 209715200 CENTRAL STORAGE PAGING RATES - IN PAGES PER SECOND										

		PAGE IN		PAGE OUT						
		-- NON SWAP --		-- TOTAL --						
CATEGORY	SWAP	BLOCK	NON BLOCK	RATE	%	SWAP	NON SWAP	RATE	%	

PAGEABLE SYSTEM AREAS (NON-VIO)										
LPA		0.00	0.01	0.01	100					
CSA		0.00	0.00	0.00	0	0.00	0.00	0		
SUM		0.00	0.01	0.01	100	0.00	0.00	0		

ADDRESS SPACES										
HIPERSPACE		0.00	0.00	0.00	0	0.00	0.00	0		
VIO		0.00	0.00	0.00	0	0.00	0.00	0		
NON-VIO	0.00	0.00	0.00	0.00	0	0.00	0.00	0.00	0	
SUM	0.00	0.00	0.00	0.00	0	0.00	0.00	0.00	0	

TOTAL SYSTEM										
HIPERSPACE		0.00	0.00	0.00	0	0.00	0.00	0		
VIO		0.00	0.00	0.00	0	0.00	0.00	0.00	0	
NON-VIO	0.00	0.00	0.01	0.01	100	0.00	0.00	0.00	0	
SUM	0.00	0.00	0.01	0.01	100	0.00	0.00	0.00	0	

SHARED										
		0.00	0.00	0.00	0	0.00	0.00	0.00	0	

PAGE MOVEMENT WITHIN CENTRAL STORAGE					18.09					
PAGE MOVEMENT TIME %					0.0					
AVERAGE NUMBER OF PAGES PER BLOCK					0.0					
BLOCKS PER SECOND					0.00					
PAGE-IN EVENTS (PAGE FAULT RATE)					0.01					

Figure 228. PAGING Report - Central Storage Paging Rates

Table 194. Fields in the Paging Activity report - Central Storage Paging Rates

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 194. Fields in the Paging Activity report - Central Storage Paging Rates (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. TOTAL % The percentage of the total page-in rate for each part of the total.
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. TOTAL % The percentage of the total page-out rate for each part of the total.
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.

Table 194. Fields in the Paging Activity report - Central Storage Paging Rates (continued)	
Field Heading	Meaning
PAGE MOVEMENT TIME %	The percentage of general purpose processor time, including normalized AAP, CBP and IIP times, spent on page movement to obtain or free a particular type of frame for a page to be fixed (that is, a frame below the 16 megabyte line in central processor storage). The calculation is the amount of processor time needed to steal the page (including the time to move the contents of the frames, but not the time to move the new contents into the frames) divided by the length of the interval.
AVERAGE NUMBER OF PAGES PER BLOCK	The average size of address space non-VIO blocks that were paged-in during the interval. It does not include swap or hiperspace pages.
BLOCKS PER SECOND	The rate of page faults for pages that were part of a block.
PAGE-IN EVENTS (PAGE FAULT RATE)	The rate of page faults for all pages in events per second, excluding VIO and Hiperspace. The rate includes pages read from DASD only, not from expanded storage.

Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*. The following table shows the overview condition names for the Overview report.

Table 195. Overview names in the Paging Activity report - Central Storage Paging Rates	
Field Heading or Meaning	Overview Name
PAGE MOVEMENT WITHIN CENTRAL STORAGE	PGMVRT

Central Storage Movement and Request Rates

The Central Storage Movement and Request Rates section provides paging information about hiperspace and VIO pages and about various types of storage requests.

PAGING ACTIVITY										
z/OS 3.1		SYSTEM ID TRX2 RPT VERSION 3.1 RMF		DATE 09/30/2021 TIME 11.20.00		INTERVAL 05.00.000 CYCLE 1.000 SECONDS		PAGE 2		
OPT = IEAOPT00 LFAREA SIZE = 209715200 CENTRAL STORAGE MOVEMENT AND REQUEST RATES - IN PAGES PER SECOND										

SYSTEM UIC: MIN = 65535		MAX = 65535		AVG = 65535						
CENTRAL STORAGE		PAGE WRITE		PAGE READ		FRAME COUNTS -----				
-----		-- RATE --		-- RATE --		-- MIN --		MAX --		AVG --
HIPERSPACE		0.00		0.00		2		2		2
VIO		0.00		0.00		0		0		0
-----		GETMAIN -----		FIXED -----		REF FAULTS ---				
STORAGE REQUESTS		REQUESTS		FRAMES		BACKED		REQ < 2GB		1ST NON-1ST
RATE		596.37		FRAMES		145.45		2.74 91.51		704.03 0.00

FRAME AND SLOT COUNTS										

(31 SAMPLES)										
CENTRAL STORAGE FRAMES		TOTAL		AVAILABLE		SQA		LPA		CSA
-----		-----		-----		-----		-----		LSQA
MIN		2,621,312		2,268,509		5,702		18,958		5,400
MAX		2,621,312		1,279,645		5,737		18,964		5,412
AVG		2,621,312		1,787,259		5,724		18,962		5,406

FIXED FRAMES		TOTAL		NUCLEUS		SQA		LPA		CSA
-----		-----		-----		-----		-----		LSQA
MIN		41,531		3,165		5,108		71		17,240
MAX		1,107,806		3,165		5,143		71		27,480
AVG		576,679		3,165		5,130		71		20,745

SHARED FRAMES / SLOTS		TOTAL		CENTRAL STORAGE		FIXED TOT		FIXED BEL		HV 1M
-----		-----		-----		-----		-----		HV 4K
MIN		1,061,428		179,907		58		0		674
MAX		1,061,529		180,008		58		0		674
AVG		1,061,516		179,995		58		0		674
LOCAL PAGE DATA SET SLOTS		TOTAL		AVAILABLE		BAD		NON-VIO		VIO
-----		-----		-----		-----		-----		-----
MIN		1,802,699		1,802,699		0		0		0
MAX		1,802,699		1,802,699		0		0		0
AVG		1,802,699		1,802,699		0		0		0
SCM PAGING BLOCKS		TOTAL		AVAILABLE		BAD		IN-USE		
-----		-----		-----		-----		-----		-----
MIN		8,388,608		8,372,612		0		15,996		
MAX		8,388,608		8,372,612		0		15,996		
AVG		8,388,608		8,372,612		0		15,996		

Figure 229. PAGING Report - Central Storage Movement and Request Rates / Frame and Slot Counts

Table 196. Fields in the Paging Activity report - Central Storage Movement and Request Rates	
Field Heading	Meaning
SYSTEM UIC (MIN, MAX, AVG)	The minimum, maximum and average system high unreferenced interval count. The maximum SYSTEM UIC value can indicate contention for central storage frames. When the SYSTEM UIC is relatively low, contention for central storage is high. Although total paging rates might vary with the type and level of workload, the unreferenced interval count is the best indicator of actual storage contention.
The Central Storage Movement and Request Rates section consists of two categories: CENTRAL STORAGE and STORAGE REQUESTS. The CENTRAL STORAGE category displays the following counts:	
PAGE WRITE RATE	HIPERSPACE Rate of hiperspace pages written to central storage. VIO Rate of VIO pages written to central storage.
PAGE READ RATE	HIPERSPACE Rate of hiperspace pages read from central storage. VIO Rate of VIO pages read from central storage.
FRAME COUNTS	MIN, MAX, and AVG of allocated frame counts. HIPERSPACE Storage frame counts allocated to hiperspace. VIO Storage frame counts allocated to VIO address space.
The STORAGE REQUESTS category displays the following counts:	
GETMAIN	REQUESTS GETMAIN request rate FRAMES BACKED Rate of pages backed during GETMAIN requests.
FIXED	REQ < 2 GB Rate of fix requests issued for storage below 2 GB (address space only). FRAMES < 2 GB Rate of pages requested to be fixed for storage below 2 GB (address space only).
REF FAULTS	1ST First page reference faults rate. NON-1ST Non-first page reference faults rate.

Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*. The following table shows the overview condition names for the Overview report.

Table 197. Overview names in the Paging Activity report - Central Storage Movement and Request Rates	
Field Heading or Meaning	Overview Name
SYSTEM UIC - MAX	MXHUIC
SYSTEM UIC - AVG	AVGHUIC
Overview names in the CENTRAL STORAGE category:	
PAGE WRITE RATE - HIPERSPACE	RSHSPW
PAGE WRITE RATE - VIO	RSVIOW
PAGE READ RATE - HIPERSPACE	RSHSPR

Table 197. Overview names in the Paging Activity report - Central Storage Movement and Request Rates (continued)

Field Heading or Meaning	Overview Name
PAGE READ RATE - VIO	RSVIOR
FRAME COUNTS - HIPERSPACE - MIN	RSHSPM
FRAME COUNTS - HIPERSPACE - MAX	RSHSPX
FRAME COUNTS - HIPERSPACE - AVG	RSHSPA
FRAME COUNTS - VIO- MIN	RSVIOM
FRAME COUNTS - VIO- MAX	RSVIOX
FRAME COUNTS - VIO- AVG	RSVIOA

Frame and Slot Counts

This section of the Paging Activity report (included in [Figure 229 on page 437](#)) shows information about the following storage related categories:

- CENTRAL STORAGE FRAMES
- FIXED FRAMES
- SHARED FRAMES
- LOCAL PAGE DATA SET SLOTS
- SCM PAGING BLOCKS

All values are presented as MIN (minimum), MAX (maximum), and AVG (average).

Table 198. Fields in the Paging Activity report - Frame and Slot Counts

Field Heading	Meaning
SAMPLES	The number of valid samples taken in this interval is shown in the upper left corner of this report section in various formats: • as (nn SAMPLES) if all samples are valid • as SAMPLES = xx VALID SAMPLES = yy VALID SAMPLES CSA/REGION = zz if there are invalid samples and therefore the number of valid samples is less than the number of samples. In this case, the number of valid samples, and the number of valid samples for CSA and REGION+SWA values is also displayed to indicate that some of the CENTRAL STORAGE and FIXED FRAMES counts are based on less data.

Table 198. Fields in the Paging Activity report - Frame and Slot Counts (continued)

Field Heading	Meaning
CENTRAL STORAGE FRAMES	TOTAL The total number of central storage frames in the system. AVAILABLE The number of central storage frames that are not in-use by the system. SQA, LPA, CSA, LSQA, REGIONS+SWA, HV SHARED, HV COMMON These columns show the number of central storage frames that are in-use by each of these areas. The CSA value includes the number of frames used by the restricted use common service area (RUCSA). The REGIONS+SWA value also includes the number of frames used by high virtual private storage. The value of the TOTAL count is not a summation of the AVAILABLE, SQA, LPA, CSA, LSQA, REGIONS+SWA, HV SHARED (high virtual SHARED), and HV COMMON (high virtual COMMON) counts at the end of the interval, but is derived by adding these counts from each valid sample and then reporting the MIN sum, MAX sum, and AVG sum for the complete set of samples. If data is not available for any of the SQA, LPA, CSA, LSQA, or REGIONS+SWA counts, the following text appears across these columns: *** NO COUNTS AVAILABLE *** If there are no valid samples for CSA and REGION+SWA values, 'NO DATA' is displayed in these columns. Note: The actual maximum or minimum value of a field might occur at a time when RMF is not sampling.
FIXED FRAMES	TOTAL The total number of central storage frames in the system that are in-use by fixed pages. NUCLEUS, SQA, LPA, CSA, LSQA, REGIONS+SWA, <16MB, 16MB-2GB These columns show the number of central storage frames that are in-use by fixed pages allocated in each of these areas. The CSA value also includes frames used by user-key common fixed pages allocated in RUCSA and by fixed pages allocated in High Virtual Common (HV COMMON). The REGIONS+SWA value also includes the number of frames used by high virtual private storage (4K, 1M, and 2G pages). The SQA value also includes fixed CSA pages. However, pageable CSA pages that have been fixed after allocation are reported in the CSA category. If data is not available for any of the SQA, LPA, CSA, LSQA, or REGIONS+SWA counts, the following text appears across these columns: *** NO COUNTS AVAILABLE *** If there are no valid samples for CSA and REGION+SWA values, 'NO DATA' is displayed in these columns.

Table 198. Fields in the Paging Activity report - Frame and Slot Counts (continued)	
Field Heading	Meaning
SHARED FRAMES / SLOTS	TOTAL The total number of central storage frames and auxiliary slots that are in-use by shared pages. CENTRAL STORAGE The total number of central storage frames that are in-use by shared pages. FIXED TOT The number of central storage frames that are in-use by shared fixed pages. FIXED BEL The number of central storage frames that are in-use by shared fixed pages allocated below 16 megabytes. HV 1M The number of 1 MB pages that are in-use by shared memory objects which are backed in central storage. HV 4K The number of central storage frames that are in-use by shared high virtual 4K pages. AUX DASD The number of shared pages backed on DASD. AUX SCM The number of shared pages backed on Storage Class Memory (SCM). Note: AUX DASD and AUX SCM are also called auxiliary storage slots.
LOCAL PAGE DATA SET SLOTS	TOTAL Total number of page data set slots. AVAILABLE Number of page data set slots that do not contain any data pages and that are available for use. BAD 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. NON-VIO Number of local page data set slots that contain pages belonging to address-space virtual storage. VIO Number of local page data set slots that contain pages for VIO data sets.
SCM PAGING BLOCKS	TOTAL The total number of 4K SCM paging blocks. AVAILABLE The number of SCM blocks that do not contain any data and are available to ASM. BAD The number of SCM blocks that do not contain any data and are unavailable for use. IN-USE The number of SCM blocks that are in-use by ASM.

Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*. The following table shows the overview condition names for the Overview report.

Table 199. Overview names in the Paging Activity report - Frame and Slot Counts	
Field Heading or Meaning	Overview Name
CENTRAL STORAGE FRAMES - CSA / RUCSA - MAX	MXCSAT
CENTRAL STORAGE FRAMES - CSA / RUCSA - AVG	AVGCSAT
FIXED FRAMES - SQA - MAX	MXSQA
FIXED FRAMES - SQA - AVG	AVGSQA

Table 199. Overview names in the Paging Activity report - Frame and Slot Counts (continued)	
Field Heading or Meaning	Overview Name
FIXED FRAMES - CSA / RUCSA - MAX	MXCSAF
FIXED FRAMES - CSA / RUCSA - AVG	AVGCSAF
LOCAL PAGE DATA SET SLOTS - VIO - MAX	MAXVIOF
LOCAL PAGE DATA SET SLOTS - VIO - AVG	AVGVIOF
FIXED FRAMES - 16MB-2GB - MIN	FXBETWM
FIXED FRAMES - 16MB-2GB - MAX	FXBETWX
FIXED FRAMES - 16MB-2GB - AVG	FXBETWA
SHARED FRAMES/SLOTS - TOTAL - AVG	SHRPT
SHARED FRAMES/SLOTS - CENTRAL STORAGE - AVG	SHRPC
SHARED FRAMES/SLOTS - FIXED TOT - AVG	SHRPF
SHARED FRAMES/SLOTS - FIXED BEL - AVG	SHRPB
SHARED FRAMES/SLOTS - HV 1 MB - AVG	SFR1MA
SHARED FRAMES/SLOTS - HV 4K - AVG	SFR4KA
SHARED FRAMES/SLOTS - AUX DASD - AVG	SHRPA
SHARED FRAMES/SLOTS - AUX SCM - AVG	SHRPASCM

Memory Objects and High Virtual Storage Frames

Figure 230 on page 442 shows a sample of the MEMORY OBJECTS AND HIGH VIRTUAL STORAGE FRAMES section if Enhanced DAT Facility 2 is available. Information about 1 MB frames and memory objects that can be backed by 1 MB frames is only available if Enhanced DAT Facility 1 is installed. Information about 2 GB frames and memory objects that are backed by 2 GB frames is only available if Enhanced DAT Facility 2 is installed.

All values in this section of the **Paging Activity** report are presented as MIN, MAX, and AVG values.

PAGING ACTIVITY									
z/OS 3.1		SYSTEM ID SYSF		DATE 09/30/2021		INTERVAL 10.29.700		PAGE 3	
OPT = IEA0PT00		RPT VERSION 3.1 RMF		TIME 10.00.00		CYCLE 1.000 SECONDS			
MEMORY OBJECTS AND HIGH VIRTUAL STORAGE FRAMES									

LFAREA		MAXIMUM							
1 MB FRAMES		512M							
2 GB FRAMES		12G							
MEMORY OBJECTS		FIXED 1M	FIXED 2G	COMMON	SHARED	SHARED 1M			

MIN		1	1	95	3	1			
MAX		2	2	96	3	1			
AVG		1	1	95	3	1			
1 MB FRAMES		-----							
		FIXED			PAGEABLE	AVAILABLE	TOTAL		

MIN		512	472	30	123	320,993	321,180		
MAX		512	482	40	125	320,995	321,180		
AVG		512	479	33	124	320,994	321,180		
2 GB FRAMES		-----							
		MAXIMUM	AVAILABLE	IN-USE					

MIN		6	4	1					
MAX		6	5	2					
AVG		6	5	1					
HIGH SHARED FRAMES		TOTAL	CENTRAL STORAGE	BACKED 1M	AUX DASD		AUX SCM		

MIN		136902.1M	605	1	0		N/A		
MAX		136902.1M	606	1	0		N/A		
AVG		136902.1M	605	1	0		N/A		
HIGH COMMON FRAMES		TOTAL	CENTRAL STORAGE	BACKED 1M	FIXED	FIXED 1M	AUX DASD	AUX SCM	

MIN		17301504	32,982	70	12,773	30	0	N/A	
MAX		17301504	35,634	80	12,773	40	0	N/A	
AVG		17301504	33,853	73	12,773	33	0	N/A	

Figure 230. PAGING Report - Memory Objects and High Virtual Storage Frames

Table 200. Fields in the Paging Activity report - Memory Objects and High Virtual Storage Frames	
Field Heading	Meaning
LFAREA MAXIMUM	Amount of real storage available for 1 MB and 2GB pages as specified in the LFAREA parameter.

Table 200. Fields in the Paging Activity report - Memory Objects and High Virtual Storage Frames (continued)

Field Heading	Meaning
MEMORY OBJECTS	FIXED 1M Number of fixed memory objects that are allocated in the system and can be backed in 1 MB frames. FIXED 2G Number of fixed memory objects that are allocated in the system and are backed in 2 GB frames. COMMON Number of memory objects allocated in the high virtual common storage of the system. SHARED Number of memory objects allocated in the high virtual shared storage of the system. SHARED 1M Number of shared memory objects that are allocated in the system and can be backed in 1 MB frames.
1 MB FRAMES	FIXED MAXIMUM Maximum number of 1 MB frames that can be used by fixed 1 MB pages. This value is specified in the LFAREA parameter. FIXED AVAILABLE Number of available 1 MB frames that can be used by fixed 1 MB pages. FIXED IN-USE Number of 1 MB frames that are in-use by fixed 1MB pages. PAGEABLE Number of 1 MB frames that are in-use by pageable 1 MB pages. AVAILABLE Number of available 1 MB frames in central storage. TOTAL Total number of 1 MB frames in central storage.
2 GB FRAMES - FIXED	MAXIMUM Maximum number of 2 GB frames that can be used by fixed 2 GB pages. This value is specified in the LFAREA parameter. AVAILABLE Number of available 2 GB frames that can be used by fixed 2 GB pages. IN-USE Number of 2 GB frames that are in-use by fixed memory objects.
HIGH SHARED FRAMES	TOTAL Size of high virtual shared area in units of 4 KB pages. CENTRAL STORAGE Number of pages from high virtual shared storage that are backed in central storage (in units of 4 KB). BACKED 1M Number of 1 MB pages that are in-use by shared memory objects which are backed in central storage. AUX DASD Number of auxiliary storage slots used for high virtual shared pages that are backed on DASD. AUX SCM Number of auxiliary storage slots used for high virtual shared pages that are backed on SCM storage.

Table 200. Fields in the Paging Activity report - Memory Objects and High Virtual Storage Frames (continued)

Field Heading	Meaning
HIGH COMMON FRAMES	TOTAL Size of high virtual common area in units of 4 KB pages. CENTRAL STORAGE Number of pages from high virtual common storage that are backed in central storage (in units of 4 KB). BACKED 1M Number of high virtual common memory 1 MB pages that are backed in central storage. FIXED Number of pages from high virtual common storage that are fixed in central storage (in units of 4 KB). FIXED 1M Number of high virtual common memory 1 MB pages that are fixed in central storage. AUX DASD Number of auxiliary storage slots used for high virtual common pages that are backed on DASD. AUX SCM Number of auxiliary storage slots used for high virtual common pages that are backed on SCM storage.

Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*. The following table shows the overview condition names for the Overview report.

Table 201. Overview names in the Paging Activity report - Memory Objects and High Virtual Storage Frames

Field Heading or Meaning	Overview Name
MEMORY OBJECTS - FIXED 1M AVG	LMOA
MEMORY OBJECTS - FIXED 2G AVG	GMOA
MEMORY OBJECTS - COMMON AVG	CMOA
MEMORY OBJECTS - SHARED AVG	SMOA
MEMORY OBJECTS - SHARED 1M AVG	SMO1MA
1 MB FRAMES - FIXED MAXIMUM AVG	LFFRTA
1 MB FRAMES - FIXED AVAILABLE AVG	LFFRNA
1 MB FRAMES - FIXED IN-USE AVG	LFRA
1 MB FRAMES - TOTAL AVG	LTFRTA
1 MB FRAMES - AVAILABLE AVG	LTFRAA
1 MB FRAMES - PAGEABLE AVG	LPFRPA
2 GB FRAMES - FIXED MAXIMUM AVG	GFFRTA
2 GB FRAMES - FIXED AVAILABLE AVG	GFFRNA
2 GB FRAMES - FIXED IN-USE AVG	GFFRUA
HIGH SHARED FRAMES - TOTAL AVG	SFRTA
HIGH SHARED FRAMES - CENTRAL STORAGE AVG	SFRA
HIGH SHARED FRAMES - BACKED 1M AVG	SFR1MA
HIGH SHARED FRAMES - AUX DASD AVG	SAUXSA
HIGH SHARED FRAMES - AUX SCM AVG	SAUXSSA
HIGH COMMON FRAMES - TOTAL AVG	CFRTA
HIGH COMMON FRAMES - CENTRAL STORAGE AVG	CFRA

Table 201. Overview names in the Paging Activity report - Memory Objects and High Virtual Storage Frames (continued)	
Field Heading or Meaning	Overview Name
HIGH COMMON FRAMES - BACKED 1M AVG	CFR1MA
HIGH COMMON FRAMES - FIXED AVG	CFFRA
HIGH COMMON FRAMES - FIXED 1M AVG	CFFR1MA
HIGH COMMON FRAMES - AUX DASD AVG	CAUXSA
HIGH COMMON FRAMES - AUX SCM AVG	CAUXSSA

Dedicated Memory

The Dedicated Memory section provides a system-level view of the Dedicated Memory usage counts when they are available.

All values in this section of the **Paging Activity** report are presented as MIN, MAX, and AVG values.

Figure 231. PAGING Report - Dedicated memory

P A G I N G A C T I V I T Y									
z/OS 3.1		SYSTEM ID RSB6		DATE 04/08/2024		INTERVAL 14.59.879		PAGE 4	
OPT = IEAOPTRS				TIME 18.59.33		CYCLE 1.000 SECONDS			
DEDICATED MEMORY									

TOTAL ASSIGNABLE		2							
USER ADDRESS SPACES									

MIN		2							
MAX		2							
AVG		2							
2G UNITS	TOTAL	ONLINE	AVAIL	REQD	MIN REQD	MIN ASGN	USED		

MIN	3	3	0	2	2	2	0		
MAX	3	3	0	2	2	2	0		
AVG	3	3	0	2	2	2	0		
1M FRAMES	FIXED	PAGEABLE							

MIN	0	1							
MAX	0	1							
AVG	0	1							
4K FRAMES	DAT	USED							

MIN	3,287	533,824							
MAX	3,287	533,824							
AVG	3,287	533,824							
***** BOTTOM OF DATA *****									

Table 202. Fields in PAGING Report - Dedicated memory	
Field Heading	Meaning
TOTAL ASSIGNABLE	Total amount of dedicated memory at system initialization that can be used by address spaces in 2G units.
USER ADDRESS SPACES	The number of address spaces using dedicated memory.

Table 202. Fields in PAGING Report - Dedicated memory (continued)

Field Heading	Meaning
2G UNITS	TOTAL The amount of online and offline dedicated memory in 2G units, including dedicated memory used by the system. ONLINE The amount of online dedicated memory in 2G units, including dedicated memory used by the system. AVAIL The amount of available dedicated memory in 2G units that can be used by address spaces. REQD The amount of requested dedicated memory in 2G units over all address spaces. MIN REQD The minimum amount of requested dedicated memory in 2G units over all address spaces. MIN ASGN The amount of assigned dedicated memory over all address spaces. USED The amount of dedicated memory used as 2G frames over all address spaces.
1M FRAMES	FIXED The amount of dedicated memory used as 1M fixed frames over all address spaces. PAGEABLE The amount of dedicated memory used as 1M pageable frames over all address spaces.
4K FRAMES	DAT The amount of dedicated memory used as 4K frames over all address spaces. USED The number of frames in 4K units backing DAT tables (system use).

Dedicated Memory: Spreadsheet and Overview reference

You can make **PAGING Report - Dedicated memory** report available in a spreadsheet, using the Spreadsheet Reporter.

The following table shows the overview condition names for the Overview report.

Field Heading or Meaning	Overview Name
AVG JOBS USING DMEM	DMJOBBSA
AVG DMEM 2 GB TOT	GDMTA
AVG DMEM 2 GB TOT ONL	GDMTOA
AVG DMEM 2 GB AVAIL	GDMAA

Field Heading or Meaning	Overview Name
AVG DMEM 2 GB REQD	GDMRA
AVG DMEM 2 GB MIN REQD	GDMMRA
AVG DMEM 2 GB ASSIGNED	GDMSA
AVG DMEM 2 GB USED	GDMUA
AVG DMEM 1 MB USED FIXED	LFDMUA
AVG DMEM 1 MB USED PGBL	LPDMUA
AVG DMEM 4 K USED	DMUA
AVG DMEM 4 K USED DAT	DMUDA

PCIE - PCIE Activity Report

The PCIE Activity Report provides statistics and performance measurements on PCI Express based functions (PCIE functions) allocated by at least one z/OS address space for a period of time within the reporting interval. A PCIE function is captured by the report if one of the following feature activities has been detected:

- RDMA (Remote Direct Memory Access) over Converged Enhanced Ethernet
- zEnterprise Data Compression (zEDC) capability using zEDC Express
- SMC-Direct over Internal Shared Memory (ISM) virtual PCIE function
- IBM zHyperlink

How to request this report

If the currently active SMFPRMxx parameter settings indicate that SMF record type 74 subtype 9 is to be collected, then RMF Monitor III gathers the data required for the PCIE Activity Report into this SMF record.

To produce this report, specify

```
REPORTS(PCIE)
```

This single-system report is only available in XML output format. Therefore, you need to specify the XPRPTS ddname in your Postprocessor job. [How to work with Postprocessor XML reports in z/OS Resource Measurement Facility User's Guide](#) provides all required information on how to produce and view XML reports.

Example URL for the DDS API

```
http://ddshost:8803/gpm/rmfpp.xml?reports=PCIE
```

Contents of the report

The [Figure 232 on page 449](#) is divided into the following sections:

- [General PCIE Activity](#)
- [Hardware Accelerator Activity](#)
- [Hardware Accelerator Compression Activity](#)
- [RoCE Activity](#)
- [Internal Shared Memory Activity](#)
- [Synchronous I/O Link Activity](#)

- Synchronous I/O Response Time Distribution

The General PCIE Activity section shows measurements for all PCIE functions partially dependent from the type of the exploited hardware feature. The measurements reflect the activity of the z/OS system on which RMF data collection took place. They comprise data rates about the communication of z/OS programs with PCIE functions by means of PCI operations that are transferring data blocks from z/OS to the PCIE function (PCI LOAD, PCI STORE, PCI STORE BLOCK, and REFRESH PCI TRANSLATIONS) as well as Read/Write Transfer data rates.

The Hardware Accelerator Activity section and the Hardware Accelerator Compression Activity section have single system scope and are leveraging the measurements displayed in the General PCIE Activity section. They are only displayed if the hardware feature zEnterprise Data Compression (zEDC) is used for compression acceleration. In this case, they display:

- common accelerator metrics, for example, total request execution time, or the amount of transferred data
- compression specific metrics, for example, the amount of compressed data and the number and throughput of compression requests
- device driver buffer statistics

The RoCE Activity section has system scope and displays a function specific view for RoCE functions. The section displays:

- interconnection metrics, for example, the physical network ID port 1 and 2
- transfer metrics, for example, read and write rate and packet based transmission metrics

The Internal Shared Memory Activity section has system scope and displays a function specific view for Internal Shared Memory (ISM) functions. The section displays:

- the physical network ID
- the write transfer rate

The Synchronous I/O Link Activity section has either system scope (metrics showing values per function) or CEC scope (metrics showing values with a CPC wide view). For IBM zHyperLinks the section displays:

- interconnection metrics, for example, the port ID, serial number, type and model of the storage controller the synchronous I/O link is connected to
- transfer metrics, for example, a data transfer rate
- metrics on requests processed, for example, the percentage of successful requests
- utilization metrics, for example, the time busy percentage

The Synchronous I/O Response Time Distribution subsection provides a response time distribution overview for read and write instructions executed on the allocated synchronous I/O links.

▼ General PCIE Activity

Function ID	Function CHID	Function Name	Function Status	Owner Job Name	Owner Address Space ID	Function Allocation Time	PCI Load Operations Rate	PCI Store Operations Rate	PCI Store Block Operations Rate	Refresh PCI Translations Operations Rate	Data Transfer Rate
0022	0540	Hardware Accelerator	Allocated	FPGHWAM	0013	900	0	143	0	16.7	
0023	05FC	Hardware Accelerator	Allocated	FPGHWAM	0013	900	0	143	0	16.7	
0080	057C	10GbE RoCE	Allocated	VTAM390	0029	900	0.907	5869	0	0.417	4629
0430	0144	8GB zHyperLink	Allocated	IOSAS	001A	900					22.8
0579	07C0	ISM	Allocated	VTAM	002B	900	0	0	428022	0	1413

▼ Hardware Accelerator Activity

Function ID	Time Busy %	Adapter Utilization	Work Units Processed Rate	Request Execution Time	Std Dev for Request Execution Time	Request Queue Time	Std Dev for Request Queue Time	Request Size
0022	0.395	0.127	524924K	27.7	7.56	15.0	48.6	47.0
0023	0.352	0.123	3335654K	24.7	5.74	12.4	27.7	47.2

▼ Hardware Accelerator Compression Activity

Function ID	Compression Request Rate	Compression Throughput	Compression Ratio	Decompression Request Rate	Decompression Throughput	Decompression Ratio	Buffer Pool Size	Buffer Pool Utilization
0022	142	3.99	1.45	0.451	0.010	0.405	16	0
0023	0.451	0.031	2.73	142	2.75	0.688	16	0

▼ RoCE Activity

Function ID	Physical Network ID Port 1	Physical Network ID Port 2	Read Transfer Rate	Write Transfer Rate	Packets Received Rate	Packets Transmitted Rate
0080	TSTNET01	TSTNET11	2629	1999	0	0

▼ Internal Shared Memory Activity

Function ID	Physical Network ID	Write Transfer Rate
0579	NETWORK7C0	1413

▼ Synchronous I/O Link Activity

Function ID	Function CHID	Port ID	Serial Number	Type-Model	Total Request Rate	Total Request Rate (CPC)	Successful Request %	Successful Request % (CPC)	Read Transfer Rate	Read Transfer Rate (CPC)	Read Transfer Ratio	Read Transfer Ratio (CPC)	Write Transfer Rate
0430	0144	1	0000000GYF81	002107-985	5291	21163	100.0	100.0	21.7	86.7	0.004	0.004	1.06

▼ Synchronous I/O Response Time Distribution

Function ID	% Read < 20usec	% Read < 30usec	% Read < 40usec	% Read < 50usec	% Read < 60usec	% Read < 70usec	% Read < 80usec	% Read < 90usec	% Read < 100usec	% Read >=100usec	% Write < 20usec	% Write < 30usec	% Write < 40usec	% Write < 50usec
0430	16.7	83.2	0.074	0.003	<.001	<.001	<.001	0.001	<.001	0.002	0	0	0	52.0

Figure 232. Postprocessor PCIE Activity Report

Table 203. Fields in the PCIE Activity Report

Field Heading	Meaning
General PCIE Activity	
Function ID	Identifier of the monitored PCIE function.
Function CHID	Physical or virtual channel identifier for the PCIE function.
Function Name	Device name for the PCIE function.
Function Status	The PCIE function status can be one of the following: Allocated The function is allocated and in use at the end of the reporting interval. Re-Allocated The function was de-allocated during the interval but has been re-allocated again. It is in use at the end of the reporting interval. De-Allocated The function was de-allocated during the interval and is unused at the end of the reporting interval. De-Allocate-Pending The function is in the process of de-allocation. Error The function is in permanent error. Unknown The function status is unknown.
Owner Job Name	Job name of the owner who allocated the PCIE function.
Owner Address Space ID	Address space ID of the owner who allocated the PCIE function.
Function Allocation Time	Time in seconds for which the PCIE function was allocated or de-allocate-pending during this interval.
PCI Load Operations Rate	Rate of PCI Load operations executed during the reporting interval. This value is not reported for synchronous I/O functions.
PCI Store Operations Rate	Rate of PCI Store operations executed during the reporting interval. This value is not reported for synchronous I/O functions.
PCI Store Block Operations Rate	Rate of PCI Store Block operations executed during the reporting interval. This value is not reported for synchronous I/O functions.
Refresh PCI Translations Operations Rate	Rate of Refresh PCI Translations operations executed during the reporting interval. This value is not reported for synchronous I/O functions.
Data Transfer Rate	The number of megabytes per second transferred to and from the PCIE function. For ISM functions, this value reports the number of megabytes transmitted to the function.
Hardware Accelerator Activity	
Time Busy %	The percentage of time that this partition kept the hardware accelerator busy.
Adapter Utilization	Utilization of the Hardware Accelerator. This value is not reported on zEC12 and zBC12 hardware.
Work Units Processed Rate	The number of work units per second that were processed by the Hardware Accelerator. This value is not reported on zEC12 and zBC12 hardware.
Request Execution Time	The average time in microseconds the hardware accelerator used to process a request.
Std Dev for Request Execution Time	The standard deviation of the request execution time.
Request Queue Time	The average queue time in microseconds that was spent for a request. This value has single system scope but is affected by activity from other partitions sharing the hardware accelerator.
Std Dev for Request Queue Time	The standard deviation of the request queue time.
Request Size	The average number of kilobytes transferred per request.

Table 203. Fields in the PCIE Activity Report (continued)	
Field Heading	Meaning
Hardware Accelerator Compression Activity	
Compression Request Rate	The number of compression requests per second.
Compression Throughput	The number of megabytes compressed per second.
Compression Ratio	The ratio between input and output bytes compressed within this interval.
Decompression Request Rate	The number of decompression requests per second.
Decompression Throughput	The number of megabytes decompressed per second.
Decompression Ratio	The ratio between input and output bytes decompressed within this interval.
Buffer Pool Size	The total size of memory in megabytes that is allocated to the buffer pool.
Buffer Pool Utilization	The average utilization of the buffer pool that z/OS kept for in-use buffers.
RoCE Activity	
Physical Network ID Port 1 and Port 2	Physical-network identifier (PNET ID) that identifies the first or second port of the RoCE / RoCE-2 device.
Read Transfer Rate	The number of megabytes per second that a RoCE / RoCE-2 device received on the external Ethernet interface. On zEC12 or zBC12, this field designates the number of megabytes per second that were transferred by DMA reads from all defined DMA address spaces to the PCIE function.
Write Transfer Rate	The number of megabytes per second transmitted from a RoCE / RoCE-2 device on the external Ethernet interface. On zEC12 or zBC12, this field designates the number of megabytes per second that were transferred by DMA writes from the PCIE function to all defined DMA address spaces.
Packets Received Rate	The number of packets per second that were received on the external Ethernet interface of the RoCE / RoCE-2 device. This value is not reported on zEC12 and zBC12 hardware.
Packets Transmitted Rate	The number of packets per second that were transmitted on the external Ethernet interface of the RoCE / RoCE-2 device. This value is not reported on zEC12 and zBC12 hardware.
Internal Shared Memory Activity	
Physical Network ID	Physical-network identifier (PNET ID) that identifies the port of the virtual PCIE function for Internal Shared Memory (ISM).
Write Transfer Rate	The number of megabytes per second transmitted on the ISM function.
Synchronous I/O Link Activity	
The values collected for the whole CPC are only reported if Global Performance Reporting is enabled in the LPAR image profile of the Hardware Management Console (HMC).	
Port ID	ID of the port on the zHyperlink card where the synchronous I/O link is attached to.
Serial Number	The serial number of the storage controller the synchronous I/O link connects to.
Type-Model	The type and model of the storage controller the synchronous I/O link connects to.
Link ID	The identifier of the synchronous I/O link that is configured in the storage controller.
Total Request Rate	The total number of synchronous I/O requests per second for this function.
Total Request Rate (CPC)	The total number of synchronous I/O requests per second for the synchronous I/O link this function is defined on.
Successful Request %	The percentage of synchronous I/O requests that completed successfully for this function.
Successful Request % (CPC)	The percentage of requests that completed successfully for the synchronous I/O link this function is defined on.
Read Transfer Rate	The number of megabytes per second that were read from the storage controller by this synchronous I/O function.
Read Transfer Rate (CPC)	The number of megabytes per second that were read from the storage controller on the synchronous I/O link this function is defined on.

Table 203. Fields in the PCIE Activity Report (continued)	
Field Heading	Meaning
Read Transfer Ratio	The number of megabytes read per request processed by this synchronous I/O function.
Read Transfer Ratio (CPC)	The number of megabytes read per request processed on the synchronous I/O link this function is defined on.
Write Transfer Rate	The number of megabytes per second that were written to the storage controller by this synchronous I/O function.
Write Transfer Rate (CPC)	The number of megabytes per second that were written to the storage controller on the synchronous I/O link this function is defined on.
Write Transfer Ratio	The number of megabytes written per request processed by this synchronous I/O function.
Write Transfer Ratio (CPC)	The number of megabytes written per request processed on the synchronous I/O link this function is defined on.
Time Busy %	The percentage of time spent on synchronous I/O processing by this synchronous I/O function.
Time Busy % (CPC)	The percentage of time spent on synchronous I/O processing on the synchronous I/O link this function is defined on.
Synchronous I/O Response Time Distribution	
% Read < (>=) n usec	Leftmost read bucket: Percentage of read samples with a response time less than n microseconds. Rightmost read bucket: Percentage of read samples with a response time greater or equal n microseconds. Other read buckets: Percentage of read samples with a response time less than n microseconds and greater or equal to the prior bucket time limit. For example: % Read <30msec = 15.3 means that 15.3 percent of the read samples had a response time of more than or equal to 20 microseconds but less than 30 microseconds.
% Write < (>=) n msec	Leftmost write bucket: Percentage of write samples with a response time less than n microseconds. Rightmost write bucket: Percentage of write samples with a response time greater or equal n microseconds. Other write buckets: Percentage of write samples with a response time less than n microseconds and greater or equal to the prior bucket time limit. For example: % Write <40msec = 10.8 means that 10.8 percent of the write samples had a response time of more than or equal to 30 microseconds but less than 40 microseconds.

Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*. The following table shows the overview condition names for the Overview report.

Table 204. Overview conditions in the PCIE Activity Report	
Field Heading or Meaning	Overview Name
PCI Load Operations Rate (Hardware Accelerator, RoCE and ISM)	PCILOAD
PCI Store Operations Rate (Hardware Accelerator, RoCE and ISM)	PCISTOR
PCI Store Block Operations Rate (Hardware Accelerator, RoCE and ISM)	PCISTBL
Refresh PCI Translations Operations Rate	PCIRPTR

Table 204. Overview conditions in the PCIE Activity Report (continued)	
Field Heading or Meaning	Overview Name
DMA Read Rate (on zEC12 or zBC12 hardware only)	PCIDMAR
DMA Write Rate (on zEC12 or zBC12 hardware only)	PCIDMAW
Number of megabytes received per second (RoCE on z13 or later and Synchronous I/O)	PCIBYTR
Number of megabytes transmitted per second (RoCE on z13 or later , ISM and Synchronous I/O)	PCIBYTT
Number of packets received per second (RoCE on z13 and later only)	PCIPAKR
Number of packets transmitted per second (RoCE on z13 and later only)	PCIPAKT
Number of work units processed per second (zEDC on z13 and later only)	PCIWUP
PCI Function Utilization (zEDC on z13 and later only)	PCIUTIL
Hardware Accelerator Activity	
Time Busy %	FPGBUSY
Request Execution Time	FPGRTIM
Request Queue Time	FPGQTIM
Request Size	FPGBYTR
Transfer Rate Total	FPGBYTS
Hardware Accelerator Compression Activity	
Compression Request Rate	FPGCORS
Compression Throughput	FPGCOBS
Compression Ratio	FPGCORT
Decompression Request Rate	FPGDCRS
Decompression Throughput	FPGDCBS
Decompression Ratio	FPGDCRT
Buffer Pool Size	FPGBPSZ
Buffer Pool Utilization	FPGBPRT
Synchronous I/O Link Activity	
Sync I/O Request Rate	SYNCTR
Sync I/O Successful Request Rate	SYNCSR

SDELAY - Serialization Delay report

In large systems, it may be difficult to detect and debug performance problems due to resource contention. System dumps or traditional performance reports may not be adequate tools to identify the address space that is causing a contention.

For this purpose, RMF provides global resource serialization (GRS) enqueue and latch performance statistics, as well as system suspend lock contention information to help users in analyzing serialization-related performance problems.

How to request this report

RMF Monitor III gathers the data required for the Serialization Delay report by default in SMF record type 72 subtype 5.

To produce this report, specify

REPORTS(SDELAY)

Note: The SDELAY report is only available as an interval report, not as a duration report.

If you do not want to use this report, you should suppress the associated SMF data collection for record type 72-5. Methods how to achieve this are listed in "Defining SMF record writing" in *z/OS Resource Measurement Facility User's Guide*.

This single-system report is only available in XML output format. Therefore, you need to specify the XPRPTS ddname in your Postprocessor job. How to work with Postprocessor XML reports in *z/OS Resource Measurement Facility User's Guide* provides all required information on how to produce and view XML reports.

Example URL for the DDS API

<http://ddshost:8803/gpm/rmfpp.xml?reports=SDELAY&sysid=SYSF>

Contents of the report

The **Serialization Delay** report provides contention information on system and address space level for different types of suspend locks, GRS latches, and GRS ENQs. Reported suspend lock types (with their abbreviations used in the report in parentheses) are: CMS lock (CMS), CMS Enqueue/Dequeue lock (CMSEQDQ), CMS Latch lock (CMSLatch), CMS SMF lock (CMSSMF), LOCAL lock (Local), and CML lock (CML).

The **Serialization Delay Report** consists of two sections:

- the **Serialization Delay Summary** (see [“Serialization Delay Summary” on page 454](#))
- the **Serialization Delay Details** (see [“Serialization Delay Details” on page 456](#))

Serialization Delay Summary

The **Serialization Delay Summary** section contains system-wide summary data for all address spaces and is divided into three subsections:

- The **System Locks** subsection displays summary data for system suspend locks.
- The **GRS Latch Set Creator** subsection displays summary data about GRS latches.
- The **GRS Enqueue** subsection displays summary data about GRS enqueue requests.

RMF Postprocessor Interval Report [System CB89] : Serialization Delay Report



RMF Version : z/OS V2R1 SMF Data : z/OS V2R1

Start : 08/10/2011-10.30.00 End : 08/10/2011-11.00.00 Interval : 30:00:000 minutes Cycle : 1000 milliseconds

Serialization Delay Summary

GRS ModeSTAR

System Locks

Lock Type	Total Contention Time	Avg Contention Time	Total Contention Count	Contention Count with QLen>1
CMS	1	0.07	13	0
CMSEQDQ	5	0.01	275	0
CMSLATCH	0		0	0
CMSSMF	1	0.02	38	4
Local	4864310	0.16	30124648	26068270
CML Owner	4929	1.28	3840	863

GRS Latch Set Creator

Total Contention Time	Avg Contention Time	Std Dev of Contention Time	Total Contention Count
10825	216	912	50

GRS Enqueue

Scope	Total Contention Time	Avg Contention Time	Std Dev of Contention Time	Total Request Count	Total Contention Count
Step	363	3.98	3.01	6158	91
System	954	0.73	0.23	15439	1300
Systems	35119854	25	32	6376290	1356971

Figure 233. SDELAY Report - Serialization Delay Summary

Table 205. Fields in the Serialization Delay Summary section

Field Heading	Meaning
GRS Mode	The operation mode of GRS: - NONE - RING - STAR
System Locks – contains system-wide summary data on system suspend locks for all address spaces.	

Table 205. Fields in the Serialization Delay Summary section (continued)

Field Heading	Meaning
Lock Type	Displays the system suspend lock type: CMS CMS lock CMSEQDQ CMS Enqueue/Dequeue lock CMSLatch CMS Latch lock CMSSMF CMS SMF lock Local LOCAL lock CML Owner CML lock owner
Total Contention Time	The total amount of time in milliseconds that a unit of work was suspended by a lock of the indicated type.
Avg Contention Time	The average amount of time in milliseconds that a unit of work was suspended by a lock of the indicated type.
Total Contention Count	The total number of times that a unit of work was suspended by a lock of the indicated type.
Contention Count with QLen>1	The total number of times that a unit of work was suspended by a lock of the indicated type when there was already at least one other unit of work suspended for the lock (that is, queue length > 1).
GRS Latch Set Creator – contains summary data about GRS latches for all address spaces.	
Total Contention Time	The total amount of time in milliseconds that latch obtain requests were suspended.
Avg Contention Time	The average amount of time in milliseconds that latch obtain requests were suspended.
Std Dev of Contention Time	The standard deviation of the total contention time in milliseconds.
Total Contention Count	The total number of suspended latch obtain requests.
GRS Enqueue – contains summary data about GRS enqueue requests for all address spaces.	
Scope	The scope of an GRS enqueue request: • STEP • SYSTEM • SYSTEMS One line is displayed for requests of a certain scope.
Total Contention Time	The total amount of time in milliseconds that the GRS ENQ requests with the specified Scope were suspended.
Avg Contention Time	The average amount of time in milliseconds that the GRS ENQ requests with the specified Scope were suspended.
Std Dev of Contention Time	The standard deviation of the Total Contention Time in milliseconds.
Total Request Count	The total number of GRS ENQ requests with the specified Scope .
Total Contention Count	The total number of GRS ENQ requests with the specified Scope that were suspended.

Serialization Delay Details

The **Serialization Delay Details** section provides the following information in four subsections:

- The **CMS Lock Details** subsection contains detail data about CMS/CMSEQDQ/CMSLatch/CMSSMF locks per address space (see [Figure 234 on page 457](#)).
- The **CML and Local Lock Details** subsection contains detail data about CML and LOCAL locks per address space (see [Figure 235 on page 458](#)).

- The **GRS Latch Details** subsection contains detail data about GRS latches (see [Figure 236](#) on page 459).
- The **GRS Enqueue Details** subsection contains detail data about GRS enqueue requests (see [Figure 237](#) on page 460).

Note: For each lock type, a maximum of the top twenty address spaces with the longest contention times are reported.




Figure 234. SDELAY Report - Serialization Delay Details - CMS Lock Details

Table 206. Fields in the Serialization Delay Details section - CMS Lock Details

Field Heading	Meaning
CMS Lock Details – contains detail data about CMS/CMSEQDQ/CMSLatch/CMSSMF locks per address space.	
Address Space ID	The hexadecimal address space identifier (ASID) of the job for which lock data was collected.
Jobname	The name of the job.
Service Class Name	The name of the service class that the job has been running in.
Service Class Period	The service class period that the job has been running in.
CMS - Total Contention Time CMSEQDQ - Total Contention Time CMSLatch - Total Contention Time CMSSMF - Total Contention Time	The total amount of time in milliseconds that a unit of work of the indicated address space was suspended on the respective lock type.
CMS - Avg Contention Time CMSEQDQ - Avg Contention Time CMSLatch - Avg Contention Time CMSSMF - Avg Contention Time	The average amount of time in milliseconds that a unit of work of the indicated address space was suspended on the respective lock type.

Table 206. Fields in the Serialization Delay Details section - CMS Lock Details (continued)

Field Heading	Meaning
CMS - Total Contention Count CMSEQDQ - Total Contention Count CMSLatch - Total Contention Count CMSSMF - Total Contention Count	The number of times that a unit of work of the indicated address space was suspended on the respective lock type.
CMS - Contention Count with QLen>1 CMSEQDQ - Contention Count with QLen>1 CMSLatch - Contention Count with QLen>1 CMSSMF - Contention Count with QLen>1	The number of times that a unit of work of the indicated address space was suspended on the respective lock type when there was already at least one other unit of work suspended for the lock.

CML and Local Lock Details

Address Space ID	Job Name	Service Class Name	Service Class Period	CML Lock Owner Total Contention Time	CML Lock Owner Avg Contention Time	CML Lock Owner Total Contention Count	CML Lock Owner Contention Count with QLen>1	Local Lock Total Contention Time	Local Lock Avg Contention Time	Local Lock Total Contention Count	Local Lock Contention Count with QLen>1	CML Lock Requestor Total Contention Time	CML Lock Requestor Avg Contention Time	CML Lock Requestor Total Contention Count
0007	GRS	SYSTEM	1	1281	0.92	1392	536	4857691	0.16	30083088	26063879			
0006	XCFAS	SYSTEM	1	2065	2.96	696	236	5069	2.75	1840	449			
0016	IOSAS	SYSTEM	1					510	0.15	3226	871	14	0.63	22
0018	IXGLOGR	SYSTEM	1					251	0.08	3082	85	359	7.97	45
003E	RMF	SYSSTC	1					208	0.02	9096	1969	14	0.60	23
0055	BPXAS	SYSSTC	1					78	0.06	1256	254			
000A	SMSVSAM	SYSTEM	1	0	0.00	3	0	72	0.02	2846	210	350	8.33	42
000C	WLM	SYSTEM	1	1	0.16	6	4	71	0.16	425	42	32	0.84	38
000B	CONSOLE	SYSTEM	1	84	0.10	825	58	53	0.05	903	100			
001F	APPC	STCLOW	1					43	0.03	1158	94			
00AB	SOAKER	STCLOW	1					29	0.00	3858	0			
0055	BPXAS	OMVLOW	2					24	0.08	288	49			
00D8	SOAKER	STCLOW	1					23	0.00	3826	0			
002A	NETVIEW	SYSSTC	1					20	0.04	466	20	54	0.12	424
0041	SOAKER	STCLOW	1					20	0.00	4021	0			

Figure 235. SDELAY Report - Serialization Delay Details - CML and Local Lock Details

Table 207. Fields in the Serialization Delay Details section - CML and Local Lock Details

Field Heading	Meaning
CML and Local Lock Details – contains detail data about CML and LOCAL locks per address space.	
Address Space ID	The hexadecimal address space identifier (ASID) of the job for which lock data was collected.
Jobname	The name of the job.
Service Class Name	The name of the service class that the job has been running in.
Service Class Period	The service class period that the job has been running in.
CML Lock Owner - Total Contention Time	The total amount of time in milliseconds that a unit of work from another address space was suspended when requesting the LOCAL lock of the indicated address space.
CML Lock Owner - Avg Contention Time	The average amount of time in milliseconds that a unit of work from another address space was suspended when requesting the LOCAL lock of the indicated address space.
CML Lock Owner - Total Contention Count	The number of times that a unit of work from another address space was suspended when requesting the LOCAL lock of the indicated address space.
CML Lock Owner - Contention Count with QLen>1	The number of times that a unit of work from another address space was suspended when requesting the LOCAL lock of the indicated address space and there was already at least one other unit of work waiting for this lock.

Table 207. Fields in the Serialization Delay Details section - CML and Local Lock Details (continued)	
Field Heading	Meaning
Local Lock - Total Contention Time	The total amount of time in milliseconds that a unit of work of the indicated address space was suspended on a LOCAL lock.
Local Lock - Avg Contention Time	The average amount of time in milliseconds that a unit of work of the indicated address space was suspended on a LOCAL lock.
Local Lock - Total Contention Count	The number of times that a unit of work of the indicated address space was suspended on a LOCAL lock.
Local Lock - Contention Count with QLen>1	The number of times that a unit of work of the indicated address space was suspended on a LOCAL lock when there was already at least one other unit of work suspended.
CML Lock Requestor - Total Contention Time	The total amount of time in milliseconds that a unit of work of the indicated address space was suspended when requesting the LOCAL lock of another address space.
CML Lock Requestor - Contention Time	The average amount of time in milliseconds that a unit of work of the indicated address space was suspended when requesting the LOCAL lock of another address space.
CML Lock Requestor - Total Contention Count	The number of times that a unit of work from this address space was suspended when requesting the LOCAL lock of another address space.
CML Lock Requestor - Contention Count with QLen>1	The number of times that a unit of work from this address space was suspended when requesting the LOCAL lock of another address space and there was already at least one other unit of work waiting for that lock.




Table 208. Fields in the Serialization Delay Details section - GRS Latch Details	
Field Heading	Meaning
GRS Latch Details – contains detail data about GRS latches.	
Address Space ID	The hexadecimal address space identifier (ASID) of the job for which lock data was collected.
Jobname	The name of the job.
Service Class Name	The name of the service class that the job has been running in.
Service Class Period	The service class period that the job has been running in.
In the following field descriptions, the term Latch Set Creator denotes statistics for latch obtain requests against latch sets created by this address space and Latch Requestor denotes statistics for latch obtain requests issued from this address space :	

Table 208. Fields in the Serialization Delay Details section - GRS Latch Details (continued)	
Field Heading	Meaning
Latch Set Creator - Total Contention Time Latch Requestor - Total Contention Time	The amount of contention time in milliseconds that was caused by latch obtain requests.
Latch Set Creator - Avg Contention Time Latch Requestor - Avg Contention Time	The average amount of contention time in milliseconds that was caused by latch obtain requests.
Latch Set Creator - Std Dev of Contention Time Latch Requestor - Std Dev of Contention Time	The standard deviation of the total contention time.
Latch Set Creator - Total Contention Count Latch Requestor - Total Contention Count	The number of times a latch obtain request was suspended.




Figure 237. SDELAY Report - Serialization Delay Details - GRS Enqueue Details

Table 209. Fields in the Serialization Delay Details section - GRS Enqueue Details	
Field Heading	Meaning
GRS Enqueue Details – contains detail data about GRS enqueue requests.	
Address Space ID	The hexadecimal address space identifier (ASID) of the job for which lock data was collected.
Jobname	The name of the job.
Service Class Name	The name of the service class that the job has been running in.
Service Class Period	The service class period that the job has been running in.
ENQ STEP - Total Contention Time ENQ SYSTEM - Total Contention Time ENQ SYSTEMS - Total Contention Time	The total amount of contention time in milliseconds that was caused by GRS ENQ requests of the indicated scope for this address space.
ENQ STEP - Avg Contention Time ENQ SYSTEM - Avg Contention Time ENQ SYSTEMS - Avg Contention Time	The average amount of contention time in milliseconds that was caused by GRS ENQ requests of the indicated scope for this address space.
ENQ STEP - Std Dev of Contention Time ENQ SYSTEM - Std Dev of Contention Time ENQ SYSTEMS - Std Dev of Contention Time	The standard deviation of the total contention time in milliseconds for GRS ENQ requests of the indicated scope for this address space.

Table 209. Fields in the Serialization Delay Details section - GRS Enqueue Details (continued)	
Field Heading	Meaning
ENQ STEP - Request Count ENQ SYSTEM - Request Count ENQ SYSTEMS - Request Count	The total number of GRS ENQ requests of the indicated scope for this address space.
ENQ STEP - Contention Count ENQ SYSTEM - Contention Count ENQ SYSTEMS - Contention Count	The total number of GRS ENQ requests of the indicated scope that were suspended for this address space.

SDEVICE - Shared Device Activity report

This section describes the Shared Device report. There are two types:

- The Shared Direct Access Device Activity Report
- The Shared Magnetic Tape Device Report

The report gives you an overall performance picture of DASD and TAPE devices that are shared between z/OS systems in a sysplex.

For each shared DASD or tape device the report contains one line for each system that has access to it. The additional system line shows the device activity contributed by all systems in the sysplex.

How to request this report

Monitor I gathers data for the DASD Activity report automatically with the default option `DEVICE(DASD)`. If you want to suppress gathering, you need to specify `DEVICE(NODASD)`.

To gather data for the TAPE Activity report, specify `DEVICE(TAPE)`.

To produce this report, specify

```
SYSRPTS(SDEVICE(options))
```

This report is also available in XML output format. [How to work with Postprocessor XML reports in z/OS Resource Measurement Facility User's Guide](#) provides all required information on how to produce and view XML reports.

Note: The report requires matching device numbers (the physical device must have the same device number on all systems), or self-defining devices to give meaningful results.

Example URL for the DDS API

```
http://ddshost:8803/gpm/rmfpp.xml?reports=SDEVICE(NMBR(2000:2200))
```

Using the information given in the report

The summary line allows you to identify a bottleneck caused by device delay in the sysplex. Furthermore, it allows you to see each systems share in the bottleneck.

The summary device activity rate and the device utilization show the total load on the device. The single-system values show the share of each system.

Shared Direct Access Device Activity report

The following example reports about a sysplex consisting of two systems (SYSD and SYSE). Only two devices are shown.

Both devices have the same device number on both systems.

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																	PAGE	1	
z/OS 3.1			SYSPLEX SYSDPLEX RPT VERSION 3.1 RMF			DATE 09/30/2021 TIME 12.30.00			INTERVAL 15.00.000 CYCLE 1.000 SECONDS										
TOTAL SAMPLES(AVG) = 900.0 (MAX) = 900.0 (MIN) = 900.0																			
DEV NUM	DEVICE TYPE	VOLUME SERIAL	PAV	SMF SYS ID	IOF SUFF	LCU	DEVICE ACTIVITY RATE	AVG RESP TIME	AVG IOSQ TIME	AVG CMR DLY	AVG DB DLY	AVG INT DLY	AVG PEND TIME	AVG DISC TIME	AVG CONN TIME	% DEV CONN	% DEV UTIL	% DEV RESV	AVG NUMBER ALLOC
02180	33909	SYST10		*ALL			0.036	0.248	0.000	0.000	0.000	0.005	0.120	0.000	0.128	0.00	0.00	0.0	0.0
			1	SYSD	00	0062	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.00	0.00	0.0	0.0
			1	SYSE	00	0062	0.036	0.248	0.000	0.000	0.000	0.000	0.005	0.120	0.000	0.128	0.00	0.00	0.0
02181	33909	SYST11		*ALL			0.071	0.264	0.000	0.006	0.000	0.000	0.130	0.000	0.134	0.00	0.00	0.0	0.0
			1	SYSD	00	0062	0.036	0.280	0.000	0.012	0.000	0.000	0.136	0.000	0.144	0.00	0.00	0.0	0.0
			1	SYSE	00	0062	0.036	0.248	0.000	0.000	0.000	0.000	0.124	0.000	0.124	0.00	0.00	0.0	0.0

Figure 238. Shared DASD Activity Report

For the field descriptions, see [Table 210 on page 462](#).

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 3.1			SYSPLEX SYSDPLEX RPT VERSION 3.1 RMF			DATE 09/30/2021 TIME 00.15.00			INTERVAL 30.00.001 CYCLE 1.000 SECONDS											
TOTAL SAMPLES(AVG) = 1800.0 (MAX) = 1800.0 (MIN) = 1800.0																				
DEV NUM	DEVICE TYPE	VOLUME SERIAL	SMF SYS ID	IOF SUFF	LCU	DEVICE ACTIVITY RATE	AVG RESP TIME	AVG IOSQ TIME	AVG CMR DLY	AVG DB DLY	AVG PEND TIME	AVG DISC TIME	AVG CONN TIME	% DEV CONN	% DEV UTIL	% DEV RESV	NUMBER OF MOUNTS	AVG MOUNT TIME	TIME DEVICE ALLOC	
00500	3490	TAP508	*ALL			9.511	42.97	1.134	11.71	2.442	27.51	7.312	7.014	6.64	13.6	0.0	4	19	57:37	
		TAP508	P40	A1	0006	0.012	451.8	227.4	39.02	131.8	217.7	0.156	6.621	0.01	0.01	0.0	0	0	13:37	
		TAP508	P50	A1	0006	9.474	40.92	0.157	11.52	1.034	26.35	7.379	7.023	6.61	13.6	0.0	4	19	12:01	
		TAP508	P60	A1	0006	0.012	152.9	45.24	33.03	29.94	101.5	0.342	5.845	0.01	0.01	0.0	0	0	15:39	
		TAP508	P70	A1	0006	0.012	1130	409.3	78.81	560.8	714.7	0.145	6.014	0.01	0.01	0.0	0	0	16:18	
00505	3490	TAP606	*ALL			9.829	54.13	3.174	14.18	4.034	34.53	9.913	6.514	6.39	16.2	0.0	8	28	1:04:04	
		TAP606	P40	A1	0006	0.012	1347	636.4	31.82	640.8	702.6	0.158	8.036	0.01	0.01	0.0	0	0	16:27	
		TAP606	P50	A1	0006	1.109	37.01	3.235	8.147	7.942	25.94	1.969	5.868	0.65	0.86	0.0	1	1:05	17:17	
		TAP606	P60	A1	0006	8.696	50.71	0.000	14.72	1.248	33.01	11.02	6.689	5.73	15.3	0.0	*	7	22	16:11
		TAP606	P70	A1	0006	0.012	2145	955.8	60.86	1061	1182	0.147	7.625	0.01	0.01	0.0	0	0	14:07	

Figure 239. Shared Magnetic Tape Device Activity Report

Field descriptions

The table [Table 210 on page 462](#) contains all report fields and the description how the values shown are calculated.

Table 210. 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.

Table 210. Fields in the Shared Device Activity Reports (continued)

Field Heading	Meaning
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	A five-digit hexadecimal device number that identifies a physical I/O device. The first digit represents the ID of the subchannel set to which the I/O device is physically configured. 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 changed during the reporting interval, it is followed by an '*'. If the device is a HyperPAV base device, the number is followed by an 'H', for example, 5.4H. The value is the average number of HyperPAV volumes (base and alias) in that interval. $\text{Average \# of HPAV devices} = \frac{\text{Accumulated \# of HPAV devices}}{\text{Number of Samples}}$
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.
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: $\text{DEVICE ACTIVITY RATE} = \frac{\text{SSCH Count}}{\text{CIL}}$ 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.

Table 210. Fields in the Shared Device Activity Reports (continued)

Field Heading	Meaning
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: $\begin{aligned} \text{AVG RESP TIME} = & \text{Avg IOSQT} \\ & + \text{AVG PEND TIME} \\ & + \text{AVG DISC TIME} \\ & + \text{AVG CONN TIME} \end{aligned}$ In the summary line, the same formula is used using the corresponding summary line fields.
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: $\text{AVG IOSQ TIME} = \frac{\text{Total IOS Queue Time}}{\text{Start Subchannel Count}}$ 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: $W_i = \frac{\text{MEC}(\text{SYS}_i)}{\text{MEC}(\text{SYS}_1) + \text{MEC}(\text{SYS}_2) + \dots + \text{MEC}(\text{SYS}_n)}$ The weighted AVG IOS queue time is therefore calculated as: $\begin{aligned} \text{WEIGHTED AVG IOS QT} = & (\text{AVG IOS QT}(\text{SYS}_1)) * W_1 + \\ & (\text{AVG IOS QT}(\text{SYS}_2)) * W_2 + \\ & \vdots \\ & (\text{AVG IOS QT}(\text{SYS}_n)) * W_n \end{aligned}$ 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. $\text{AVG CMR DLY} = \frac{\text{Initial Command Response Time}}{\text{MEC}}$
AVG DB DLY	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: $\text{AVG DB DLY} = \frac{\text{Total DB DLY Time}}{\text{MEC}}$ 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.
AVG INT DLY	The average interrupt delay time in units of milliseconds encountered for I/O requests to this device. For each I/O request, the time is measured from when the I/O operation is complete to when the operating system begins to process the status. $\text{AVG INT DLY} = \frac{\text{Device Interrupt Delay Time}}{\text{MEC}}$

Table 210. 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: $\text{AVG PEND TIME} = \frac{\text{Total Pending Time}}{\text{MEC}}$ 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: $\text{AVG DISC TIME} = \frac{\text{Total Disconnect Time}}{\text{MEC}}$ 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: $\text{AVG CONN TIME} = \frac{\text{Total Connect Time}}{\text{MEC}}$ 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: $\% \text{ DEV CONN} = 100 * \frac{\text{Device Connect Time}}{\text{CIL}}$ In the summary line, this field contains the sum of the single system % DEV CONN values.
% 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: $\% \text{ DEV UTIL} = 100 * \frac{\text{Device Connect Time} + \text{Disconnect Time}}{\text{CIL}} + 100 * \frac{\text{UTL}}{\text{STS}}$ In the summary line, this field contains the sum of the single system % DEV UTIL values.

Table 210. Fields in the Shared Device Activity Reports (continued)

Field Heading	Meaning
% 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: $\% \text{ DEV RESV} = 100 * \frac{\text{DRP} * \text{SIL}}{\text{STS} * \text{CIL}}$ 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: $\text{AVG NUMBER ALLOC} = \frac{\text{Num Alloc} * \text{System Interval Length}}{\text{STS} * \text{Common Interval Length}}$ In the summary line, this field contains the sum of the single system AVG NUMBER ALLOC values.
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 increments 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: $\text{AVG MOUNT TIME} = \frac{\text{MTP} * \text{SIL}}{\text{STS} * \text{MTC}}$ 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: $\text{AVG MOUNT TIME} = \frac{\text{AVG Mount Time}(\text{SYS1}) * \text{MTC}(\text{SYS1}) + \dots}{\sum \text{All System MTC Counts}}$ This field is reported only for magnetic tape devices.

Table 210. Fields in the Shared Device Activity Reports (continued)	
Field Heading	Meaning
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. $\text{TIME DEVICE ALLOC} = \frac{\text{ALC} * \text{SIL}}{\text{STS}}$ 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 *z/OS Resource Measurement Facility User's Guide*.

Table 211. Overview names in the Shared DASD Activity Report	
Field Heading or Meaning	Overview Name
Percent reserved	DR
Percent mount pending	DMTPEND
Percent device utilization	DVUTL
Device activity rate	DART
Average connect time	DCTAVG
Average disconnect time	DDTAVG
Average pending time	DPTAVG
Average IOS queue time	DQTAVG
Average response time	DRTAVG
Average device busy delay time	DBDL
Average initial command response delay time	CMRDL
Average interrupt delay time	INTDL

TRACE - Trace Activity report

The Trace Activity report provides information from various trace variables.

The report shows snapshots of each of the specified variables along with timing information. RMF trace treats values collected as unsigned binary integers. See *z/OS Data Gatherer User's Guide* for a description of how to specify the TRACE options. A description of the trace variables is in Table 213 on page 469.

Note: Monitor I gathers and reports all trace variables the way they are provided by the system. Monitor I cannot influence the format, and does not perform any calculation.

How to request this report

To gather data for this report, specify as a Monitor I gatherer option

```
TRACE(variable [,options list])
```

To produce this report, specify

```
REPORTS(TRACE)
```

Note: The TRACE report is only available as an interval report, not as a duration report.

Using the information given in the report

You can use the Trace report to monitor the SRM multiprogramming level (MPL) adjustment or monitor the contention detected and handled by the system. You can see how the system handles contention by tracing the following variables: RCVUICA, RCVCPUA and RCVPTR.

Contents of the report

RMF reports all trace variables that contain invalid data on a separate report page.

The number of lines in the report is based on the cycle and interval values specified when the session is started. For example, if you specify CYCLE(250) and INTERVAL(60M), RMF takes approximately 14,400 samples. Based on a constant 60 samples/set and 1 line/set, the report contains 240 lines of data for each field.

The number of samples per set is determined by a constant located in the first halfword of CSECT ERBMFTTB in load module ERBMFMFC (in SYS1.LINKLIB). Any value in the range 1 to 32,767 is valid and can be changed by the system programmer. If changed to zero, RMF overrides it with 1. If changed to a value less than 0 (a negative number), RMF defaults to 32,767 (X'7FFF'). If the values specified for cycle, interval, and samples per set would result in the number of sets exceeding this limit, RMF suspends trace sampling for the remainder of the interval.

TRACE ACTIVITY														PAGE 1	
z/OS 3.1		SYSTEM ID SYS1		DATE 09/30/2021		INTERVAL 14.59.946									
		RPT VERSION 3.1 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	* RCVPAVRT	* RCVPTR	* SMCABFLS	* MCVFRONT	* MCVSTCR1	* 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 240. TRACE Report

Table 212. 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.

Table 212. Fields in the Trace Activity Report (continued)	
Field Heading	Meaning
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 213. Variables in the Trace Activity Report	
Variable	Value
ASMERRS	Bad slots on local page data sets
ASMIORQC	Count of I/O requests completed and returned to RSM
ASMIORQR	Count of I/O requests received by I/O control
ASMNVSC	Total local slots allocated for non-VIO private area pages
ASMSLOTS	Total local slots (sum of slots in open local page data sets)
ASMVSC	Total local slots allocated for VIO private area pages
CCVCPUCT	Number of online logical processors (threads)
CCVENQCT	Number of users non-swappable for enqueue reasons
CCVRBSTD	Recent base time of day
CCVRBSWT	Recent base system wait time
CCVUTILP	System CPU utilization
LSCTCNT	Current number of logically swapped users for terminal wait
LSCTMTE	Maximum think time allowed for logical swap candidate
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

Table 213. Variables in the Trace Activity Report (continued)	
Variable	Value
RCEAEC	Total number of expanded storage E frames currently on the ESTE queue
RCEAECLO	Available expanded storage low threshold
RCEAECOK	Available expanded storage satisfactory threshold
RCEAFC	Total number of frames currently on all available frame queues
RCEAFCLO	Available central storage low threshold
RCEAFCOK	Available central storage satisfactory threshold
RCEBELFX	Total number of fixed pages below the 16M line in central storage, which is the sum of page-fixed LSQA, SQA (excluding reserved SQA), and V=R allocated pages
RCECOMPI	Number of common area pages paged-in
RCECOMPO	Number of common area pages paged-out
RCEDFRS	Number of times a deferred frame allocation has been satisfied
RCEESINU	Number of in-use expanded storage frames
RCEESREA	Number of non-VIO pages read from expanded storage
RCEESWRT	Number of pages written to expanded storage frames
RCEHSPDM	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
RCENVBEL	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 213. 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
RMCARSSC	Central storage shortage swap count
RMCATISC	Terminal input swap count
RMCATOSC	Terminal output swap count
RMCATSSC	Count of transition swaps
RMCAUSSC	Unilateral swap out count
RMCASSC	Auxiliary storage shortage swap count
RMCTTRPC	Number of pages used for transaction elements
SMCABFLS	Number of records lost because of a shortage of buffers
SMCABFWT	Number of buffers written
SMCACNBF	Current number of buffers
SMCADSCT	Number of records lost because of a full data set
SMCANMFL	Current number of full buffers
SMCARCWT	Number of records written

Standard deviation output

The Trace Activity report gives an account of standard deviation as an exponential (E-format) number. This number expresses the standard deviation for a set (sub-interval) as a number between 0.000 and 9.999, raised to a power of ten.

Example

1.123E +01 is the same as 1.123 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.

VSTOR - Virtual Storage Activity report

The Virtual Storage Activity report provides information about the use of virtual storage that can help your installation manage its use of virtual storage.

The report is formatted into the following five sections, each with a separate subheading:

- COMMON STORAGE SUMMARY (see [“Common Storage Summary section” on page 473](#))
- COMMON STORAGE DETAIL (see [“Common Storage Detail section” on page 475](#))
- PRIVATE AREA SUMMARY (see [“Private Area Summary section” on page 476](#))
- PRIVATE AREA DETAIL (see [“Private Area Detail section” on page 479](#))
- HIGH VIRTUAL MEMORY USAGE (see [“High Virtual Memory Usage section” on page 479](#))

Most of the information you need to begin managing virtual storage appears in the common storage summary report. When this report indicates a problem, you can request one of the three more comprehensive reports for additional information. The structure and fields of the different reports are described in [“Contents of the report” on page 473](#).

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 with the GETMAIN macro instruction. **Allocated storage** is any block of at least 4K that contains any storage obtained with the GETMAIN macro instruction. Thus, for the purposes of the report, free storage within a 4K block assigned to a subpool is allocated storage. Both free storage and allocated storage are reported as a multiple of 4K on the reports.

Using the information given in the report

Information on virtual storage use is particularly helpful in the process of long-term measurements. It helps you, for example, understand your current use of virtual storage, see the relationship between increased use of your system and increased demands on virtual storage, and predict future constraints before they occur. This ability to predict a future constraint is useful for the virtual storage resource because actions that can relieve a virtual storage constraint generally require significant time to plan and implement. The report can also help you determine the effect of any actions, such as moving a large application above the 16-megabyte line or installing products that take advantage of expanded addressing.

The information in the report can help you identify any expansion of SQA into CSA and set appropriate size values for CSA, RUCSA, and SQA at IPL time. You can use the report to verify the cost (in increased PLPA inter-module space) of any pack lists your installation uses to reduce PLPA paging.

How to request this report

Monitor I gathers data for this report automatically with the default option VSTOR(S). See *z/OS Data Gatherer User's Guide* for details. If you want to suppress gathering, specify NOVSTOR.

To produce this report, specify

```
REPORTS(VSTOR(S))  
REPORTS(VSTOR(D))  
REPORTS(VSTOR(D[,jobname1,jobname2,...]))  
REPORTS(VSTOR(jobname1[,jobname2,...]))
```

This report is also available in XML output format. [How to work with Postprocessor XML reports in z/OS Resource Measurement Facility User's Guide](#) provides all required information on how to produce and view XML reports.

Example URLs for the DDS API

```
http://ddshost:8803/gpm/rmfpp.xml?reports=VSTOR
http://ddshost:8803/gpm/rmfpp.xml?reports=VSTOR(D)
http://ddshost:8803/gpm/rmfpp.xml?reports=VSTOR(D,rmf)
```

Contents of the report

All size data values are reported in bytes. The size is followed by a K (indicating the number of kilobytes the value represents) unless the size is greater than 9999K. When the size is greater than 9999K, the size is followed by an M (indicating the number of megabytes the value represents). Because peak values are especially important when analyzing virtual storage use, the minimum, maximum, and average values are reported whenever useful, and the minimum and maximum values are time-stamped.

Data gathering considerations

To minimize overhead, RMF does not sample virtual storage data at every cycle. RMF takes one sample of virtual storage data for every ten RMF cycles. For example, if the RMF cycle is one second, RMF samples virtual storage data every ten seconds. In this case, RMF provides time stamps (accurate to within a ten-second range) for each minimum and maximum value on the report. The time stamp shows the time when RMF first observed the minimum or maximum value in the sample.

Common Storage Summary section

The common storage summary section enables you to measure the use of virtual storage with minimal overhead. It contains the information you need to understand your current use of virtual storage. If you archive the data, you can use differences over time to predict a problem or constraint before it becomes critical. It also helps you to verify the size values set for CSA and SQA at IPL time and determine if you are using common storage effectively. Because RMF does not sample virtual storage data at every cycle, the value reported for NUMBER OF SAMPLES is less than the number of cycles.

VIRTUAL STORAGE ACTIVITY										PAGE 1	
z/OS 3.1			SYSTEM ID TRX1		DATE 09/30/2021		INTERVAL 03.31.164				
			RPT VERSION 3.1 RMF		TIME 16.56.01		CYCLE 1.000 SECONDS				
COMMON STORAGE SUMMARY											
NUMBER OF SAMPLES 24											
STATIC STORAGE MAP			ALLOCATED CSA/SQA								
AREA	ADDRESS	SIZE	----- BELOW 16M -----		----- EXTENDED (ABOVE 16M) -----						
EPVT	22B00000	1493M	MIN	MAX	AVG	MIN	MAX	AVG			
ECSA	E7C3000	323M	SQA 396K 16.56.02	1196K 16.56.41	1062K	33.5M 16.56.02	97.5M 16.56.52	86.8M			
EMLPA	0	0K	CSA 392K 16.56.02	892K 16.56.41	808K	45.6M 16.56.02	49.7M 16.57.21	49.0M			
EFLPA	E7C0000	12K									
EPLPA	A96F000	62.3M	ALLOCATED CSA BY KEY								
ESQA	1DB6000	140M	0	144K 16.56.02	144K	32.7M 16.56.02	32.7M 16.56.02	32.7M			
ENUC	1000000	13.7M	1	128K 16.56.02	128K	1896K 16.56.02	1896K 16.56.02	1896K			
----- 16 MEG BOUNDARY -----			2	40K 16.56.02	40K	44K 16.56.02	44K 16.56.02	44K			
NUCLEUS	FD3000	180K	3	0K 16.56.02	0K	0K 16.56.02	0K	0K			
SQA	E3E000	1620K	4	0K 16.56.02	0K	0K 676K 16.56.02	676K 16.56.02	676K			
PLPA	C5B000	1932K	5	4K 16.56.02	4K	2840K 16.56.02	2840K 16.56.02	2840K			
FLPA	C5A000	4K	6	76K 16.56.02	76K	7760K 16.57.41	7788K 16.57.21	7768K			
MLPA	0	0K	7	0K 16.56.02	0K	0K 16.56.02	0K	0K			
CSA	800000	4456K	8-F	0K 16.56.02	500K 16.56.41	416K	0K 16.56.02	4096K 16.56.41	3413K		
PRIVATE	2000	8184K									
PSA	0	8K	SQA EXPANSION INTO CSA								
			0K 16.56.02	0K	0K	0K 16.56.02	0K	0K			
PLPA INTERMODULE SPACE - 2K IN PLPA AND 307K IN EPLPA											
PLPA SPACE REDUNDANT WITH MLPA/FLPA - 0K IN PLPA AND 11K IN EPLPA											

		MIN	MAX	AVG	MIN	MAX	AVG				
CSA											
FREE PAGES (BYTES)		3564K 16.56.41	4064K 16.56.02	3647K	274M 16.57.21	278M 16.56.02	274M				
LARGEST FREE BLOCK		3040K 16.56.02	3040K 16.56.02	3040K	255M 16.56.02	255M 16.56.02	255M				
ALLOCATED AREA SIZE		392K 16.56.02	892K 16.56.41	808K	45.7M 16.56.02	49.7M 16.56.41	49.0M				
SQA											
FREE PAGES (BYTES)		424K 16.56.41	1224K 16.56.02	557K	42.2M 16.56.52	106M 16.56.02	52.9M				
LARGEST FREE BLOCK		300K 16.56.41	896K 16.56.02	399K	38.2M 16.56.41	102M 16.56.02	48.9M				
ALLOCATED AREA SIZE		724K 16.56.02	1524K 16.56.41	1390K	140M 16.56.02	140M 16.56.02	140M				
MAXIMUM POSSIBLE USER REGION - 8184K BELOW AND 1484M ABOVE											
DEFINED SIZE OF RUCSA - 1M BELOW AND 23M ABOVE											

VIRTUAL STORAGE ACTIVITY										PAGE	2
z/OS 3.1		SYSTEM ID TRX1		DATE 09/30/2021		INTERVAL 03.31.164					
		RPT VERSION 3.1 RMF		TIME 16.56.01		CYCLE 1.000 SECONDS					
COMMON STORAGE DETAIL											
ALLOCATED CSA BY SUBPOOL		BY KEY (BELOW 16 MEG)		ALLOCATED SQA BY SUBPOOL (BELOW 16M)							
SUBPOOL 227		SUBPOOL 228		SUBPOOL 231		SUBPOOL 241					
-----		MINIMUM		-----							
0	28K 16.56.02	4K 16.56.02	4K 16.56.02	108K 16.56.02	226	24K 16.56.02	24K 16.56.02	24K			
1		4K 16.56.02		124K 16.56.02	239	44K 16.56.02	44K 16.56.02	44K			
2				40K 16.56.02	245	328K 16.56.02	1128K 16.56.41	994K			
3											
4											
5				4K 16.56.02							
6	12K 16.56.02	56K 16.56.02	4K 16.56.02	4K 16.56.02							
7											
8-F		0K 16.56.02	0K 16.56.02								
ALL	40K 16.56.02	64K 16.56.02	8K 16.56.02	280K 16.56.02							
-----		MAXIMUM		-----							
0	28K 16.56.02	4K 16.56.02	4K 16.56.02	108K 16.56.02							
1		4K 16.56.02		124K 16.56.02							
2				40K 16.56.02							
3											
4											
5				4K 16.56.02							
6	12K 16.56.02	56K 16.56.02	4K 16.56.02	4K 16.56.02							
7											
8-F		200K 16.56.41	300K 16.56.41								
ALL	40K 16.56.02	264K 16.56.41	308K 16.56.41	280K 16.56.02							
-----		AVERAGE		-----							
0	28K	4K	4K	108K							
1		4K		124K							
2				40K							
3											
4											
5				4K							
6	12K	56K	4K	4K							
7											
8-F		166K	250K								
ALL	40K	230K	258K	280K							

Figure 241. VSTOR report - Common Storage Summary

Table 214. 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. Note that restricted use common service area (RUCSA) storage is allocated with storage key 8 - F. Important: SQA-to-CSA expansion is counted under SQA/ESQA statistics.

Table 214. Fields in the Virtual Storage Activity Report - Common Storage Summary (continued)	
Field Heading	Meaning
SQA EXPANSION INTO CSA	The MIN, MAX, and AVG size of any expansion of SQA into CSA. SQA does not expand into RUCSA.
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. Note that the largest CSA free block can either be in CSA or RUCSA.
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. Important: Storage allocated for SQA-to-CSA expansion is counted under CSA/ECSA.
MAXIMUM POSSIBLE USER REGION	The largest size specified on the REGION= JCL parameter that this system can satisfy (assuming a minimal number of DD statements). RMF determines this value, reported for below and above the 16-megabyte line, by examining its own private area; it calculates the size by finding the difference between the bottom of its allocated area at RMF initialization and the start of the user region. There is no guarantee that a particular job can obtain a region of the reported size. For a job that runs under an initiator (as opposed to a job that runs as a started task), the storage that the initiator obtains, reduces the size of the region that the job can obtain. The number of DD statements in the JCL can also reduce the size of the region. The reported region size, however, can indicate whether a job with a known region requirement is likely to obtain the region it requires under the system conditions reflected in the report.
DEFINED SIZE OF RUCSA	The size that is defined in the IEASYSxx parmlib member for RUCSA and extended RUCSA.

Overview reference

Table 215. Overview names in the Virtual Storage Activity Report	
Field Heading or Meaning	Overview Name
SQA EXPANSION INTO CSA	SQAE
LARGEST FREE BLOCK - MIN, CSA / RUCSA	CSAFB
LARGEST FREE BLOCK - MIN, SQA	SQAFB
ALLOCATED AMOUNT BELOW 16M - MAX, CSA / RUCSA	CSAA
FREE BYTES BELOW 16M - MIN, CSA / RUCSA	CSAFP

Common Storage Detail section

The optional common storage detail section contains additional information about the use of CSA and SQA below the 16-megabyte line.

VIRTUAL STORAGE ACTIVITY										PAGE	2
z/OS 3.1		SYSTEM ID TRX1		DATE 09/30/2021		INTERVAL 05.00.000					
		RPT VERSION 3.1 RMF		TIME 10.35.00		CYCLE 1.000 SECONDS					
COMMON STORAGE DETAIL											
ALLOCATED CSA BY SUBPOOL		BY KEY (BELOW 16 MEG)				ALLOCATED SQA BY SUBPOOL (BELOW 16M)					
SUBPOOL 227		SUBPOOL 228		SUBPOOL 231		SUBPOOL 241		SUBPOOL		MIN	
								MAX		AVG	
-----		MINIMUM		-----							
0	28K 11.20.25	4K 11.20.25	4K 11.20.25	112K 11.20.25	124K 11.20.25	226	24K 11.20.25	24K 11.20.25	24K		
1		4K 11.20.25		124K 11.20.25	40K 11.20.25	239	36K 11.20.25	36K 11.20.25	36K		
2						245	224K 11.20.25	224K 11.20.25	224K		
3											
4											
5				4K 11.20.25							
6	12K 11.20.25	52K 11.20.25	12K 11.20.25	8K 11.20.25							
7											
8-F											
ALL	40K 11.20.25	60K 11.20.25	16K 11.20.25	288K 11.20.25							
-----		MAXIMUM		-----							
0	28K 11.20.25	4K 11.20.25	4K 11.20.25	112K 11.20.25	124K 11.20.25						
1		4K 11.20.25		124K 11.20.25	40K 11.20.25						
2											
3											
4											
5				4K 11.20.25							
6	12K 11.20.25	52K 11.20.25	12K 11.20.25	8K 11.20.25							
7											
8-F											
ALL	40K 11.20.25	60K 11.20.25	16K 11.20.25	288K 11.20.25							
-----		AVERAGE		-----							
0	28K	4K	4K	112K	124K						
1		4K		124K	40K						
2											
3											
4											
5				4K							
6	12K	52K	12K	8K							
7											
8-F											
ALL	40K	60K	16K	288K							

Figure 242. VSTOR report - Common Storage Detail

Table 216. Fields in the Virtual Storage Activity Report - Common Storage Detail Section	
Field Heading	Meaning
ALLOCATED CSA BY SUBPOOL BY KEY	The matrix that presents MINIMUM, MAXIMUM, and AVERAGE use of CSA by subpools 227, 228, 231, and 241 broken down by storage key and summed for ALL keys.
ALLOCATED SQA BY SUBPOOL	The MIN, MAX, and AVG values for subpools 226, 239, and 245.

Private Area Summary section

The optional private area summary section presents information about how a specific address space is using its private virtual storage. RMF uses the job name you specify when you request the report to identify the address space. In choosing a job, note that gathering data for a specific address space requires additional RMF overhead. Note that RMF can gather private area data only when a job is active at the beginning of the interval, and various conditions can limit RMF's ability to report complete private area data. These conditions, and the actions RMF takes, are described later in this section under [“Partial private area data” on page 478](#).

In general, RMF can gather meaningful data only for jobs that run for a relatively long period of time. Note that you cannot monitor the master scheduler address space.

Note: Measuring virtual storage activity for a specific job may have significant impact on the performance of the job. When requesting a VSTOR report, system address spaces, such as CATALOG, VTAM, Db2, IMS or other, should be specified as *jobname* only for a short period of time when diagnosing a special performance situation.




Figure 243. VSTOR report - Private Area Summary

Table 217. Fields in the Virtual Storage Activity Report - Private Area Summary	
Field Heading	Meaning
JOBNAME and PROGRAM NAME	The job analyzed in the report. RMF takes the program name from the PGM= parameter on the exec statement.
NUMBER OF SAMPLES	The number of samples RMF used to generate the data in the report. If the job was swapped out at a time when RMF tried to sample virtual storage data, this number will be less than the number of samples reported for the common storage summary report.
REGION REQUESTED	The values specified for the REGION= parameter on the JOB or EXEC JCL statement for the job step or the system default used for the job step.
REGION ASSIGNED	The region assigned to the job by installation control for virtual storage below and above the 16-megabyte line. This value limits the amount of storage that a job can obtain by issuing a variable length GETMAIN. (When a job issues a variable length GETMAIN, the amount of storage obtained is the difference between REGION ASSIGNED and the top of the allocated area, assuming that the largest free block is there.)

Private Storage Map

The information reported under Private Storage Map defines significant boundaries within the private area and shows the space between them. There is a separate map for storage below the 16-megabyte line and for extended storage (above the 16-megabyte line). Each map identifies:

Table 218. 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 with 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.

Table 218. Fields in the Virtual Storage Activity Report - Private Storage Map (continued)	
Field Heading	Meaning
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 219. Fields in the Virtual Storage Activity Report - Bottom Half	
Field Heading	Meaning
FREE PAGES (BYTES)	The MIN, MAX, and AVG values, in bytes, for the amount of free storage.
LARGEST FREE BLOCK	The MIN, MAX, and AVG values, in bytes, for the size of the largest free block in the available amount of free storage.
PAGES ALLOCATED (IN BYTES)	The MIN, MAX, and AVG values for the amount of allocated storage. 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 jobs that are running at least one interval. You can, of course, monitor other jobs, but there are some conditions that mean RMF can collect little or no data. These conditions, which are related to the way virtual storage reporting works, are:

1. RMF searches for any requested jobs at the beginning of each interval. If it does not find a job, it does not monitor the job during the interval. In this case, RMF issues a message to the operator and produces a report. The report, however, contains no data; instead, the following message appears:

```
JOB WAS NOT ACTIVE AT THE BEGINNING OF THIS INTERVAL
```

RMF continues to search for the job at the beginning of each interval. When it finds the job, it deletes the message, monitors the job, and produces a report.

If a job begins and ends within a single RMF interval, RMF cannot monitor its use of virtual storage.

2. If a job that RMF is monitoring terminates and is then restarted, the report for the interval in which it terminated includes data only up to the point when the job terminated. RMF resumes its monitoring of the restarted job at the beginning of the interval following the interval during which the job was restarted.

- If a job RMF is monitoring is swapped out at the time RMF takes a sample of virtual storage data, RMF does not cause a swap-in; it skips the sample for that job. Thus, the number of samples for a swappable job may be less than expected. If a job is swapped out every time RMF tries to take a sample during an interval, RMF reports no data for that interval.

Private Area Detail section

The optional **Private Area Detail** section provides information about the number of bytes of allocated blocks by area below the 16-megabyte line. In the header, the job name and the memory limit in bytes for this address space is displayed.

VIRTUAL STORAGE ACTIVITY									
z/OS 3.1		SYSTEM ID SYSF		START 09/30/2021-10.14.33		INTERVAL 000.15.00		PAGE 4	
		RPT VERSION 3.1 RMF		END 09/30/2021-10.29.33		CYCLE 1.000 SECONDS			
PRIVATE AREA DETAIL									
JOB NAME - RMF		MEMORY LIMIT - 16384P							
NUMBER OF BYTES OF		ALLOCATED BLOCKS		BY AREA (BELOW 16 MEG)					
SUBPOOL (AREA)		MIN		MAX		AVG			
230		72K 10.15.01		72K 10.15.01		72K			
236 (SWA)		84K 10.15.01		84K 10.15.01		84K			
237 (SWA)		24K 10.15.01		24K 10.15.01		24K			
255 (LSQA)		28K 10.15.01		28K 10.15.01		28K			
USER REGION									
0		32K 10.15.01		32K 10.15.01		32K			
7		4K 10.15.01		4K 10.15.01		4K			
9		4K 10.15.01		4K 10.15.01		4K			
10		4K 10.15.01		4K 10.15.01		4K			
11		4K 10.15.01		4K 10.15.01		4K			
252 (REENTRANT)		552K 10.15.01		552K 10.15.01		552K			
HIGH VIRTUAL MEMORY USAGE (ABOVE 2GB)									
BYTES		MIN		MAX		AVG		PEAK	
PRIVATE		10.00M 10.15.01		10.00M 10.15.01		10.00M		9.000M	
SHARED		0 10.15.01		0		0		0	
COMMON		22.00M 10.15.01		22.00M 10.15.01		22.00M		22.00M	
MEMORY OBJECTS									
PRIVATE		6 10.15.01		6 10.15.01		6			
SHARED		0 10.15.01		0		0			
COMMON		20 10.15.01		20 10.15.01		20			
FIXED 1 MB		2 10.15.01		2 10.15.01		2			
FIXED 2 GB		0 10.15.01		0		0			
SHARED 1 MB		0 10.15.01		0		0			
1 MB FRAMES									
FIXED		4 10.15.01		4 10.15.01		4			
PAGEABLE		0 10.15.01		0		0			
2 GB FRAMES									
FIXED		0 10.15.01		0		0			

Figure 244. VSTOR report - Private Area Detail and High Virtual Memory Usage (above 2GB)

Table 220. Fields in the Virtual Storage Activity report - Private Area Detail section	
Field Heading	Meaning
JOB NAME	The job analyzed in the report.
MEMORY LIMIT	The memory limit in bytes for this address space.
NUMBER OF BYTES OF ALLOCATED BLOCKS BY AREA (BELOW 16 MEG)	
SUBPOOL (AREA) / USER REGION	The MIN, MAX, and AVG values for the number of bytes of allocated blocks during the report interval, broken down by subpool and by area (LSQA, SWA).

High Virtual Memory Usage section

The optional **High Virtual Memory Usage** section provides information about the high virtual memory usage above the 2-gigabyte line.

See [Figure 244 on page 479](#) for an example of the High Virtual Memory Usage section.

Table 221. Fields in the Virtual Storage Activity report - High Virtual Memory Usage section	
Field Heading	Meaning
BYTES	The MIN, MAX, and AVG values for the number of bytes in PRIVATE , SHARED , and COMMON memory objects allocated with the indicated job as the owner. In addition, the PEAK useable storage since the start of the job are shown. Unlike MIN, MAX, and AVG, the PEAK values report usable storage only. Hidden storage, such as guard areas, are not included so that the reported PEAK value may be smaller than the MIN, MAX, or AVG values.

Table 221. Fields in the Virtual Storage Activity report - High Virtual Memory Usage section (continued)	
Field Heading	Meaning
MEMORY OBJECTS	The MIN, MAX, and AVG values for the number of PRIVATE , SHARED , and COMMON memory objects allocated with the indicated job. If Enhanced DAT Facility is installed, fixed memory objects and shared memory objects that are backed in 1 MB frames are reported. If Enhanced DAT Facility 2 is installed, fixed memory objects that are backed in 2 GB frames are reported.
1 MB FRAMES	The MIN, MAX, and AVG values for the number of 1 MB frames that are used by FIXED and PAGEABLE memory objects with the indicated job as the owner (only available with Enhanced DAT Facility). The PAGEABLE value also includes 1 MB frames that are used by DREF memory objects. Frames that are either used by shared 1 MB pages or to satisfy 4 KB space requests on a constrained system are not included.
2 GB FRAMES	The MIN, MAX, and AVG values for the number of 2 GB frames that are used by FIXED memory objects with the indicated job as the owner (only available with Enhanced DAT Facility 2).

WLMGL - Workload Activity report

The Workload Activity report (WLMGL) can be used to request a variety of reports, as shown in [“The WLMGL option list”](#) on page 480.

How to request this report

Monitor I gathers data for this report automatically. If you want to suppress gathering, you need to specify NOWKLD.

To produce this report, specify

```
SYSRPTS(WLMGL(options))
```

This report is also available in XML output format. How to work with Postprocessor XML reports in z/OS Resource Measurement Facility User's Guide provides all required information on how to produce and view XML reports.

Example URL for the DDS API

```
http://ddshost:8803/gpm/rmfpp.xml?reports=WLMGL(SCLASS(BATCH))
```

The WLMGL option list

You can select among various types of WLMGL reports by specifying the SYSRPTS WLMGL options:

SCPER

All service class periods found for a service class.

This report contains detailed information about:

- Subsystem delays
- Response time goals vs actuals
- General execution delays
- Response time distribution.

SCLASS

Summary of data for all service class periods defined for a service class.

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 and tenant report classes defined in a service policy.

RCPER

All periods found for a report class or tenant report class. The report has the same structure and information as the Service Class Period report.

RTD

Response Time Distribution is displayed in WLMGL Service/Report Class Period reports (RTD is default, NORTD otherwise). This suboption can only be specified together with suboptions SCPER or RCPER; otherwise, it is ignored.

In addition, you have the **SYSNAM** option to select systems to be included in the report. All systems must belong to the same sysplex.

Contents of the report

The following sections describe the various WLMGL report types as introduced in [“The WLMGL option list” on page 480](#). [“Field descriptions of all reports” on page 490](#) describes the report headers and explains all of the fields displayed by these reports.

Note that for all report types, the reporting interval is finished in case of a policy change or a policy refresh. A policy refresh is initiated either by the operator or by the Workload Manager component itself when a IEAOPTxx parameter change or a processor speed change occurs. Examples for a processor speed change are the begin or end of a boost period or when the CPC starts or ends to run in power-save or cycle-steering mode.

Service Class Period report

To request a Service Class Period report, specify:

```
SYSRPTS(WLMGL(SCPER(service_class)))
```

For example, to produce a report for all service class periods of service class STCLOW, specify:

```
SYSRPTS(WLMGL(SCPER(STCLOW)))
```

[Figure 245 on page 482](#) shows a report for service class STCLOW where the first service class period (PERIOD=1) is defined with an execution velocity goal. For service class periods with an execution velocity goal, the WLMGL report displays one tabular response time distribution for each system in the sysplex.

W O R K L O A D A C T I V I T Y												PAGE 1
z/OS 3.1		SYSplex SVplex3		DATE 03/12/2024		INTERVAL 15.00.002		MODE = GOAL				
		RPT VERSION 3.1 RMF		TIME 04.45.00								
POLICY ACTIVATION DATE/TIME 02/25/2024 09.00.02												
SERVICE CLASS PERIODS												
POLICY=BASEPOL		WORKLOAD=STC_WLD		SERVICE CLASS=STCLOW		RESOURCE GROUP=NONE		PERIOD=1 IMPORTANCE=3				
				CRITICAL								
				BATCH INITIATOR AI-INFUSED=YES								
-TRANSACTIONS--		TRANS-TIME HHH.MM.SS.FFFFFFFF		TRANS-APPL%----		CP--CBPCP--IIPCP----		CBP----		IIP		
AVG 179.62		ACTUAL 19.22.775534		TOTAL 28.64		0.00 0.02		0.00 0.99		--ENCLAVES--		
MPL 179.62		EXECUTION 19.22.306471		MOBILE 0.00		N/A 0.00		N/A 0.00		AVG ENC 0.00		
ENDED 9		QUEUED 469061		CATEGORYA 0.00		N/A 0.00		N/A 0.00		REM ENC 0.00		
END/S 0.01		R/S AFFIN 0		CATEGORYB 0.00		N/A 0.00		N/A 0.00		MS ENC 0.00		
*SWAPS 4378		INELIGIBLE 0										
EXCTD 0		CONVERSION 0										
		STD DEV 49.41.350310										

---SERVICE---		SERVICE TIME		---APPL %--		---PROMOTED---		---DASD I/O--		---STORAGE----		
IOC 236673		CPU 260.471		CP 28.76		BLK 0.000		SSCHRT 925.7		AVG 4886.40		
CPU 17355K		SRB 2.798		IIPCP 0.02		ENQ 0.000		RESP 0.5		TOTAL 877714.6		
MSD 0		RCT 0.595		IIP 0.99		CRM 0.000		CONN 0.3		SHARED 11.89		
SRB 185343		IIT 0.478		CBPCP 0.00		LCK 1.069		DISC 0.0		HSP 0.0		
TOT 17777K		HST 0.000		CBP 0.00		SUP 0.000		Q+PEND 0.1				
/SEC 19752		IIP 5.491						IOSQ 0.1				
ABSRPTN 110		CBP 0.000										
TRX SERV 110												
GOAL: EXECUTION VELOCITY 25.0% VELOCITY MIGRATION: NO I/O MGMT 80.1% INIT MGMT 80.1%												

SYSTEM		RESPONSE TIME		EX VEL		PERF INDX		AVG ADRSP		--EXEC USING%--		
										C CPU CBP IIP I/O		
										TOT I/O		
										----- EXEC DELAYS % -----		
										--USING%--		
										CRY CNT		
										UNK IDL		
										CRY CNT		
										% QUI		
*ALL		--N/A--		80.1		0.3		272.1		0.1 0.0 0.0 0.3		
CB8A				36.5		0.7		38.0		0.0 0.0 0.0 0.0		
CB8B				76.0		0.3		39.0		0.1 0.0 0.0 0.0		

-----RESPONSE TIME DISTRIBUTIONS-----												
SYSTEM: CB8A ---INTERVAL: 00.15.00.000 ---MRT CHANGES: 0 ---												
---TIME-----												
HH.MM.SS.FFFFFFFF CUM TOTAL IN BUCKET CUM TOTAL IN BUCKET												
---# TRANSACTIONS--- ---% TRANSACTIONS---												
<= 00.00.23.165500 3 3 100 0.0 <= 00.00.05.656000 1 1 100 0.0												
<= 00.00.27.798600 3 0 100 0.0 <= 00.00.06.787200 1 0 100 0.0												
<= 00.00.32.431700 3 0 100 0.0 <= 00.00.07.918400 1 0 100 0.0												
<= 00.00.37.064800 3 0 100 0.0 <= 00.00.09.049600 1 0 100 0.0												
<= 00.00.41.697900 3 0 100 0.0 <= 00.00.10.180800 1 0 100 0.0												
<= 00.00.46.331000 3 0 100 0.0 <= 00.00.11.312000 1 0 100 0.0												
<= 00.00.50.964100 3 0 100 0.0 <= 00.00.12.443200 1 0 100 0.0												
<= 00.00.55.597200 3 0 100 0.0 <= 00.00.13.574400 1 0 100 0.0												
<= 00.01.00.230300 3 0 100 0.0 <= 00.00.14.705600 1 0 100 0.0												
<= 00.01.04.863400 3 0 100 0.0 <= 00.00.15.836800 1 0 100 0.0												
<= 00.01.09.496500 3 0 100 0.0 <= 00.00.16.968000 1 0 100 0.0												
<= 00.01.32.662000 3 0 100 0.0 <= 00.00.22.624000 1 0 100 0.0												
<= 00.03.05.324000 3 0 100 0.0 <= 00.00.45.248000 1 0 100 0.0												
> 00.03.05.324000 3 0 100 0.0 > 00.00.45.248000 1 0 100 0.0												

Figure 245. WLMGL - Service Class Period report - with execution velocity goal for Period 1 and response time distributions

[illegible]Chapter 5. Long-term overview reporting with the Postprocessor **483**

Figure 247 on page 484 shows a report for service class CICSLOW with subsystem delay data included, where the first period is defined with a percentile response time goal.

Note: For transaction service classes of subsystem work managers, such as CICS or IMS, no service consumption and execution delays are reported. This data is reported with the service classes for the regions.

WORKLOAD ACTIVITY															PAGE 1
z/OS 3.1		SYSPLEX SVPLEX3		RPT VERSION 3.1 RMF		DATE 03/12/2024		TIME 23.15.00		INTERVAL 15.00.003		MODE = GOAL			
POLICY ACTIVATION DATE/TIME 02/25/2024 09.00.02															
----- SERVICE CLASS PERIOD -----															
POLICY=BASEPOL		WORKLOAD=CICS_WLD		SERVICE CLASS=CICSLOW		RESOURCE GROUP=NONE		PERIOD=1		IMPORTANCE=3					
				CRITICAL		=CPU									

-TRANSACTIONS--		TRANS-TIME HHH.MM.SS.FFFFFF		TRANS-APPL%		--CP--CBPCP--IIPCP--CBP--IIP		--ENCLAVES--							
AVG		0.00 ACTUAL		50085		TOTAL		0.04 0.00 0.00 0.00 0.00		AVG ENC		0.00			
MPL		0.00 EXECUTION		605292		MOBILE		0.00 N/A 0.00 N/A 0.00		REM ENC		0.00			
ENDED		364 QUEUED		0		CATEGORYA		0.00 N/A 0.00 N/A 0.00		MS ENC		0.00			
END/S		0.40 R/S AFFIN		0		CATEGORYB		0.00 N/A 0.00 N/A 0.00							
#SWAPS		0 INELIGIBLE		0											
EXCTD		3 CONVERSION		0											
		STD DEV		156398											

SUB P		RESP TIME		--ACTIVE-- READY IDLE		----- STATE SAMPLES BREAKDOWN (%) -----		----- STATE -----							
TYPE		(%)		SUB APPL		MISC I/O TIME		SWITCHED SAMPL (%)							
CICS		BTE 5044K		0.0 0.0 0.0 16.6		72.9 0.0 10.5		LOCAL SYSPL REMOT							
CICS		EXE 4.1		100 0.0 0.0 0.0		0.0 0.0 0.0		0.0 0.0 0.0							
DB2		BTE 0.0		0.0 0.0 0.0 0.0		0.0 0.0 0.0		0.0 0.0 0.0							
DB2		EXE 9.6		28.6 0.0 0.0 0.0		42.9 28.6 0.0		0.0 0.0 0.0							

GOAL: RESPONSE TIME 000.00.00.500 FOR 80%															

SYSTEM		RESPONSE TIME		EX PERF											
		ACTUAL%		VEL% INDX											
*ALL		94.8		N/A 0.5											
CB8A		95.0		N/A 0.5											
CB8C		98.9		N/A 0.5											

-----RESPONSE TIME DISTRIBUTION-----															
-----TIME-----		---# TRANSACTIONS---		---% TRANSACTIONS---		0 10 20 30 40 50 60 70 80 90 100									
HH.MM.SS.FFFFFF		CUM TOTAL		IN BUCKET		CUM TOTAL		IN BUCKET		!.....!					

Figure 247. WLMGL - Service Class Period report - with subsystem data and response time distribution for response time goal

Figure 248 on page 485 shows the Enclave Transaction fields reported for a service class period for which transaction data for enclaves has been collected.

Note: These fields will be visible in the report only if the corresponding information is collected.

```

----- SERVICE CLASS PERIODS
POLICY=BASEPOL    WORKLOAD=DDF    SERVICE CLASS=DDF    RESOURCE GROUP=*NONE    PERIOD=1 IMPORTANCE=4
                    CRITICAL      =NONE
                    BATCH INITIATOR AI-INFUSED=NO

-TRANSACTIONS-- TRANS-TIME HHH.MM.SS.FFFFFFFF TRANS-APPL%----CP-IIPCP/AAPCP-IIP/AAP ---ENCLAVES---
AVG 55.00 ACTUAL 82501 TOTAL 9.58 0.03 19.83 AVG ENC 55.00
MPL 55.00 EXECUTION 80303 MOBILE 0.00 0.00 0.00 REM ENC 0.00
ENDED 1224457 QUEUED 2198 CATEGORYA 0.00 0.00 0.00 MS ENC 0.00
END/S 680.14 R/S AFFIN 0 CATEGORYB 0.00 0.00 0.00
#SWAPS 0 INELIGIBLE 0
EXCTD 0 CONVERSION 0
      STD DEV 2.201152

----SERVICE---- SERVICE TIME ---APPL %--- --PROMOTED-- --DASD I/O--- ---STORAGE--- -PAGE-IN RATES-
IOC 0 CPU 518.399 CP 9.58 BLK 0.013 SSCHRT 1583 AVG 0.00 SINGLE 0.0
CPU 50886K SRB 0.000 IIPCP 0.03 ENQ 0.000 RESP 0.0 TOTAL 0.00 BLOCK 0.0
MSO 0 RCT 0.000 IIP 19.83 CRM 0.000 CONN 0.0 SHARED 0.00 SHARED 0.0
SRB 0 IIT 0.000 AAPCP 0.00 LCK 0.000 DISC 0.0 HSP 0.0
TOT 50886K HST 0.000 AAP N/A SUP 0.000 Q+PEND 0.0
/SEC 28265 IIP 345.952 N/A SUP 0.000 IOSQ 0.0
ABSRPTN 514 AAP N/A
TRX SERV 514

-ENCLAVE TRANS- ENC-TR-TM HHH.MM.SS.FFFFFFFF
AVG 0.76 EXECUTION 6.947060
END/S 0.11 EX STD DEV 21.254062
ENC DEL 1
TR/ENC 197.00

```

Figure 248. WLMGL - Service Class Period report with transaction data for enclaves

Service Class report

To request this report, specify:

```
SYSRPTS(WLMGL(SCLASS(service_class)))
```

For example, to produce the report shown in Figure 249 on page 485 , specify:

```
SYSRPTS(WLMGL(SCLASS(STCLOW)) )
```

The report can contain an additional part SERVICE CLASSES BEING SERVED if address spaces in the service class in this report are doing work for transactions that were classified to another service class.

```

WORKLOAD ACTIVITY
z/OS 3.1          SYSPLEX UTCPLXJ8      DATE 02/25/2024    INTERVAL 30.27.072  MODE = GOAL    PAGE 1
RPT VERSION 3.1 RMF LEVEL 143          TIME 08.59.33

POLICY ACTIVATION DATE/TIME 02/18/2024 10.16.40

----- SERVICE CLASS(ES)
POLICY=WLMPOL01    WORKLOAD=BATCH    SERVICE CLASS=BATI1V70    RESOURCE GROUP=*NONE
                    CRITICAL      =NONE
                    BATCH INITIATOR AI-INFUSED=NO
                    DESCRIPTION   =BATCH Imp=1, VEL=70

-TRANSACTIONS-- TRANS-TIME HHH.MM.SS.FFFFFFFF TRANS-APPL%----CP-IIPCP/AAPCP-IIP/AAP ---ENCLAVES---
AVG 118.34 ACTUAL 38.884371 TOTAL 5.36 0.00 0.00 AVG ENC 0.00
MPL 118.34 EXECUTION 38.071067 MOBILE 0.00 0.00 0.00 REM ENC 0.00
ENDED 6 QUEUED 813303 CATEGORYA 0.00 0.00 0.00 MS ENC 0.00
END/S 0.00 R/S AFFIN 0 CATEGORYB 0.00 0.00 0.00
#SWAPS 0 INELIGIBLE 0
EXCTD 0 CONVERSION 349525
      STD DEV 3.390334

----SERVICE---- SERVICE TIME ---APPL %--- --PROMOTED-- --DASD I/O--- ---STORAGE--- -PAGE-IN RATES-
IOC 0 CPU 349.172 CP 19.45 BLK 0.000 SSCHRT 304.8 AVG 1511.11 SINGLE 0.0
CPU 19220K SRB 4.171 IIPCP 0.00 ENQ 0.000 RESP 0.1 TOTAL 178828.1 BLOCK 0.0
MSO 0 RCT 0.009 IIP 0.00 CRM 0.000 CONN 0.1 SHARED 118.36 SHARED 0.0
SRB 249974 IIT 2.067 AAPCP 0.00 LCK 1.957 DISC 0.0 HSP 0.0
TOT 19470K HST 0.000 AAP N/A SUP 0.000 Q+PEND 0.0
/SEC 10817 IIP 0.000 N/A SUP 0.000 IOSQ 0.0
ABSRPTN 90 AAP N/A
TRX SERV 90

POWER USAGE
-----PROCESSOR-----I/O-----MEMORY-----
SYSTEM -AVG POWR- --ENERGY-- -ENRG/MIN- -AVG POWR- --ENERGY-- -ENRG/MIN- -AVG POWR- --ENERGY-- -ENRG/MIN-
ALL 104.07 190152.21 6244.49 0.11 192.45 6.32 0.15 275.34 9.04
SYS1 32.29 58998.29 1937.47 0.02 38.98 1.28 0.03 50.06 1.64
SYS2 71.78 131154.23 4307.03 0.08 153.47 5.04 0.11 200.25 6.58

-----SERVICE CLASSES BEING SERVED-----
II390%P7

```

Figure 249. WLMGL - Service Class Report

If transaction data for enclaves has been collected for the service class, the Enclave Transaction fields will be displayed similar to the Service Class Period report. For information, see Figure 248 on page 485.

Note: These fields will be visible in the report only if the corresponding information is collected.

Workload Group report

To request this report, specify:

```
SYSRPTS(WLMGL(WGROUP(workload_group)))
```

For example, to produce the report for a workload group called OMVS_WLD, specify:

```
SYSRPTS(WLMGL(WGROUP(OMVS_WLD)))
```

The report has the same layout as a Service Class report, but all service classes associated with the workload OMVS_WLD are combined in a workload summary.

Workload Group and Service Class Period report

To request this report, specify:

```
SYSRPTS(WLMGL(WGPER(workload_group)))
```

Each service class associated with the specified workload group is listed with its service class periods and the defined and achieved GOALS for the periods.

For example, to produce the report for a workload group called STC_WLD with its associated service classes STCHIGH and STCLOW, specify:

```
SYSRPTS(WLMGL(WGPER(STC_WLD)))
```

W O R K L O A D A C T I V I T Y																																																																	
z/OS 3.1	SYSPLX SVPLX3	DATE 09/30/2021	INTERVAL 15.00.003	MODE = GOAL	PAGE 1																																																												
	RPT VERSION 3.1 RMF	TIME 23.15.00																																																															
POLICY ACTIVATION DATE/TIME 09/14/2019 09.00.02																																																																	
-----SERVICE CLASS(ES)																																																																	
POLICY=BASEPOL	WORKLOAD=STC_WLD	SERVICE CLASS=STCDISC	RESOURCE GROUP=NONE																																																														
		CRITICAL DESCRIPTION	=NONE	=For discretionary started tasks																																																													
-TRANSACTIONS--	TRANS-TIME HHH.MM.SS.FFFFFF	TRANS-APPL%	CP--CBPCP--IIPCP--CBP--IIP	---ENCLAVES---																																																													
AVG 12.03	ACTUAL	0	TOTAL 7.57	0.00	0.00	0.00	0.00	0.00	AVG ENC	0.00																																																							
MPL 12.03	EXECUTION	0	MOBILE	0.00	N/A	0.00	N/A	0.00	REM ENC	0.00																																																							
ENDED 0	QUEUED	0	CATEGORYA	0.00	N/A	0.00	N/A	0.00	MS ENC	0.00																																																							
END/S 0.00	R/S AFFIN	0	CATEGORYB	0.00	N/A	0.00	N/A	0.00																																																									
#SWAPS 16	INELIGIBLE	0																																																															
EXCTD 0	CONVERSION	0																																																															
	STD DEV	0																																																															
-----SERVICE-----																																																																	
IOC 0	CPU 67.913	CP 7.57	BLK 0.000	SSCHRT 0.0	AVG 420.55	SINGLE	0.0																																																										
CPU 4654K	SRB 0.242	IIPCP 0.00	ENQ 0.023	RESP 0.0	TOTAL 5057.37	BLOCK	0.0																																																										
MSO 0	RCT 0.000	IIP 0.00	CRM 0.000	CONN 0.0	SHARED 0.00	SHARED	0.0																																																										
SRB 16614	IIT 0.000	CBPCP 0.00	LCK 0.025	DISC 0.0		HSP	0.0																																																										
TOT 4671K	HST 0.000	CBP 0.00	SUP 0.000	Q+PEND 0.0																																																													
/SEC 5190	IIP 0.000		IOSQ 0.0																																																														
ABSRPTN 432	CBP 0.000																																																																
TRX SERV 432																																																																	
PER IMPORTANCE																																																																	
1	DISC	INDX	-TRANSACTIONS--	-NUMBER-	%-	-GOAL-	-RESPONSE TIME-	-ACTUAL-	TOTAL	-EX VEL%-	GOAL	ACT	TOTAL	-EXE--																																																			
				0	0	DISCRETIONARY				40.6			0.4	DELAY%																																																			
				0	0									0.6																																																			
POLICY=BASEPOL																																																																	
WORKLOAD=STC_WLD	SERVICE CLASS=STCHI	RESOURCE GROUP=NONE																																																															
		CRITICAL DESCRIPTION	=NONE	=High priority for system address																																																													
-TRANSACTIONS--	TRANS-TIME HHH.MM.SS.FFFFFF	TRANS-APPL%	CP--CBPCP--IIPCP--CBP--IIP	---ENCLAVES---																																																													
AVG 59.00	ACTUAL	0	TOTAL 0.57	0.00	0.00	0.00	0.00	0.00	AVG ENC	0.00																																																							
MPL 59.00	EXECUTION	0	MOBILE	0.00	N/A	0.00	N/A	0.00	REM ENC	0.00																																																							
ENDED 0	QUEUED	0	CATEGORYA	0.00	N/A	0.00	N/A	0.00	MS ENC	0.00																																																							
END/S 0.00	R/S AFFIN	0	CATEGORYB	0.00	N/A	0.00	N/A	0.00																																																									
#SWAPS 50604	INELIGIBLE	0																																																															
EXCTD 0	CONVERSION	0																																																															
	STD DEV	0																																																															
-----SERVICE-----																																																																	
IOC 115	CPU 4.824	CP 0.62	BLK 0.000	SSCHRT 3.0	AVG 23346.32	SINGLE	0.0																																																										
CPU 333756	SRB 0.731	IIPCP 0.00	ENQ 0.000	RESP 0.2	TOTAL 1377428	BLOCK	0.0																																																										
MSO 0	RCT 0.000	IIP 0.00	CRM 0.000	CONN 0.2	SHARED 24.00	SHARED	0.0																																																										
SRB 50604	IIT 0.000	CBPCP 0.00	LCK 0.003	DISC 0.0		HSP	0.0																																																										
TOT 384475	HST 0.000	CBP 0.00	SUP 0.000	Q+PEND 0.1																																																													
/SEC 427	IIP 0.000		IOSQ 0.0																																																														
ABSRPTN 7	CBP 0.000																																																																
TRX SERV 7																																																																	
PER IMPORTANCE																																																																	
1	2	INDX	-TRANSACTIONS--	-NUMBER-	%-	-GOAL-	-RESPONSE TIME-	-ACTUAL-	TOTAL	-EX VEL%-	GOAL	ACT	TOTAL	-EXE--																																																			
				0	0					35	23.0		0.0	DELAY%																																																			
				0	0									0.1																																																			

Figure 250. WLMGL - Workload Group with associated service classes

If the transaction data for enclaves has been collected for the workload group or any of the service class periods, the Enclave Transaction fields will be displayed similar to the Service Class Period report. For information, see [Figure 248 on page 485](#).

Note: These fields will be visible in the report only if the corresponding information is collected.

Policy Summary report

To request this report, specify:

```
SYSRPTS(WLMGL(POLICY(policy)))
```

For example, to produce a report for a policy named HOLIDAY, specify:

```
SYSRPTS(WLMGL(POLICY(HOLIDAY)))
```

Except for the separation line, a Policy Summary report has the same layout as the Workload Group report, and combines all workload groups associated to the policy.

Report Class report

To request this report, specify:

```
SYSRPTS(WLMGL(RCLASS(report_class)))
```

For example, to produce a report for a report class called BATCH, specify:

```
SYSRPTS(WLMGL(RCLASS(BATCH)))
```

Optionally, classification rules can assign incoming work to either a report class or to a tenant report class. Report classes are for additional reporting data, across service classes, or for monitoring special work. Tenant report classes have the same purpose but can additionally be subject to WLM capping if the tenant report class is associated with a tenant resource group that has assigned CPU or memory consumption limits.

```

COMMAND INPUT ==>>                                SCROLL ==> CSR
***** TOP OF DATA ***** WORKLOAD ACTIVITY *****
z/OS V3R1      SYSPLX RSPLEX01      DATE 09/21/2022      INTERVAL 15.00.261      MODE = GOAL      PAGE 1
SOME RECS CONVERTED TO V3R1      TIME 23.44.33

POLICY ACTIVATION DATE/TIME 09/21/2022 08.07.10

----- REPORT CLASS(ES) -----
POLICY=SAMPLE01      REPORT CLASS=CICSRPT
DESCRIPTION =CICS Report Group

-TRANSACTIONS-- TRANS-TIME HHH.MM.SS.FFFFFF TRANS-APPL%----CP-IIPCP/AAPCP-IIP/AAP ---ENCLAVES---
AVG      0.00 ACTUAL      196725 TOTAL      0.00      0.00      0.00 AVG ENC      0.00
MPL      0.00 EXECUTION      0 MOBILE      0.00      0.00      0.00 REM ENC      0.00
ENDED      4 QUEUED      0 CATEGORYA      0.00      0.00      0.00 MS ENC      0.00
END/S      0.00 R/S AFFIN      0 CATEGORYB      0.00      0.00      0.00
#SWAPS      0 INELIGIBLE      0
EXCTD      0 CONVERSION      0
STD DEV      250914

POLICY=SAMPLE01      REPORT CLASS=MXIPROD
DESCRIPTION =MXI Production

-TRANSACTIONS-- TRANS-TIME HHH.MM.SS.FFFFFF TRANS-APPL%----CP-IIPCP/AAPCP-IIP/AAP ---ENCLAVES---
AVG      18.01 ACTUAL      0 TOTAL      4.08      0.00      0.42 AVG ENC      0.00
MPL      18.01 EXECUTION      0 MOBILE      0.00      0.00      0.00 REM ENC      0.00
ENDED      0 QUEUED      0 CATEGORYA      0.00      0.00      0.00 MS ENC      0.00
END/S      0.00 R/S AFFIN      0 CATEGORYB      0.00      0.00      0.00
#SWAPS      0 INELIGIBLE      0
EXCTD      0 CONVERSION      0
STD DEV      0

----SERVICE----- SERVICE TIME ---APPL %--- --PROMOTED-- --DASD I/O--- ---STORAGE--- --PAGE-IN RATES--
IOC      0 CPU      41.351 CP      4.09 BLK      0.000 SSCHRT      26.6 AVG      9983.78 SINGLE      0.0
CPU      3821K SRB      1.884 IIPCP      0.00 ENQ      0.000 RESP      0.2 TOTAL      179789.9 BLOCK      0.0
MSO      0 RCT      0.000 IIP      0.42 CRM      0.000 CONN      0.2 SHARED      27.01 SHARED      0.0
SRB      173974 IIT      0.064 AAPCP      0.00 LCK      0.433 DISC      0.0 HSP      0.0
TOT      3995K HST      0.000 AAP      N/A SUP      0.000 Q+PEND      0.0
/SEC      443K IIP      6.474
ABSRPTN      246 AAP      N/A
TRX SERV      246

POWER USAGE
-----PROCESSOR----- I/O-----MEMORY-----
SYSTEM -AVG POWR- --ENERGY-- -ENRG/MIN- -AVG POWR- --ENERGY-- -ENRG/MIN- -AVG POWR- --ENERGY-- -ENRG/MIN-
ALL      104.07 93694.51 6244.49 0.11 94.83 6.32 0.15 135.67 9.04
SYS1      32.29 29070.48 1937.47 0.02 19.21 1.28 0.03 24.67 1.64
SYS2      71.78 64624.19 4307.03 0.08 75.62 5.04 0.11 98.67 6.58

WORKLOAD ACTIVITY

z/OS V3R1      SYSPLX RSPLEX01      DATE 09/21/2022      INTERVAL 15.00.261      MODE = GOAL      PAGE 2
SOME RECS CONVERTED TO V3R1      TIME 23.44.33
F1=HELP F2=SPLIT F3=END F4=RETURN F5=RFIN D F6=RCHANGE F7=UP F8=DOWN F9=SWAP F10=LEFT F11=RIGHT F12=RETRIEVE

```

Figure 251. WLMGL - Report Class Report

The report has nearly the same contents as the Service Class report. There are only two differences:

- the workload group name cannot be associated to a report class and is therefore not reported.
- the tenant identifier that is associated with the tenant resource group is reported for tenant report classes but not for report classes. The same applies to the tenant resource group name.

Report Class Period report

To request this report, specify:

```
SYSRPTS(WLMGL(RCPER(report_class))
```

The report has nearly the same contents as the Service Class Period report, but reporting of subsystem delays and response time distribution data is possible only for homogeneous report class periods.

All goal-related data in the report is based on the corresponding service class period.

A report class period is called **homogeneous** if all its transactions are being assigned to the same service class period.

Example: You classify all TSO users to run in service class TSOPROD and distinguish the departments for reporting purposes in report classes 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 is created automatically for each interval. This page contains the following information:

- Information about the installation of the service definition
- Service definition coefficients and normalization factors for special purpose processors
- Information about each system, including the system ID, parmlib member suffix, service units/second, effective capacity percentage, and the interval start time and length
- Resource group and tenant resource group definitions
- Resource group actuals for each service class
- Tenant resource group actuals for each tenant report 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																
z/OS 3.1		SYSPLEX UTCPLXCB				DATE 09/30/2021				INTERVAL 14.59.999		MODE = GOAL		PAGE 2		
		RPT VERSION 3.1 RMF				TIME 03.00.00										
POLICY ACTIVATION DATE/TIME 09/14/2019 09.00.04																
- SERVICE POLICY PAGE -																
SERVICE DEFINITION: CMBSVDEF WLM BASEPOL COMBAT for WBG																
INSTALL DATE: 09/07/2019 13.47.34 INSTALLED BY: RDOWNNEY																
POLICY: BASEPOL WLM BASEPOL COMBAT for WBG																
DISCRETIONARY GOAL MANAGEMENT: YES																
DYNAMIC ALIAS MANAGEMENT: YES																
I/O PRIORITY MANAGEMENT: YES																
-SERVICE DEFINITION COEFFICIENTS-																
		IOC		CPU		SRB		MSO		-NORM FACTORS-		AAP IIP				
		0.0		1.0		1.0		0.0000		1.0000		1.0000				
SYSTEMS																
---ID---		OPT	SU/SEC	CAP%	--TIME--	INTERVAL	--BOOST--	---ID---		OPT	SU/SEC	CAP%	--TIME--	INTERVAL	--BOOST--	
CB8E		CB	59259.3	100	03.00.00	00.14.59	Inactive	CB89		CB	68965.5	100	03.00.00	00.14.59	Inactive	
RESOURCE GROUPS																
--NAME--		TYPE		-----DESCRIPTION-----			-SYSTEM-		---CPU CONSUMPTION---			-----CPU CAPACITY-----			-----MEMORY-----	
									#CPS MSU SU/SEC			MIN MAX DEFINED AS			USAGE LIMIT	
HWTRG1		TRG		Hardware Container TRG #1					0.22 27 14K							
							CB8E		0.11 14 7K						98M	
							CB89		0.11 13 8K						102M	
				-----REPORT CLASSES			HWTRC1		0.17 21 11K							
							HWTRC3		0.05 6 3K							
HWTRG2		TRG		Hardware Container TRG #2					0.21 25 13K							
							CB8E		0.11 13 6K						111M	
							CB89		0.10 12 7K						131M	
				-----REPORT CLASSES			HWTRC2		0.21 25 13K							
RGPTYP1		RG		V1r12 LI1370 testing, type 1					0.00 0 0			5 100 SU/SEC				
							CB8E		0.00 0 0						0	
							CB89		0.00 0 0						0	
				-----SERVICE CLASSES			RGRP1		0.00 0 0							
RGPTYP3		RG		v1r12 li1370 testing, type 3					0.00 0 0			0.01 1.00 NUMBER OF CPs				
							CB8E		0.00 0 0						0	
							CB89		0.00 0 0						0	
				-----SERVICE CLASSES			RGRP3		0.00 0 0							
							RGRP3B		0.00 0 0							
							RGRP3C		0.00 0 0							

Figure 252. WLMGL Report - Service Policy Page

Field descriptions of all reports

Table 222. Fields in the Workload Activity Report	
Field Heading	Meaning
CRITICAL	This indication reports on the specification in the service policy, not whether SRM is using or ignoring it. This information can be found in Monitor II (ARD report) and in Monitor III (for example, DELAY report). CPU CPU critical CPUI Service Class Period implicitly CPU critical CPUI* First period of Service Class implicitly CPU critical STORAGE Storage critical CPU + STORAGE Both CPU and storage critical CPUI+STORAGE Both implicitly CPU and storage critical (Service Class Period) CPUI*+STORAGE Both first period implicitly CPU and storage critical (Service Class) NONE Neither storage nor CPU critical
BATCH INITIATOR AI-INFUSED	This field indicates whether the batch initiator for this service class is AI-infused. The values can be: YES or NO.
HONOR PRIORITY=NO	This indication is reported for service classes that do not allow to overflow to regular CPs when there is insufficient zIIP or zAAP capacity.
I/O PRIORITY GROUP=HIGH	This indication is reported for service classes assigned to I/O priority group HIGH in the active service policy.
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/S 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.

Table 222. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
TRANS-TIME	Transaction time in HHH.MM.SS.FFFFFFF 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 AFFIN, 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 AFFIN 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.
Transaction Application Time % (TRANS-APPL%) When transaction processor usage is reported to WLM through IWM4RPT or IWM4MNTF services, the consumed service units are accounted to the transaction service or report classes, and deducted from the region's service and report classes. If the number of transactions is very small and a single transaction reports high processor times, it can occur that processor times become negative. In such a case RMF displays asterisk (*).	
TOTAL	Total percentage of the processor time used by transactions running on the different processor types. CP Total percentage of general purpose processor time used by transactions. CBP Total percentage of zCBP processor time used by transactions. CBPCP Total percentage of general purpose processor time used by transactions eligible to run on zCBP processors. IIPCP or IIPCP/AAPCP Total percentage of general purpose processor time used by transactions eligible to run on specialty processors. IIP or IIP/AAP Total percentage of specialty processor time used by transactions.

Table 222. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
MOBILE CATEGORYA CATEGORYB	Percentage of the processor time used by transactions classified with reporting attribute MOBILE / CATEGORYA / CATEGORYB running on the different processor types. MOBILE / CATEGORYA / CATEGORYB is a subset of TOTAL. CP Percentage of general purpose processor time used by transactions classified with reporting attribute MOBILE / CATEGORYA / CATEGORYB. IIPCP or IIPCP/AAPCP Percentage of general purpose processor time used by transactions classified with reporting attribute MOBILE / CATEGORYA / CATEGORYB, eligible to run on specialty processors. IIP or IIP/AAP Percentage of specialty processor time used by transactions classified with reporting attribute MOBILE / CATEGORYA / CATEGORYB. Note: N/A for zCBP processors
ENCLAVES	Average number of enclaves. 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 The 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.
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). CONN Average DASD connection time of the transactions in this group, as reported by the channel measurement subsystem. DISC Average DASD disconnect time of the transactions in this group, as reported by the channel measurement subsystem. Q+PEND Average DASD wait time (queue time + pending time) of the transactions in this group. This does not include IOSQ time, as reported by the channel measurement subsystem. IOSQ Average time the transactions in this group spent on the IOS queue, based on sampled delays.

Table 222. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
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. ABSRPTN Absorption rate at which service is used while transactions are resident in main storage. This is the total service divided by the transaction residency time. 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.
SERVICE TIME	This category is made up of the following: CPU Task and preemptible-class SRB (enclave) time in seconds consumed on general purpose and special purpose processors. 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. AAP zAAP service time in seconds. CBP zCBP service time in seconds. IIP zIIP service time in seconds. Note: 1. If special purpose processors are running faster than general purpose processors, AAP, CBP and IIP times are not normalized. 2. Normalized AAP, CBP and IIP times are included in CPU time.

Table 222. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
APPL%	Percentage of the processor time used by transactions running on the different processor types. The calculation is: $\text{APPL\%} = \frac{\text{Processor time used}}{\text{Interval length} * \text{Multithreading maximum capacity factor}} * 100$ CP Percentage of the processor time used by transactions running on general purpose processors in the service or report class period. The calculation of the processor time is based on the time values displayed under field heading SERVICE TIME. Processor time used = CPU + SRB + RCT + IIT + HST – (AAPNF * AAP) – (CBPNF * CBP) – (IIPNF * IIP) The AAP, CBP and IIP times may be normalized to general purpose processor time from a faster zAAP, zCBP or zIIP where AAPNF, CBPNF and IIPNF are the zAAP, zCBP and zIIP normalization factors. They can be found under field heading NORM FACTORS on the POLICY page of the Workload Activity report. AAPCP Percentage of the processor time used by zAAP eligible transactions running on general purpose processors. This is a subset of APPL% CP. CBPCP Percentage of the processor time used by zCBP eligible transactions running on general purpose processors. This is a subset of APPL% CP. IIPCP Percentage of the processor time used by zIIP eligible transactions running on general purpose processors. This is a subset of APPL% CP. AAP Percentage of the processor time used by transactions executed on zAAPs in the service or report class period. CBP Percentage of the processor time used by transactions executed on zCBPs in the service or report class period. IIP Percentage of the processor time used by transactions executed on zIIPs in the service or report class period. Notes: 1. APPL% shows processor utilization based on uniprocessor capacity. This means that the values can exceed 100% in systems with more than one processor. 2. If the multithreading mode is set to 1, a multithreading maximum capacity factor of 1 is used for the APPL% calculation. 3. The interval length in a sysplex is the common interval length. 4. In a sysplex, the values for seconds and CPU time percentages are meaningful only if all processors have the same speed and the multithreading mode is the same on all systems. You can use the SYSRPTS WLMGL SYSNAM option to select only a subset of the systems to be included in the report. 5. AAPCP, zCBPCP or IIPCP may report values greater than zero even if no special purpose processors are configured or if they are varied offline, because the PROJECTCPU option is specified in the active IEAOPT Parmlib member. This information can be used to understand the benefit of adding special purpose processors to your system.

Table 222. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
PROMOTED	CPU time in seconds that transactions in this group were running at a promoted dispatching priority, separated by the reason for the promotion: BLK CPU time in seconds consumed while the dispatching priority of work with low importance was temporarily raised to help blocked workloads ENQ CPU time in seconds consumed while the dispatching priority was temporarily raised by enqueue management because the work held a resource that other work needed. CRM CPU time in seconds consumed while the dispatching priority was temporarily raised by chronic resource contention management because the work held a resource that other work needed LCK In HiperDispatch mode, the CPU time in seconds consumed while the dispatching priority was temporarily raised to shorten the lock hold time of a local suspend lock held by the work unit. SUP CPU time in seconds consumed while the dispatching priority for a work unit was temporarily raised by the z/OS supervisor to a higher dispatching priority than assigned by WLM.
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. SHARED Total number of shared storage pages allocated to resident ASIDs.
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. Note: Enclave transaction residency time is not included in the calculation of these values because there is no paging on behalf of enclaves.
MOBILE	Percentage of the processor time used by transactions classified with reporting attribute MOBILE running on the different processor types. MOBILE is a subset of TOTAL. CP Percentage of general purpose processor time used by transactions classified with reporting attribute MOBILE. AAP/IIP ON CP Percentage of general purpose processor time used by transactions classified with reporting attribute MOBILE, eligible to run on specialty processors. AAP/IIP Percentage of specialty processor time used by transactions classified with reporting attribute MOBILE.

Table 222. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
PROCESSOR	The Power Consumption report includes three columns of values: AVG POWER The average power usage by all CPU resources allocated to this workload over the time interval, shown in watts. ENERGY The amount of energy consumed by all CPU resources allocated to this workload over the time interval, shown in watt-seconds. ENG/MINUTE The average amount of energy consumed by all CPU resources allocated to this workload during one minute, shown in watts-seconds.
I/O	The Power Consumption report includes three columns of values: AVG POWER The average power usage by all I/O resources allocated to this workload over the time interval, shown in watts. ENERGY The amount of energy consumed by all I/O resources allocated to this workload over the time interval, shown in watt-seconds. ENG/MINUTE The average amount of energy consumed by all I/O resources allocated to this workload during one minute, shown in watts-seconds.
MEMORY	The Power Consumption report includes three columns of values: AVG POWER The average power usage by all storage resources allocated to this workload over the time interval, shown in watts. ENERGY The amount of energy consumed by all storage resources allocated to this workload over the time interval, shown in watt-seconds. ENG/MINUTE The average amount of energy consumed by all storage resources allocated to this workload during one minute, shown in watts-seconds.
Service Classes being Served	
SERVICE CLASSES BEING SERVED	This section is only available if address spaces are doing work for transactions that were classified to another service class. The name of each service class being served by the reported service class (see name in the separation line) is displayed.
Work Manager/Resource Manager	
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 222. 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 processed 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.

Table 222. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
STATE SAMPLES BREAKDOWN (%) - WAITING FOR	STATE SAMPLES BREAKDOWN (%) - continued. This category presents up to fifteen named delay reason states having the highest non-zero values. These are sorted by the sum of the BTE and EXE rows in each column. The values of the remaining less important states will be accumulated, if applicable, and presented as delay reason OTHR. Here 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. BPMI waiting for I/O resulting from a Db2 buffer pool miss. TYPn/TYnn The generic delay state defined by the subsystem (possible values: TYP1 - TYP9 and TY10 - TY15). If the subsystem uses the WLM service IWM4MGDD (Define Descriptions for Generic Delay States) to provide a description for a generic delay state, RMF displays a legend with the delay state description. If the subsystem did not use the IWM4MGDD service, the legend is omitted. For further explanation of the generic delay state types please refer to the subsystem documentation.

Table 222. 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 and actual values If measurement data for systems in a sysplex is available, this section starts with an *ALL line showing the average or cumulative values for the sysplex. The *ALL line is followed by one line for each system.	
GOAL	This line shows the goal specified in the WLM service policy for a service class period. For information about available goal types refer to <i>z/OS MVS Planning: Workload Management</i>. In a report for a homogeneous report class period, the goal of the corresponding service class period is printed. For heterogeneous report class periods, N/A is printed.
VELOCITY MIGRATION	The following two values are only provided for periods with an execution velocity goal: I/O MGMT / NO I/O MGMT I/O Priority Management The I/O MGMT of Velocity Migration is bidirectional, which means that the reported value of the I/O MGMT is associated with the I/O Priority Management setting. If I/O Priority Management is set to YES, the reported value of I/O MGMT is the Execution Velocity if WLM I/O Management is disabled. Similarly, if I/O Priority Management is set to NO, the reported value of I/O MGMT is the Execution Velocity value if WLM I/O Management is enabled. When the WLM I/O Management field is set to YES, the Velocity Migration is displayed in a NO I/O MGMT field. Similarly, when the WLM I/O Management field is set to NO, the Velocity Migration is displayed in an I/O MGMT field. You can refer to current definition in the Service Policy page (see Figure 252 on page 489). 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%. If WLM batch initiator management is enabled in the JES settings, 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.
RESPONSE TIME	This column either shows: • for an AVG response time goal: the measured average response times • for a percentile response time goal: the percentages of the transactions that met the response time goal • for an execution velocity goal, a system or a discretionary goal: N/A • for heterogeneous report class periods: N/A
EX VEL %	The execution velocity measures the portion of the acceptable processor and storage delays relative to the total execution time. For details about the execution velocity, see “Common Monitor III report measurements” on page 10 .
PERF INDX	The performance index for a period represents how close a period came to reaching the goal (PI is 1.0 if goal is reached), and how much this period suffered versus its goal. See Table 88 on page 197 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.

Table 222. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
EXEC USING%	The following using samples are measured as percentages of the total samples: CPU Standard CP using samples. This value includes using samples of zAAP, zIIP and zCBPwork executing on general purpose processors (standard CPs). AAP zAAP using samples. CBP zCBP using samples. IIP zIIP using samples. I/O I/O using samples. Note: 1. Use the APCUSGP (AAP on CP Using%) overview condition to retrieve the using samples of zAAP work executing on general purpose processors (standard CPs). 2. Use the IPCUSGP (IIP on CP Using%) overview condition to retrieve the using samples of zIIP work executing on standard CPs. 3. Use the CBCUSGP (CBP on CP Using%) overview condition to retrieve the using samples of zCBP work executing on general purpose processors (standard CPs).

Table 222. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
EXEC DELAYS %	General execution delays included in TOT (total). Each dispatchable unit sampled can increase one of the CPU or paging delay samples. Besides the TOT value, only the seven highest values contributing to TOT will be shown. The remaining less important values will be accumulated and presented as OTH. TOT 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. AAP zAAP-eligible work is delayed because it is waiting for a processor that can run zAAP work. CBP zCBP-eligible work is delayed because it is waiting for a processor that can run zCBP work. IIP zIIP-eligible work is delayed because it is waiting for a processor that can run zIIP work. 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 “CONN” on page 492 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. CAP 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. SIN Swap-In delay. Swap-In has started but not completed. MPL MPL delay. Ready but swap-in has not started. Q MPL Queue MPL - work is waiting for a server address space or batch initiator. SRV PRV Private area paging delay for a server address space. SRV VIO VIO paging delay for a server address space. SRV SHS Hiperspace paging delay for a server address space. SRV SIN Swap-in delay for a server address space. SRV MPL MPL delay for a server address space. AUX PRV Auxiliary paging from private. AUX COM Auxiliary paging from common. AUX XME Auxiliary paging from cross memory. AUX VIO Auxiliary paging from VIO. AUX SHS Auxiliary paging from standard hiperspaces. AUX EHS 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 222. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
USING%	Percentage of using states: CRY Crypto using state — a TCB or SRB was found to be using an adjunct processor (AP). CNT Contention using state - work is holding resources.
DELAY %	The following states are NOT included in the TOTAL EXECUTION DELAYS. UNK State is unknown. The address space or enclave was not found to be using or delayed for any sampled resource, but z/OS has not been notified that it is idle. IDL Idle state. Work is in STIMER wait, TSO terminal wait, APPC wait, OMVS input or output wait, or an initiator is waiting for work. CRY Crypto delay state — a TCB or SRB was found to be waiting for an AP or a processor feature queue. CNT Contention delay state - work is waiting for resources.
% QUI	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 (for service/report class periods with a response time goal only)	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 sysplex-wide view on 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. # 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 % TRANSACTIONS 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

Table 222. Fields in the Workload Activity Report (continued)

Field Heading	Meaning
RESPONSE TIME DISTRIBUTIONS (for service/report class periods with an execution velocity goal only)	WLM maintains counts of how many transactions were completed within a particular time. The average of all response times (midpoint) measured for each service/report class period is split into 14 response time buckets where: • bucket 1 covers the gap from 0 to half the midpoint • buckets 2 to 11 cover the gap between half to 1.5 times the midpoint evenly divided • bucket 12 covers two times the midpoint • bucket 13 covers four times the midpoint • bucket 14 covers the gap from four times the midpoint to infinity For each system in the sysplex, the report presents a tabular representation of: • the number of total (ended) transactions • the response time • the number of total (ended) transactions in percent. SYSTEM System Name INTERVAL Measurement interval (Time since last midpoint change) MRT CHANGES Number of midpoint changes during the SMF interval TIME Response time associated to this bucket # 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 % TRANSACTIONS Percentage CUM TOTAL Cumulative percentage of transactions so far IN BUCKET Percentage of transactions associated to the bucket

Table 223. 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 DEFINITION COEFFICIENTS	Service definition 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. Starting with z/OS 3.1, service definition coefficients are preset to the following values and cannot be changed: IOC = 0, CPU = 1, SRB = 1, and MSO = 0.

Table 223. Fields in the WLMGL Report - POLICY (continued)

Field Heading	Meaning
NORM FACTORS	Normalization factors for special purpose processors: AAP Normalization factor for zAAP. Multiply zAAP service times or service units with this value to calculate the CP equivalent value. CBP Normalization factor for zCBP. Multiply zCBP service times or service units with this value to calculate the CP equivalent value. IIP Normalization factor for zIIP. Multiply zIIP service times or service units with this value to calculate the CP equivalent value.
POLICY	Policy name and description.
DISCRETIONARY GOAL MANAGEMENT	Discretionary goal management allows certain types of work, when overachieving their goals, to be capped in order to give discretionary work a better chance to run. YES If the active WLM service policy allows discretionary goal management (WLM default). NO If the active WLM service policy disallows discretionary goal management.
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 Nominal capacity rating in service units per second per online CPU. CAP% Percentage of effective capacity available to the CPU. The value is 100, if the machine is working at its full nominal capacity. If the machine is working in power-save mode or cycle-steering mode, the value is less than 100. If the nominal or effective processor capacity cannot be determined, N/A is reported. TIME Begin time of the interval for this system. INTERVAL Interval length for this system in HH.MM.SS. BOOST Indicates whether boost was inactive or active at the end of the interval during IPL, shutdown or recovery. If active, the boost type is shown: I zIIP boost. S Speed boost. A zIIP and speed boost were both active.

Table 223. 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, a memory limit per system, or a combination of both. The report includes resource group name and description as: NAME Name of resource group or tenant resource group. TYPE Type of resource group • RG: Resource group • TRG: Tenant resource group with associated work that cannot run on CBP • CBP: Tenant resource group with associated work that can run on CBP DESCRIPTION Description of resource group or tenant resource group. SYSTEM Name of a MVS system that is belonging to the sysplex. CPU CONSUMPTION #CPS CPU consumption in terms of number of general purpose processors. The scope of the reported value can be: • this resource group • a particular system in the sysplex • a service class that is associated with this resource group • a tenant report class that is associated with this tenant resource group CPU CONSUMPTION MSU CPU consumption in terms of MSU. The scope of the reported value can be: • this resource group • a particular system in the sysplex • a service class that is associated with this resource group • a tenant report class that is associated with this tenant resource group CPU CONSUMPTION SU/SEC CPU consumption in terms of unweighted CPU and SRB service units per second. The scope of the reported value can be: • this resource group • a particular system in the sysplex • a service class that is associated with this resource group • a tenant report class that is associated with this tenant resource group If an asterisk (*) is placed behind the reported total value, WLM capping considers service units consumed on general purpose and zCBP processors as well as specialty processors. In such a case, the reported value reflects the actual general purpose, zCBP and specialty processor consumption. Otherwise, only general purpose and zCBP processor consumption is reported. CPU CAPACITY Resource group capacity limits. For tenant resource groups with associated work that can run on CBP, the minimum and maximum of service applies to general purpose and zCBP processors. MIN Minimum amount of service that the resource group should receive if demand exists. MAX Maximum amount of service that the resource group should be allowed to consume. DEFINED AS The method how the resource group's capacity is defined: • SU/SEC: in unweighted CPU and SRB service units per second across the sysplex • % LPAR SHARE: as percentage of the LPAR share on each system in the sysplex • NUMBER OF CPs: as number of general purpose processors on each system in the sysplex • MSU: as MSU/h capacity across the sysplex MEMORY USED Amount of memory used by this resource group on a particular system. MEMORY LIMIT Resource group memory limit for each system in the sysplex. Memory consumption of work running in address spaces associated with this resource group is limited to this size.

Table 223. Fields in the WLMGL Report - POLICY (continued)

Field Heading	Meaning
ENCLAVE TRANS	-AVG The average number of executing enclave bundled transactions during the interval. -END/S Enclave transactions per second. -ENC DEL Number of times transaction data provided with enclave delete. -TR/ENC Average number of enclave transactions per enclave. For illustration, see Figure 248 on page 485. Note: These fields will be visible in the report only if the corresponding information is collected.
ENC-TR-TM	-EXECUTION Average execution time of enclave transactions. -EX STD DEV Standard deviation of enclave transaction execution time. For illustration, see Figure 248 on page 485. Note: These fields will be visible in the report only if the corresponding information is collected.

Spreadsheet and Overview reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*. The following table shows the exception and overview condition names for the Overview report.

Table 224. Exception and Overview names in the Workload Activity Report

Field Heading or Meaning	Overview Name
TRANSACTIONS -	
- AVG	TRANSAVG
- MPL	TRANSMPL
- ENDED	TRANSTOT
- END/S	TRANS
- #SWAPS	SPERTRA
TRANS TIME -	
- ACTUAL	RTIMETOT
- EXECUTION	RTIME
- QUEUED	RTIMEQUE
- R/S AFFIN	TRANSADT
- INELIGIBLE	TRANSIQT
- CONVERSION	TRANSCVT
TRANS-APPL% -	
- TOTAL CP	TAPPLCP
- TOTAL IIP/AAP ON CP	TAPPLOCP
- TOTAL IIP/AAP	TAPPLSP
- TOTAL CBP	TAPPLCB
- MOBILE CP	MAPPLCP
- MOBILE IIP/AAP ON CP	MAPPLOCP
- MOBILE IIP/AAP	MAPPLSP
- CATEGORYA CP	AAPPLCP

Table 224. Exception and Overview names in the Workload Activity Report (continued)

Field Heading or Meaning	Overview Name
- CATEGORYA IIP/AAP ON CP	AAPPLOCP
- CATEGORYA IIP/AAP	AAPPLSP
- CATEGORYB CP	BAPPLCP
- CATEGORYB IIP/AAP ON CP	BAPPLOCP
- CATEGORYB IIP/AAP	BAPPLSP
ENCLAVES	
- AVG ENC	ENCAVG
- REM ENC	ENCREM
- MS ENC	ENCMS
SERVICE -	
Note: The Overview output is a rate, not a count. For details see <i>z/OS Resource Measurement Facility User's Guide</i> .	
- IOC	IOSRV
- CPU	CPUSRV
- MSO	MSOSRV
- SRB	SRBSRV
- TOT	TOTSRV
- ABSRPTN	ABSRPTN
- TRX SERV	TRXSERV
SERVICE TIME -	
- CPU	TCBSEC/TCBPER
- SRB	SRBSEC/SRBPER
- RCT	RCTSEC
- IIT	IITSEC
- HST	HSTSEC
- AAP	AAPSEC/AAPNSEC
- IIP	IIPSEC/IIPNSEC
- CBP	CBPSEC/CBPNSEC
APPL % -	
- CP	APPLSEC/APPLPER
- AAPCP	APPLIFCP/APPLAPCP
- IIPCP	APPLIPCP
- CBPCP	APPLCBCP
- AAP	APPLAAP
- IIP	APPLIIP
- CBP	APPLCBP
- AAP time on general purpose processors	AAPCPSEC
- CBP time on general purpose processors	CBPCPSEC
- Total number of EXCPs	EXCP
- EXCP rate	EXCPRT

Table 224. Exception and Overview names in the Workload Activity Report (continued)

Field Heading or Meaning	Overview Name
PROMOTED -	
- BLK	PROMSEC/PROMPER
- ENQ	EPROMSEC/EPROMPER
- CRM	CPROMSEC/CPROMPER
- LCK	LPROMSEC/LPROMPER
- SUP	SPROMSEC/SPROMPER
DASD I/O -	
- SSCHRT	SSCHRT
- RESP	RESP
- CONN	CONN
- DISC	DISC
- Q+PEND	QPEND
- IOSQ	IOSQ
STORAGE -	
- TOTAL	STOTOT
- CENTRAL	STOCEN
- SHARED	STOSHR
PAGE-IN RATES -	
- SINGLE	SINGLE
- BLOCK	BLOCK
- SHARED	SHARED
- HSP	HSP
EX VEL %	EXVEL
PERF INDEX	PI
USING% -	
- CPU	CPUUSGP
- AAP	AAPUSGP
- IIP	IIPUSGP
- CBP	CBPUSGP
- I/O	IOUSGP
- AAP on CP	APCUSGP
- IIP on CP	IPCUSGP
- CBP on CP	CBCUSGP
EXECUTION DELAYS % -	
- CPU	CPUDLYP
- AAP	AAPDLYP
- IIP	IIPDLYP
- CBP	CBPDLYP
- I/O	IODLYP

Table 224. Exception and Overview names in the Workload Activity Report (continued)

Field Heading or Meaning	Overview Name
- CAPP	CAPP
- SWIN	SWINP
- MPL	MPLP
- QMPL	QUEUEP
- Total Server Delays	SERV
- Total Storage Delays	STOP
DLY% -	
- UNKN	UNKP
- IDLE	IDLEP
CRYPTO% -	
- DLY	CRYDLYP, CAPDLYP, APDLYP, FQDLYP
- HSP	HSP
ENCLAVE TRANS	
-AVG	ENCTRAVG
-END/S	ENCTRRTE
-ENC DEL	ENCTRDEL
-TR/ENC	ENCTRENC
ENC-TR-TM	
-EXECUTION	ENCTRTEX
-EX STD DEV	ENCTRES

XCF - Cross-System Coupling Facility Activity report

The Cross-System Coupling Facility Activity report shows the XCF data from one system's processing in a sysplex. To better understand the traffic on corresponding outbound and inbound signalling paths, you might have to run RMF reports on two or more systems.

How to request this report

Monitor III gathers data for this report automatically. If you want to suppress gathering, you have to disable writing SMF record type 74.2.

To produce this report, specify

```
REPORTS(XCF)
```

This report is also available in XML output format. How to work with Postprocessor XML reports in z/OS *Resource Measurement Facility User's Guide* provides all required information on how to produce and view XML reports.

Example URL for the DDS API

```
http://ddshost:8803/gpm/rmfpp.xml?reports=XCF
```

Contents of the report

The XCF Activity report is divided into three sections:

- XCF Usage by System
- XCF Usage by Member
- XCF Path Statistics

The **Usage by System** section gives information about messages sent to and received from each remote system in the sysplex, broken down by transport class. Use this section to check the class lengths and message buffer space parameters. For a sample of the XCF Activity Usage by System section, see [Figure 253 on page 510](#).

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 254 on page 511](#).

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 255 on page 512](#).

Note: If the XCF system, path, or member becomes inactive during the RMF interval, the appropriate counters will be reinitialized. This is indicated in the report by the message *COUNTS RESET.

X C F A C T I V I T Y																
z/OS 3.1			SYSTEM ID TRX1				DATE 09/30/2021				INTERVAL 15.00.000				PAGE 1	
			RPT VERSION 3.1 RMF				TIME 09.15.00				CYCLE 1.000 SECONDS					
XCF USAGE BY SYSTEM																

REMOTE SYSTEMS										LOCAL						
-----										-----						
OUTBOUND FROM TRX1										INBOUND TO TRX1			TRX1			
-----										-----			-----			
TO SYSTEM	TRANSPORT CLASS	BUFFER LENGTH	REQ OUT	-----% SML	BUFFER % FIT	-----% BIG	% OVR	ALL PATHS UNAVAIL	REQ REJECT	FROM SYSTEM	REQ IN	REQ REJECT	TRANSPORT CLASS	REQ REJECT		
TRX2	BIG	40,892	92	100	0	0	0	0	0	TRX2	843	0	BIG	0		
	DB2	956	0	0	0	0	0	0	0				DB2	0		
	DEFAULT	956	720	0	100	0	0	0	0				DEFAULT	0		
	FEWFAST	956	0	0	0	0	0	0	0				FEWFAST	0		
	JES2	956	0	0	0	0	0	0	0				JES2	0		
	TCCONS	956	8	0	63	38	100	8	0				TCCONS	0		
	TCGRS	956	36	0	100	0	0	0	0				TCGRS	0		
	TCOPC	956	0	0	0	0	0	0	0				TCOPC	0		
	TCRMF	62,464	151	30	70	0	0	0	0				TCRMF	0		
	TCVLF	956	0	0	0	0	0	0	0				TCVLF	0		

TOTAL			1,007							TOTAL			843			

Figure 253. XCF Activity Report - Usage by System

Table 225. Fields in the XCF Activity Report - Usage by System	
Field Heading	Meaning
TO SYSTEM	The name of the system to which the messages were sent.
TRANSPORT CLASS	The name of the transport class used by XCF for the message transfer. There is one line in the report for each transport class for each target system. Note: If you request the XCF Activity report as a duration report, the Postprocessor inserts an * into this column, if the transport class has not been active during the complete duration interval.
BUFFER LENGTH	The internally defined message buffer size XCF uses for the transport class. The buffer size is the maximum length of the messages which can be contained in the buffers currently being used for this transport class.
REQ OUT	The total number of messages that XCF accepted for delivery to the system in the indicated transport class.
% SML	The percentage of messages sent that could have used a smaller BUFFER LENGTH than their defined BUFFER LENGTH.
% FIT	The percentage of messages sent that fit the defined BUFFER LENGTH.

Table 225. Fields in the XCF Activity Report - Usage by System (continued)	
Field Heading	Meaning
% 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.
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 the system sending the message. There is one line in the report for each system in the sysplex.
REQ IN	For each system, the total number of messages that were received from the remote system in the sysplex.
TRANSPORT CLASS	The name of the transport class that XCF uses for the message transfer. There is one line in the report for each transport class for each target system.

X C F A C T I V I T Y									
z/OS 3.1		SYSTEM ID TRX1		DATE 09/30/2021		INTERVAL 15.00.000		PAGE 2	
		RPT VERSION 3.1 RMF		TIME 09.15.00		CYCLE 1.000 SECONDS			
XCF USAGE BY MEMBER									
MEMBERS COMMUNICATING WITH TRX1					MEMBERS ON TRX1				
GROUP	MEMBER	SYSTEM	REQ FROM TRX1	REQ TO TRX1	GROUP	MEMBER	REQ OUT	REQ IN	
BOETRX1	BOETRX1\$TRX2	TRX2	199	172	BOETRX1	BOETRX1\$TRX1	199	172	
TOTAL			199	172	TOTAL		199	172	
COFVLFNO	TRX2	TRX2	0	0	COFVLFNO	TRX1	0	0	
TOTAL			0	0	TOTAL		0	0	
CTTXNGRP	CTTXN###TRX2	TRX2	0	0	CTTXNGRP	CTTXN###TRX1	0	0	
TOTAL			0	0	TOTAL		0	0	
ESCM	M245	TRX2	0	0					
TOTAL			0	0					
EZBTCPCS	TRX2TCPIP	TRX2	0	0	EZBTCPCS	TRX1TCPIP	0	0	
TOTAL			0	0	TOTAL		0	0	
					IDAVQUI0	IDAVQUI0TRX1	0	0	
					TOTAL		0	0	
					IGWXSGIS	N200711006204502	0	0	
					TOTAL		0	0	
ISTCFS01	IPSVX\$\$\$DEIBMIPS	TRX2	0	0	ISTCFS01	IPSVW\$\$\$DEIBMIPS	0	0	
TOTAL			0	0	TOTAL		0	0	
ISTXCF	IPSVX\$\$\$DEIBMIPS	TRX2	140	113	ISTXCF	IPSVW\$\$\$DEIBMIPS	140	113	
TOTAL			140	113	TOTAL		140	113	
					IXCL000F	M415	0	0	
					TOTAL		0	0	

Figure 254. XCF Activity Report - Usage by Member

Table 226. 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											
z/OS 3.1		SYSTEM ID RS21		DATE 09/30/2021		INTERVAL 15.00.00000		PAGE 6			
		RPT VERSION 3.1 RMF		TIME 09.15.00		CYCLE 1.000 SECONDS					
TOTAL SAMPLES = 899						XCF PATH STATISTICS					

OUTBOUND FROM RS21											

TO	T FROM/TO	TRANSPORT	REQ	AVG Q							
SYSTEM	Y DEVICE, OR	CLASS	OUT	LNTH	AVAIL	BUSY	RETRY				
RS22	S IXC_DEFAULT	_XCFMGD	2,847	0.00	2,847	0	0				
	S IXC_DEF8K	_XCFMGD	1,538	0.00	1,538	0	0				
RS23	S IXC_DEFAULT	DEFAULT	3,503	0.00	3,503	0	0				
	S IXC_DEF8K	DEF8K	1,013	0.00	361	652	0				
RS25	S IXC_DEFAULT	_XCFMGD	5,834	0.00	5,834	0	0				
	S IXC_DEF8K	_XCFMGD	1,024	0.00	1,024	0	0				
RS26	S IXC_DEFAULT	_XCFMGD	1,993	0.00	1,993	0	0				
	S IXC_DEF8K	_XCFMGD	1,031	0.00	1,031	0	0				
RS27	S IXC_DEFAULT	_XCFMGD	2,475	0.00	2,471	4	0				
	S IXC_DEF8K	_XCFMGD	879	0.00	871	8	0				
RS28	S IXC_DEFAULT	_XCFMGD	3,960	0.00	3,914	46	0				
	S IXC_DEF8K	_XCFMGD	1,577	0.00	1,533	44	0				
TOTAL			27,674								

INBOUND TO RS21											

FROM	T FROM/TO	REQ	BUFFERS	TRANSFER	AVG NO	----- USAGE -----					
SYSTEM	Y DEVICE, OR	IN	UNAVAIL	TIME	BUF TIME	UTIL%	TIME	IN USE	SIGNALS		
RS22	S IXC_DEFAULT	2,768	0	0.093	0.000	25	74.684	2,714	2,755		
						50	0.607	2	6		
						75	0.438	1	3		
						100	0.013	1	4		
	S IXC_DEF8K	1,272	0	0.065	0.000	25	88.844	702	971		
						50	5.552	56	138		
						75	2.815	24	75		
						100	4.640	13	88		
RS23	S IXC_DEFAULT	4,254	0	0.057	0.000	25	121.493	4,101	4,254		
						50	0.000	0	0		
						75	0.000	0	0		
						100	0.000	0	0		
	S IXC_DEF8K	4,858	0	0.075	0.000	25	89.708	1,886	2,239		
						50	16.065	333	704		
						75	20.021	449	1,398		
						100	12.643	130	592		
RS25	S IXC_DEFAULT	6,168	0	0.056	0.000	25	170.613	5,974	6,123		
						50	2.756	12	30		
						75	0.115	1	3		
						100	1.697	3	12		

Figure 255. XCF Activity Report - Path Statistics (Coupling Facility and Channel-to-Channel)

Table 227. Fields in the XCF Activity Report - XCF Path Statistics	
Field Heading	Meaning
OUTBOUND	
TO SYSTEM	The name of the system to which the messages are sent.
TYP	Path type C Channel-to-channel S Coupling facility structure

Table 227. Fields in the XCF Activity Report - XCF Path Statistics (continued)

Field Heading	Meaning
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 (for example, 1B62 TO 1C62 in Figure 255 on page 512). Coupling facility: The coupling facility structure name (for example, IXCPLEX_PATH1 in Figure 255 on page 512). UNK* TO appears in this field when one or more paths are unknown at the beginning of the interval.
TRANSPORT CLASS	The name of the transport class XCF uses for the message transfer. There is one line in the report for each transport class for each target system.
REQ OUT	The number of attempts made to send a message over the indicated outbound signalling path.
AVG Q LENGTH	The average number of messages queued for data transfer over each outbound signalling path. The calculation is: $\text{AVG Q LENGTH} = \frac{\text{\# Message Entries for this Device}}{\text{\# Samples}}$
AVAIL	The number of times the signalling path was selected while available to immediately transfer a message.
BUSY	The number of times XCF selected a signalling path while a message was already in the process of being transferred.
RETRY	The number of times XCF initialized the signalling path.
INBOUND	
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 LENGTH 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.
TRANSFER TIME	The average I/O transfer time in milliseconds for the most recently received signals. This metric is only available for inbound paths.
AVG NO BUF TIME	Average time (in milliseconds) this path had a no-inbound-buffer impact condition.
UTIL%	The percent utilization this report row represents. Typical values are 25, 50, 75, and 100.
IN USE TIME	Total time (in milliseconds) this path was in use at the indicated percent utilization.
IN USE	The number of times this path was in use at the indicated percent utilization.
SIGNALS	The number of signals received at the indicated percent utilization.

Exception report

An Exception report presents a summary of the values that exceeded installation-defined thresholds over a specific period of time.

How to request this report

Exception reports are generated from data contained in SMF records built during a Monitor I or a Monitor III data gatherer session. RMF compares the threshold values specified in the exception control statements with the computed value in the appropriate SMF record field. If the threshold is exceeded, RMF writes a line in the exception report.

You define the contents by one or more EXCEPT statements (see *z/OS Resource Measurement Facility User's Guide* for details), as follows:

```
EXCEPT(option)
```

Contents of the report

Figure 256 on page 514 and Figure 257 on page 515 show sample exception reports.

Example:

If you want to produce an Exception report to display all intervals between midnight and 8 a.m. (off-shift) when the utilization of one processor was equal to or below 3%, you can use the following report option statements:

```
ETOD(0000,0800)
EXCEPT(CPU(CPUBSY,LE,3))
EXCEPT(CPU0(CPUBSY(0),LE,3))
EXCEPT(CPU1(CPUBSY(1),LE,3))
EXCEPT(CPU2(CPUBSY(2),LE,3))
EXCEPT(CPU3(CPUBSY(3),LE,3))
EXCEPT(CPU4(CPUBSY(4),LE,3))
EXCEPT(CPU5(CPUBSY(5),LE,3))
```

Note: The sample report assumes a 6-way processor. The first EXCEPT statement reflects the average utilization for all processors. The other EXCEPT statements reflect the average utilization for the specified processor.

R M F E X C E P T I O N R E P O R T										PAGE	001
z/OS 3.1		SYSTEM ID		MVS1		START		09/30/2021-03.00.00		INTERVAL 00.30.00	
		RPT VERSION 3.1		RMF		END		09/30/2021-08.00.00		CYCLE 1.000 SECONDS	
NUMBER OF	INTERVALS 42			TOTAL LENGTH OF	INTERVALS 05.00.00						
CONTROL	DATE	TIME	INT	EXCEPTION	ACTUAL						
STMT NAME	MM/DD	HH.MM.SS	HH.MM.SS	THRESHOLD	VALUE	EXCEPTION DESCRIPTION	NAME VALUE				
CPU1	09/30	03.00.00	00.30.00	LE 3	2.453	PERCENT CPU BUSY	-1				
CPU2	09/30	03.00.00	00.30.00	LE 3	2.315	PERCENT CPU BUSY	-2				
CPU3	09/30	03.00.00	00.30.00	LE 3	2.301	PERCENT CPU BUSY	-3				
CPU4	09/30	03.00.00	00.30.00	LE 3	2.391	PERCENT CPU BUSY	-4				
CPU5	09/30	03.00.00	00.30.00	LE 3	2.858	PERCENT CPU BUSY	-5				
CPU1	09/30	03.30.00	00.30.00	LE 3	2.637	PERCENT CPU BUSY	-1				
CPU2	09/30	03.30.00	00.30.00	LE 3	2.502	PERCENT CPU BUSY	-2				
CPU3	09/30	03.30.00	00.30.00	LE 3	2.165	PERCENT CPU BUSY	-3				
CPU4	09/30	03.30.00	00.30.00	LE 3	2.279	PERCENT CPU BUSY	-4				
CPU1	09/30	04.00.00	00.30.00	LE 3	2.690	PERCENT CPU BUSY	-1				
CPU2	09/30	04.00.00	00.30.00	LE 3	2.233	PERCENT CPU BUSY	-2				
CPU3	09/30	04.00.00	00.30.00	LE 3	2.350	PERCENT CPU BUSY	-3				
CPU4	09/30	04.30.00	00.30.00	LE 3	2.191	PERCENT CPU BUSY	-4				
CPU1	09/30	04.30.00	00.30.00	LE 3	2.830	PERCENT CPU BUSY	-1				
CPU2	09/30	04.30.00	00.30.00	LE 3	2.345	PERCENT CPU BUSY	-2				
CPU3	09/30	04.30.00	00.30.00	LE 3	2.373	PERCENT CPU BUSY	-3				
CPU4	09/30	04.30.00	00.30.00	LE 3	2.274	PERCENT CPU BUSY	-4				
CPU1	09/30	05.00.00	00.30.00	LE 3	2.783	PERCENT CPU BUSY	-1				
CPU2	09/30	05.00.00	00.30.00	LE 3	2.505	PERCENT CPU BUSY	-2				
CPU3	09/30	05.00.00	00.30.00	LE 3	2.257	PERCENT CPU BUSY	-3				
CPU4	09/30	05.00.00	00.30.00	LE 3	2.405	PERCENT CPU BUSY	-4				
CPU4	09/30	05.30.00	00.30.00	LE 3	2.548	PERCENT CPU BUSY	-4				
CPU3	09/30	06.30.00	00.30.00	LE 3	2.811	PERCENT CPU BUSY	-3				
CPU4	09/30	06.30.00	00.30.00	LE 3	2.861	PERCENT CPU BUSY	-4				
CPU4	09/30	07.00.00	00.30.00	LE 3	2.988	PERCENT CPU BUSY	-4				

Figure 256. Exception Report - Low CPU Utilization

Example:

If you want to produce an Exception report to display all intervals between 3 a.m. and 6 a.m., when the utilization of one processor was greater than or equal to 10%, you can use the following report option statements:

```
ETOD(0300,0600)
EXCEPT(CPU0(CPUBSY(0),GE,10))
```

R M F E X C E P T I O N R E P O R T									
z/OS 3.1				SYSTEM ID MVS1		START 09/30/2021-03.00.00		INTERVAL 00.30.00	
				RPT VERSION 3.1 RMF		END 09/30/2021-06.00.00		CYCLE 1.000 SECONDS	
NUMBER OF		INTERVALS 6		TOTAL LENGTH OF		INTERVALS 03.00.00			
CONTROL	DATE	TIME		INT	EXCEPTION	ACTUAL			
STMT NAME	MM/DD	HH.MM.SS	HH.MM.SS	THRESHOLD	VALUE	EXCEPTION DESCRIPTION	NAME VALUE		
CPU0	09/30	03.00.00	00.30.00	GE 10	22.453	PERCENT CPU BUSY	-0		
CPU0	09/30	03.30.00	00.30.00	GE 10	22.637	PERCENT CPU BUSY	-0		
CPU0	09/30	04.00.00	00.30.00	GE 10	22.690	PERCENT CPU BUSY	-0		
CPU0	09/30	04.30.00	00.30.00	GE 10	22.830	PERCENT CPU BUSY	-0		
CPU0	09/30	05.00.00	00.30.00	GE 10	22.783	PERCENT CPU BUSY	-0		
CPU0	09/30	05.30.00	00.30.00	GE 10	21.383	PERCENT CPU BUSY	-0		

Figure 257. Exception Report - CPU Utilization

Heading fields

The heading fields for an Exception report identify the type of operating system, the release number and level of the operating system, the four-character SMF system ID of the system at system generation, and the RMF report level. The START field shows the date and time when the first interval in the reporting period began. The END field shows the date and time when the last interval ended. The date is in the form mm/dd/yy, and the time is in the form hh.mm.ss. The INTERVAL field shows the average length of the RMF measurement interval during the reporting period, in the form hh.mm.ss. The CYCLE field shows the length of the sampling cycle during the reporting period. When all SMF records have the same cycle length, that value is reported. When different cycle lengths are encountered, the Postprocessor sets the CYCLE field equal to the average of all cycle lengths encountered.

Note: When an Exception report consists of more than one page, the heading fields are repeated for each page. The START, END, CYCLE, and INTERVAL fields reflect the contents of the data in the entire report.

Data fields

Two fields precede the data fields. NUMBER OF INTERVALS indicates the number of RMF measurement intervals included in the reporting period. TOTAL LENGTH OF INTERVALS indicates (in the form hh.mm.ss) the total of the reporting period.

When an Exception report consists of more than one page, the NUMBER OF INTERVALS field and the TOTAL LENGTH OF INTERVALS field reflect the contents of the entire report.

The meaning of each field in the report is described in the following text.

Table 228. Fields in the Exception Report	
Field Heading	Meaning
CONTROL STMT NAME	The control statement name, if it 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:

RMF OVERVIEW REPORT												
z/OS 3.1		SYSTEM ID SYS1		START 09/30/2021-08.02.26		INTERVAL 00.29.59		PAGE 001				
		RPT VERSION 3.1 RMF		END 09/30/2021-18.02.26		CYCLE 1.000 SECONDS						
NUMBER OF INTERVALS 20		TOTAL LENGTH OF INTERVALS 09.59.51										
DATE	TIME	INT	TS0001IO	TS0002IO	TS0004IO	TS0009IO	TS0013IO	TS0001RT	TS0002RT	TS0004RT	TS0009RT	TS0013RT
MM/DD	HH.MM.SS	HH.MM.SS										
09/30	08.02.26	00.29.59					4.309				28.329	
09/30	09.02.26	00.29.59			18.592				26.368			
09/30	11.32.26	00.30.00									27.843	
09/30	12.02.26	00.29.59				6.458	9.962			33.949		
09/30	12.32.26	00.30.00									26.709	
09/30	13.02.26	00.30.00										29.427
09/30	14.02.26	00.30.00							32.073			26.733
09/30	14.32.26	00.29.59	4.826							34.448		
09/30	15.02.26	00.29.59		3.134								
09/30	15.32.26	00.30.00								28.152		
09/30	16.02.26	00.30.00		13.403								
09/30	16.32.26	00.30.00										26.077

Figure 258. Overview Report - Exception Version

The reporting range covers 20 intervals, but you see in [Figure 258 on page 516](#) that only intervals with at least one exception value are listed in the report.

Table 229. 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.

Table 229. Fields in the Overview Report (continued)	
Field Heading	Meaning
TIME HH.MM.SS	The starting time of the interval.
INT MM.SS	The length of the interval.
FFFFFFF	The field name is the exception-condition name of the corresponding EXCEPT statement. All columns appear in the report in the same order as the EXCEPT statements are given as Postprocessor input.

Overview report

You can use the Overview report for:

- Improved summary reporting - you can tailor the report according to your requirements
- Creating overview records - you can use the records as input for the RMF Spreadsheet Reporter or any other spreadsheet application

How to request this report

The scope of data that can be reported in the Overview report depends on the data being gathered by Monitor I gatherer sessions.

You define the contents by one or more OVW statements (see *z/OS Resource Measurement Facility User's Guide* for details), as follows:

OVW(option)

Note: For records based on single-system reports, the option EXCEPT is still valid, but it is recommended to use OVW in general.

To produce the report, specify

```
OVERVIEW(REPORT)
```

To create data records either for spreadsheet processing or other applications, specify:

```
OVERVIEW(RECORD)
```

This report is also available in XML output format. How to work with Postprocessor XML reports in *z/OS Resource Measurement Facility User's Guide* provides all required information on how to produce and view XML reports.

Example URL for the DDS API

```
http://ddshost:8803/gpm/rmfpp.xml?overview=(DATA01(CADSTG(SSID(0600),
DEVN(06F3))), (DB2PRD(CADRT(DEVN(0722),SSID(0700))),
(RHT0050(CASRHT(SSID(0050))))
```

Report description

Internally, the Overview report and the Exception report use the same technology. Therefore, the scope of data that can be reported is the same in both reports. You can find a list of all possible values in [Overview and exception conditions](#) in *z/OS Resource Measurement Facility 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 224](#) on page 506.

You cannot specify generic exception classes such as DASD without a qualifier. This is valid for exception reporting, but not for overview reporting. Here, you have to specify explicitly a qualifier, which for DASD could be either a device address or a volume serial number.

Due to the above described technology, you have to use the ETOD statement if you want to specify explicitly the time range for the Overview report.

Improved summary reporting

The Summary report provides performance data that summarize system activity for each interval within the reporting period (see [“Summary report” on page 521](#)). The contents of the report cannot be modified.

The Overview report allows you to select the performance data you want to be 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:

RMF OVERVIEW REPORT											PAGE 001
z/OS 3.1			SYSPLX ID UTCPLXHD		START 09/30/2021-10.00.00		INTERVAL 00.10.00				
			RPT VERSION 3.1 RMF		END 09/30/2021-14.00.02						
NUMBER OF INTERVALS 24			TOTAL LENGTH OF INTERVALS 04.00.02								
DATE	TIME	INT	SERVUNIT	RTIMEP1	RTIMEP2	RTIMEP3	PIP1	PIP2	TRXP1	TRXP2	TRXP3
MM/DD	HH.MM.SS	HH.MM.SS									
09/30	10.00.00	00.10.01	28406	0.229	14.491	18.916	0.763	2.415	12.40	0.82	0.10
09/30	10.10.00	00.10.00	32696	0.264	6.071	20.415	0.880	0.812	15.52	1.00	0.09
09/30	10.20.00	00.10.02	34245	0.261	4.425	17.877	0.870	0.738	15.66	1.02	0.11
09/30	10.30.00	00.10.00	34943	0.270	5.695	19.504	0.900	0.949	15.03	1.02	0.12
09/30	10.40.00	00.10.00	34830	0.313	12.344	20.806	1.043	2.057	15.29	1.02	0.11
09/30	10.50.00	00.10.01	18360	0.373	6.619	23.352	1.243	1.103	9.34	0.54	0.06
09/30	11.00.00	00.10.03	29893	3.928	25.872	57.726	13.093	4.312	12.47	0.90	0.10
09/30	11.10.00	00.10.00	35164	0.251	8.943	15.304	0.837	1.491	15.48	1.00	0.13
09/30	11.20.00	00.10.00	33544	0.236	4.563	17.671	0.787	0.761	15.50	1.03	0.11
09/30	11.30.00	00.10.00	35057	0.257	5.829	13.389	0.857	0.972	15.73	1.01	0.11
09/30	11.40.00	00.10.01	35811	0.252	10.049	15.257	0.840	1.675	15.40	1.07	0.13
09/30	11.50.00	00.10.00	35419	0.248	4.507	17.461	0.827	0.751	15.24	1.03	0.13
09/30	12.00.00	00.10.01	35902	0.437	6.883	18.944	1.457	1.147	15.12	1.02	0.13
09/30	12.10.00	00.10.00	36967	0.247	9.635	16.407	0.823	1.606	16.10	1.06	0.12
09/30	12.20.00	00.10.02	36024	0.260	4.552	18.229	0.867	0.759	15.15	1.02	0.13
09/30	12.30.00	00.10.00	36296	0.263	5.072	20.555	0.877	0.845	15.29	1.06	0.12
09/30	12.40.00	00.10.00	35129	0.262	10.237	16.135	0.873	1.706	15.37	1.02	0.12
09/30	12.50.00	00.10.00	35355	0.274	8.098	15.106	0.913	1.350	15.31	1.06	0.12
09/30	13.00.00	00.10.00	36936	0.213	3.833	10.036	0.710	0.639	16.00	1.09	0.12
09/30	13.10.00	00.10.00	36919	0.182	5.205	9.323	0.607	0.868	16.17	1.08	0.13
09/30	13.20.00	00.10.00	31098	0.208	5.420	9.599	0.693	0.903	14.14	0.87	0.12
09/30	13.29.59	00.10.01	34909	1.184	13.976	25.991	3.947	2.329	14.20	1.00	0.13
09/30	13.40.00	00.10.00	34807	0.234	3.865	11.235	0.780	0.644	15.75	1.03	0.11
09/30	13.50.00	00.10.02	31773	0.227	8.372	12.955	0.757	1.395	14.00	0.93	0.11

Figure 259. Overview Report - Summary Version

Creating Overview records

You can also create records for further processing with the Spreadsheet Reporter or other applications either on the host system or on your workstation.

The Spreadsheet Reporter provides full support for converting SMF dump data, Postprocessor listings and Overview records into spreadsheets. You can use it to create and submit Postprocessor jobs directly on the workstation without a logon to the host system, and you will receive the data in the correct format back to the workstation. In addition, it provides sample spreadsheets to help you in presenting and analyzing performance data at a glance. You can find a detailed description in *z/OS Resource Measurement Facility User's Guide*.

You get one record for each reported interval with the same information as in the printed *logical* line (this can be several *physical* lines on several pages if you define more than 11 exceptions) by specifying OVERVIEW(RECORD) .

A record can contain a maximum of 253 exceptions.

If you want to get both the report and the records, you can combine both control statements into OVERVIEW(REPORT , RECORD) .

It is recommended to use this version of the OVERVIEW statement, it provides the capability to check whether you really get the data that you expect.

For each report, the Postprocessor creates one Overview Header record and several (one for each interval) Overview Data records. If the input data for the Postprocessor consists of records for several systems, you get a set of records for each system.

The record mapping macro for all Overview records is ERBOVREC.

Overview header record

Table 230. 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 231. Overview Header Record - Header Section. One per record.				
Offsets	Format	Length	Name	Description
0	0 binary	2	OVRLN	Record length (for records on the workstation)
2	2 binary	2	OVRSGMT	

Table 231. Overview Header Record - Header Section. One per record. (continued)					
Offsets		Format	Length	Name	Description
4	4	EBCDIC	8	OVRTYPE	RMFOVREC (eye catcher)
12	C	EBCDIC	1	OVRVLV	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: ZVvrrmm (z/OS)
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	OVRCLGT	Total length for one SYSID: length of header record + (length of data records * number of data records) This length does not include the prefix section of the records.
64	40	binary	4	OVRCOLO	Offset to first report column names section
68	44	binary	4	OVRCOLN	Number of report column names sections
72	48	binary	4	OVRCOLL	Length of one report column names section
78	4C	binary	4	OVRHDRN	Overview header record counter
80	50	binary	4	OVRDATRN	Number of data records
84	54	EBCDIC	8	OVRSPID	Sysplex Id (for sysplex records)

Table 232. Overview Header Record - Report Column Names Section. One per interval.					
Offsets		Format	Length	Name	Description
0	0	EBCDIC	8	OVRNAME	Report column name
8	8	EBCDIC	8	OVRCOND	OVERVIEW condition name
16	10	EBCDIC	53	OVRQUAL	OVERVIEW condition qualifier
69	45	EBCDIC	1	*	Reserved
70	46	EBCDIC	2	OVROPER	Exception operator
72	48	EBCDIC	9	OVRTHV	Exception threshold value
81	51	EBCDIC	1	*	Reserved

Overview data record

Table 233. 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 234. Overview Data Record - Data Section. One per record.					
Offsets		Format	Length	Name	Description
0	0	binary	2	OVRDLEN	Record length (for records on the workstation)
2	2	binary	2	OVRDSGMT	Zero
4	4	binary	4	OVRDATO	Offset to first report data section
8	8	binary	4	OVRDATN	Number of report data sections
12	C	binary	4	OVRDATL	Length of one report data section
16	10	EBCDIC	5	OVRRIDAT	Reporting interval date: MM/DD
21	15	EBCDIC	1	*	Reserved
22	16	EBCDIC	8	OVRRITME	Reporting interval time hh.mm.ss
30	1E	EBCDIC	1	*	Reserved
31	1F	EBCDIC	8	OVRRIINT	Reporting interval length: hh.mm.ss
39	27	EBCDIC	1	*	Reserved

Table 235. Overview Data Record - Report Data Section. One per exception.					
Offsets		Format	Length	Name	Description
0	0	EBCDIC	1	*	Reserved
1	1	EBCDIC	9	OVRREAL	Report actual data

Summary report

Summary reports provide a high-level view of system activity. A summary report can consist of one line of data that summarizes system activity for each interval within the reporting period, a single line of data giving summary totals for all RMF intervals within the reporting period, or both interval summary data lines and a total summary data line. Figure 260 on page 522 shows a sample summary report, including both an interval summary line for each RMF measurement interval and a total summary line for all of the intervals.

How to request this report

To produce this report, specify

```
SUMMARY(INT | TOT)
```

Generating a report

The data shown in a Summary report is derived from the SMF records created by any Monitor I sessions that ran during the reporting period. You specify the type of reporting required on Postprocessor control statements. For a description of Postprocessor control statements, see *z/OS Resource Measurement Facility 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 311](#) 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.

R M F S U M M A R Y R E P O R T															PAGE 001		
z/OS 3.1			SYSTEM ID MVS1			START 09/30/2021-08.02.26			INTERVAL 00.29.59								
			RPT VERSION 3.1 RMF			END 09/30/2021-12.02.26			CYCLE 1.000 SECONDS								
NUMBER OF INTERVALS 8																	
DATE	TIME	INT	CPU	DASD	TOTAL	LENGTH	OF	INTERVALS	03.59.56	STC	ASCH	ASCH	OMVS	OMVS	SWAP	DEMAND	
MM/DD	HH.MM.SS	MM.SS	BUSY	RESP	RATE	JOB	JOB	TSO	TSO	MAX	AVE	MAX	AVE	MAX	AVE	RATE	PAGING
09/30	08.02.26	29.59	22.6	12	465.9	4	2	64	51	76	73	0	0	0	0	0.00	0.00
09/30	08.32.26	30.00	28.9	17	642.2	7	4	89	77	74	71	0	0	0	0	0.00	0.00
09/30	09.02.26	29.59	26.0	19	566.7	8	5	103	97	73	71	0	0	0	0	0.00	0.00
09/30	09.32.26	30.00	27.8	12	583.5	7	4	103	100	75	70	0	0	0	0	0.00	0.00
09/30	10.02.26	29.59	25.5	20	542.6	8	6	107	102	70	68	0	0	0	0	0.00	0.00
09/30	10.32.26	29.59	43.7	11	545.3	7	4	106	99	73	70	0	0	0	0	0.00	0.00
09/30	11.02.26	30.00	48.6	13	470.1	10	6	98	94	72	69	0	0	0	0	0.00	0.00
09/30	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 260. 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 236. 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 general purpose processors. Special purpose processors are not included in the calculation. For systems running in a PR/SM environment, this value is the LPAR busy time percentage.	
DASD RESP	The average number of milliseconds required to complete an I/O request on all direct access storage devices included in the report.	
DASD RATE	The activity per second for all direct access storage devices included in the report. The value reported corresponds to an accumulation of each DEVICE ACTIVITY RATE field in the Direct Access Device Activity report.	
TAPE RATE	The activity per second for all magnetic tape devices included in the report. The value reported corresponds to an accumulation of each DEVICE ACTIVITY RATE field in the Magnetic Tape Device Activity report.	
JOB MAX/AVE	The maximum and average number of batch jobs that were active during each measurement interval. The values reported corresponds to the MAX/AVE number of BATCH address spaces in the CPU Activity report.	
TSO MAX/AVE	The maximum and average number of TSO/E sessions that were active during each measurement interval. The values reported corresponds to the MAX/AVE number of TSO/E address spaces in the CPU Activity report.	
STC MAX/AVE	The maximum and average number of started tasks and mount tasks that were active during each measurement interval. The value reported corresponds to the MAX/AVE number of STC address spaces in the CPU Activity report.	
ASCH MAX/AVE	The maximum/average number of APPC/MVS transaction scheduler (ASCH) address spaces that were active during each measurement interval. The value reported corresponds to the MAX/AVE number of ASCH address spaces in the CPU Activity report.	

Table 236. Fields in the Summary Report (continued)	
Field Heading	Meaning
OMVS MAX/AVE	The maximum/average number of OMVS address spaces that were active during each measurement interval. The value reported corresponds to the MAX/AVE number of OMVS address spaces in the CPU Activity report.
SWAP RATE	The number of swaps per second for each interval. The value reported corresponds to the sum of the AUX STOR TOTAL and the EXP STOR TOTAL fields in the SWAP PLACEMENT ACTIVITY section of the Monitor I Paging Activity report.
DEMAND PAGING	The number of demand paging requests per second for each interval. This is the demand paging rate from DASD (page fault rate).
TOTAL/ AVERAGE	The single line that reports total summary data. The line contains either the average rate of events over the reporting period (or page), or the maximum number of events during any of the measurement intervals included in the reporting period (or page).

Spreadsheet reference

You can make this report available in a spreadsheet, using the Spreadsheet Reporter. For details, see *z/OS Resource Measurement Facility User's Guide*.

Appendix A. Accessibility

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

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

Notices

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

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

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

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

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

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

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

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

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

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

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

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

*IBM Corporation
Site Counsel
2455 South Road*

Poughkeepsie, NY 12601-5400
USA

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Terms and conditions for product documentation

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

Applicability

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

Personal use

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

Commercial use

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

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

Rights

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

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

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

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

IBM Online Privacy Statement

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

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

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

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

Policy for unsupported hardware

Various z/OS elements, such as DFSMSdfp, JES2, JES3, and MVS, contain code that supports specific hardware servers or devices. In some cases, this device-related element support remains in the product even after the hardware devices pass their announced End of Service date. z/OS may continue to service element code; however, it will not provide service related to unsupported hardware devices. Software problems related to these devices will not be accepted for service, and current service activity will cease if a problem is determined to be associated with out-of-support devices. In such cases, fixes will not be issued.

Minimum supported hardware

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

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

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

Programming interface information

This book is intended to help the customer to use RMF reports, and contains a detailed description of all reports.

This book documents intended Programming Interfaces that allow the customer to write programs to obtain the services of RMF. This information is identified where it occurs, either by an introductory statement to a topic or section or by this marking:

Programming Interface Information
End Programming Interface Information

Trademarks

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

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

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

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

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

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

Glossary

This glossary contains chiefly definitions of terms used in this book, but some more general RMF and MVS terms are also defined.

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

APPC/MVS

Advanced program-to-program communication

ASCH address space

APPC transaction scheduler address space

AS

Address space

address space

That part of MVS main storage that is allocated to a job.

auxiliary storage (AUX)

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

background session

In RMF, a monitor session that is started and controlled from the operator console. Contrast with *interactive session*

balanced systems

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

basic mode

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

bottleneck

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

callable services

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

captured storage

See shared page group.

capture ratio

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

central processor (CP)

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

central processor complex (CPC)

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

channel path

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

CICS

Customer Information Control System

CIM provider

A CIM provider is the link between the CIM server and the system interfaces. It allows the CIM server to access and manage the resources. Each CIM provider exposes the resources it represents in a standard way, using a small number of classes from the CIM schema or derived from the CIM schema. RMF monitoring providers are CIM providers implemented by RMF.

contention

Two or more incompatible requests for the same resource. For example, contention occurs if a user requests a resource and specifies exclusive use, and another user requests the same resource, but specifies shared use.

coupling facility

See *Cross-system Extended Services/Coupling Facility*.

CP

Central processor

criteria

Performance criteria set in the WFEX report options. You can set criteria for all report classes (PROC, SYSTEM, TSO, and so on).

CPU speed

Measurement of how much work your CPU can do in a certain amount of time.

cross-system coupling facility (XCF)

A component of MVS that provides functions to support cooperation between authorized programs running within a *sysplex*.

Cross-system Extended Services/Coupling Facility (XES/CF)

Provides services for MVS systems in a *sysplex* to share data on a coupling facility (CF).

CS

Central storage

Customer Information Control System (CICS)

An IBM licensed program that enables transactions entered at remote terminals to be processed concurrently by user-written application programs. It includes facilities for building, using, and maintaining data bases.

cycle

In RMF, the time at the end of which one sample is taken. Varies between 50 ms and 9999 ms. See also *sample*.

data sample

See *sample*

DCM

See *Dynamic Channel Path Management*

delay

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

direct access storage device (DASD)

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

DLY

Delay

DP

Dispatching priority

dynamic channel path management

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

EMIF

ESCON multiple image facility

enclave

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

EPDM

Enterprise Performance Data Manager/MVS

execution velocity

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

exception reporting

In RMF, the reporting of performance measurements that do not meet user-defined criteria. Shows potential performance problems explicitly, thus avoiding the need for constant monitoring.

generalized trace facility (GTF)

A service program that records significant system events, such as supervisor calls and start I/O operations, for the purpose of problem determination.

GO mode

In RMF, the Monitor III mode in which the screen is updated with the interval you specified in your session options. The terminal cannot be used for anything else when it is in GO mode. See also *mode*.

graphic mode

In RMF Monitor III, the mode which presents the performance data from the system in graphic format using the GDDM product. Contrast with *tabular mode*.

GTF

generalized trace facility

high-speed buffer (HSB)

A cache or a set of logically partitioned blocks that provides significantly faster access to instructions and data than provided by central storage.

HS

hiperspace

HSB

High-speed buffer

HSM

Hierarchical Storage Manager

IBM Z Application Assist Processor (zAAP)

A special purpose processor configured for running Java programming on selected zSeries machines.

IBM Z Integrated Information Processor (zIIP)

A special purpose processor designed to help free-up general computing capacity and lower overall total cost of computing for selected data and transaction processing workloads for business intelligence (BI), ERP and CRM, and selected network encryption workloads on the mainframe.

IMS

Information Management System

Information Management System (IMS)

A database/data communication (DB/DC) system that can manage complex databases and networks. Synonymous with IMS/VS.

interactive session

In RMF, a monitor display-session that is controlled from the display terminal. Contrast with *background session*.

JES

Job Entry Subsystem

LCU

Logical control unit. Logical control units are also called 'Control Unit Headers' (CUH). For details about LCU/CUH please refer to the applicable *IBM Z Input/Output Configuration Program User's Guide for ICP IOCP* (SB10-7037).

logically partitioned (LPAR) mode

A central processor mode that is available on the Configuration frame when using the PR/SM feature. It allows an operator to allocate processor unit hardware resources among logical partitions. Contrast with *basic mode*.

logical partition (LP)

A subset of the processor hardware that is defined to support an operating system. See also *logically partitioned (LPAR) mode*.

LP

Logical partition

LPAR

Logically partitioned (mode)

LPAR cluster

An LPAR cluster is the subset of the systems that are running as LPARs on the same CEC. Based on business goals, WLM can direct PR/SM to enable or disable CP capacity for an LPAR, without human intervention.

migration rate

The rate (pages/second) of pages being moved from expanded storage through central storage to auxiliary storage.

mintime

The smallest unit of sampling in Monitor III. Specifies a time interval during which the system is sampled. The data gatherer combines all samples gathered into a set of samples. The set of samples can be summarized and reported by the reporter.

mode

Monitor III can run in various modes: GO mode (see *GO mode*) and STOP mode, which is the default mode. See also *graphic mode* and *tabular mode*.

MPL

Multiprogramming level

OMVS

Reference to z/OS UNIX System Services

partitioned data set (PDS)

A data set in direct access storage that is divided into partitions, called members, each of which can contain a program, part of a program, or data.

PDS

partitioned data set

performance management

The activity which monitors and allocates data processing resources to applications according to goals defined in a service level agreement or other objectives.

The discipline that encompasses collection of performance data and tuning of resources.

PR/SM

Processor Resource/Systems Manager

Processor Resource/Systems Manager (PR/SM)

The feature that allows the processor to run several operating systems environments simultaneously and provides logical partitioning capability. See also *LPAR*.

range

The time interval you choose for your report.

Resident time

The time the address space was swapped in, in units of seconds.

RMF monitoring provider

see CIM provider

sample

Once in every cycle, the number of jobs waiting for a resource, and what job is using the resource at that moment, are gathered for all resources of a system by Monitor III. These numbers constitute one sample.

SCP

System control program

seek

The DASD arm movement to a cylinder. A seek can range from the minimum to the maximum seek time of a device. In addition, some I/O operations involve multiple imbedded seeks where the total seek time can be more than the maximum device seek time.

service class

In Workload Manager, a subdivision of a *workload*. Performance goals and capacity boundaries are assigned to service classes.

service level agreement (SLA)

A written agreement of the information systems (I/S) service to be provided to the users of a computing installation.

Service Level Reporter (SLR)

An IBM licensed program that provides the user with a coordinated set of tools and techniques and consistent information to help manage the data processing installation. For example, SLR extracts information from SMF, IMS, and CICS logs, formats selected information into tabular or graphic reports, and gives assistance in maintaining database tables.

service rate

In the system resources manager, a measure of the rate at which system resources (services) are provided to individual jobs. It is used by the installation to specify performance objectives, and used by the workload manager to track the progress of individual jobs. Service is a linear combination of processing unit, I/O, and main storage measures that can be adjusted by the installation.

shared page groups

An address space can decide to share its storage with other address spaces using a function of RSM. As soon as other address spaces use these storage areas, they can no longer be tied to only one address space. These storage areas then reside as *shared page groups* in the system. The pages of shared page groups can reside in central, expanded, or auxiliary storage.

SLA

service level agreement

SLIP

serviceability level indication processing

SLR

Service Level Reporter

SMF

System management facility

SMF buffer

A wrap-around buffer area in storage, to which the z/OS Data Gatherer writes performance data, and from which the Postprocessor extracts data for reports.

speed

See *workflow*

SRB

Service request block

SRM

System resource manager

SSCH

Start subchannel

system control program (SCP)

Programming that is fundamental to the operation of the system. SCPs include MVS, VM, and VSE operating systems and any other programming that is used to operate and maintain the system. Synonymous with *operating system*.

sysplex

A complex consisting of a number of coupled MVS systems.

tabular mode

In RMF, the mode in which Monitor III displays performance data in the form of lists. Contrast with *graphic mode*.

TCB

Task control block

threshold

The exception criteria defined on the report options screen.

throughput

A measure of the amount of work performed by a computer system over a period of time, for example, number of jobs per day.

TPNS

Teleprocessing network simulator

TSO

Time Sharing Option, see *Time Sharing Option/Extensions*

Time Sharing Option Extensions (TSO/E)

In MVS, a time-sharing system accessed from a terminal that allows user access to MVS system services and interactive facilities.

UIC

Unreferenced interval count

uncaptured time

CPU time not allocated to a specific address space.

using

Jobs getting service from hardware resources (PROC or DEV) are *using* these resources.

velocity

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

VTOC

Volume table of contents

workflow

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

The workflow of resources indicates how efficiently users are being served.

workload

A logical group of work to be tracked, managed, and reported as a unit. Also, a logical group of service classes.

WLM

Workload Manager

XCF

Cross-system coupling facility

XES/CF

See *Cross-system Extended Services/Coupling Facility*.

zAAP

see IBM Z Application Assist Processor.

zIIP

see IBM Z Integrated Information Processor.

Index

Special Characters

- % CU BUSY (control unit busy %) field
 - in Monitor III IOQUEUE report [117](#)
- % Delayed for (percent delayed for) field
 - in Monitor III DELAY report [69](#)
- % DP BUSY (director port busy %) field
 - in Monitor III IOQUEUE report [117](#)
- % Frames (percentage frames) field
 - in Monitor III STORR report [174](#)
- % of Trx (percentage of transactions) field
 - in Monitor III SYSRTD report [192](#)

A

- abnormally ended transactions field
 - in Monitor III SYSWKM report [209](#)
- accessibility
 - contact IBM [525](#)
- action field
 - in Monitor III WFEX report options action panel [222](#)
- activated at field
 - in Monitor III SYSSUM report [197](#)
- active frames field
 - in Monitor III STORF report [169](#)
- active percentage field
 - in Monitor III DEVJ report (Active) [123](#)
- active policy field
 - in Monitor III SYSSUM report [197](#)
- activity rate field
 - in the Monitor III DEVN report (Rate) [77](#)
- actual field
 - in Monitor III SYSWKM report [208](#)
- address space identifier (ASID) field
 - in Monitor III STORC report [165](#)
- Address Space Resource Data (ARD/ARDJ)
 - Monitor II report
 - command [256](#)
 - contents of [257](#)
 - field descriptions [257](#)
 - how to request [256](#)
 - purpose of [256](#)
- Address Space SRM Data (ASRM/ASRMJ)
 - Monitor II report
 - command [264](#)
 - contents of [265](#)
 - how to request [264](#)
 - purpose of [264](#)
- Address Space State Data (ASD/ASDJ)
 - Monitor II report
 - command [260](#)
 - contents of [260](#)
 - how to request [260](#)
 - purpose of [260](#)
- address spaces serving this service class field
 - in Monitor III SYSWKM report [211](#)
- Advanced Encryption Standard (AES) [369](#), [370](#), [373](#)
- AES (Advanced Encryption Standard) [369](#), [370](#), [373](#)
- alert field
 - in Monitor III WFEX report definition and criteria panel [225](#)
- amount of common storage not released field
 - in Monitor III STORCR report [167](#)
- amount used field
 - in Monitor III STORC report [165](#)
- AP [502](#)
- APPL% [494](#)
- Application Assist Processor, *See* IBM System z Application Assist Processor (zAAP)
- ARD/ARDJ (Address Space Resource Data)
 - Monitor II report
 - command [256](#)
 - contents of [257](#)
 - field descriptions [257](#)
 - how to request [256](#)
 - purpose of [256](#)
- ASD/ASDJ (Address Space State Data)
 - Monitor II report
 - command [260](#)
 - contents of [260](#)
 - how to request [260](#)
 - purpose of [260](#)
- ASID (address space identifier) field
 - in Monitor III STORC report [165](#)
 - in Monitor III SYSWKM report [211](#)
- ASM [431](#)
- ASRM/ASRMJ (Address Space SRM Data)
 - Monitor II report
 - command [264](#)
 - contents of [265](#)
 - how to request [264](#)
 - purpose of [264](#)
- assistive technologies [525](#)
- automatic customization
 - where to specify [230](#)
 - workflow/exception (WFEX) options [217](#)
- auxiliary slots field
 - in Monitor III STORF report [169](#)
- Auxiliary Storage Manager [431](#)
- available at end of range field
 - in Monitor III STORC report [164](#)
- average active users field
 - in Monitor III STORR report [176](#)
- average CPU utilization field
 - in Monitor III SYSINFO report (average CPU UTIL) [183](#)
 - in Monitor III WFEX report (Average CPU Util) [218](#)
- average CPU Utilization field

- average CPU Utilization field (*continued*)
 - in the Monitor III SYSTREND report (CPU %) [204](#)
- average CSA to SQA conversion field
 - in Monitor III STORC report [163](#)
- average execution time (Avg. Exec. Time) field
 - in Monitor III SYSWKM report [209](#)
- average number delayed for field
 - in Monitor III STORS report [178](#)
 - in Monitor III SYSINFO report [185](#)
- average response time (Avg. Resp. Time) field
 - in Monitor III SYSRTD report [193](#)
 - in Monitor III SYSSUM report (active time) [201](#)
 - in Monitor III SYSSUM report (queue time) [201](#)
 - in Monitor III SYSSUM report (total time) [201](#)
 - in Monitor III SYSWKM report [208](#)
- average SRB time field
 - in the Monitor III SYSTREND report [204](#)
- average TCB time field
 - in the Monitor III SYSTREND report [204](#)
- average use summary field
 - in Monitor III STORC report [164](#)
- average users delayed field
 - in Monitor III DEVJ report [123](#)

B

- blocked workload
 - threshold for promotion [356](#)
- buffer counts by pool [152](#)

C

- C (class) field
 - in Monitor III STOR report [160](#)
 - in Monitor III STORF report [168](#)
 - in Monitor III WFEX report options action panel [222](#)
- CACHDET (cache activity details)
 - how to request [32](#)
 - Monitor III report
 - field descriptions [33](#)
- CACHDET (Cache Detail)
 - Monitor III report
 - command abbreviation [19](#)
 - purpose of [32](#)
- CACHE (cache subsystem activity)
 - Postprocessor report
 - Subsystem Summary report [314](#)
- CACHE (Cache Subsystem Activity)
 - Postprocessor report
 - command [312](#)
 - device-level reporting [320](#)
 - how to request [312](#)
 - purpose of [312](#)
 - report levels [313](#)
 - spreadsheet and overview reference [321](#)
 - Subsystem Device Overview report [318](#)
 - Subsystem Status report [315](#)
 - subsystem summary report [313](#)
 - subsystem-level reporting [314](#)
 - top-20 device lists [313](#)
- cache activity details
 - Monitor III report (CACHDET)
 - field descriptions [33](#)

- cache activity summary
 - Monitor III report (CACHSUM)
 - field descriptions [37](#)
- Cache Detail
 - Monitor III report
 - command abbreviation [19](#)
 - purpose of [32](#)
- Cache Detail Report
 - how to request [32](#)
- Cache Subsystem Activity (CACHE)
 - Postprocessor report
 - command [312](#)
 - device-level reporting [320](#)
 - how to request [312](#)
 - purpose of [312](#)
 - report levels [313](#)
 - spreadsheet and overview reference [321](#)
 - Subsystem Device Overview report [318](#)
 - Subsystem Overview report [315](#)
 - subsystem summary report [313](#)
 - subsystem-level reporting [314](#)
 - top-20 device lists [313](#)
- Cache Summary (CACHSUM)
 - Monitor III report
 - command abbreviation [19](#)
 - purpose of [36](#)
- CACHSUM (cache activity summary)
 - Monitor III report
 - field descriptions [37](#)
 - how to request [36](#)
- CACHSUM (Cache Summary)
 - Monitor III report
 - command abbreviation [19](#)
 - purpose of [36](#)
- CAP 502
 - capacity field
 - in Monitor III SYSSUM report (actual) [202](#)
 - in Monitor III SYSSUM report (max) [202](#)
 - in Monitor III SYSSUM report (min) [202](#)
- Capacity report
 - contents and field descriptions [57](#)
- capping delay (Capp) field
 - in Monitor III SYSWKM report [211](#)
- central storage frames field
 - PAGING report [438](#)
- central storage movement rates [437](#)
- central storage request rates [437](#)
- Central Storage/Processor/SRM (SRCS)
 - Monitor II report
 - command [301](#)
 - contents of [302](#)
 - field descriptions [302](#)
 - how to request [301](#)
 - purpose of [301](#)
- CF (Coupling Facility Activity)
 - Postprocessor report
 - activity section [332](#)
 - CF to CF activity section [339](#)
 - command [324](#)
 - contents of [324](#)
 - how to request [324](#)
 - purpose of [324](#)
 - spreadsheet and overview reference [340](#)

- CF (Coupling Facility Activity) (*continued*)
 - Postprocessor report (*continued*)
 - subchannel activity section [336](#)
 - usage summary section [325](#)
- CF Systems report (Monitor III CFSYS report) [19](#), [49](#)
- CFACT (coupling facility activity)
 - Monitor III report
 - field descriptions [42](#)
 - report options panel [46](#)
 - sample [40](#)
- CFACT (Coupling Facility Activity)
 - Monitor III report
 - command [39](#)
 - command abbreviation [19](#)
 - contents of [40](#)
 - how to request [39](#)
 - purpose of [39](#)
 - special considerations [39](#)
- CFOVER (Coupling Facility Overview)
 - Monitor III report
 - command [47](#)
 - command abbreviation [19](#)
 - how to request [47](#)
 - purpose of [47](#)
- CFOVER Report
 - contents of [47](#)
- CFSYS (coupling facility systems)
 - Monitor III report
 - field descriptions [50](#)
- CFSYS (Coupling Facility Systems)
 - Monitor III report
 - command [49](#)
 - command abbreviation [19](#)
 - how to request [49](#)
 - purpose of [49](#)
- CHAN (Channel Path Activity)
 - Postprocessor report
 - command [342](#)
 - contents of [342](#)
 - dynamic channel path management (DCM) [341](#)
 - how to request [342](#)
 - purpose of [341](#)
 - spreadsheet and overview reference [346](#)
- CHANNEL (Channel Path Activity)
 - Monitor II report
 - command [267](#)
 - how to request [267](#)
 - purpose of [267](#)
 - Monitor III report
 - command [53](#)
 - command abbreviation [19](#)
 - customizing with Monitor III Utility [55](#)
 - how to request [53](#)
 - purpose of [52](#)
- Channel Path
 - generation [54](#)
- Channel Path Activity (CHAN)
 - Postprocessor report
 - command [342](#)
 - contents of [342](#)
 - dynamic channel path management (DCM) [341](#)
 - how to request [342](#)
 - purpose of [341](#)
 - spreadsheet and overview reference [346](#)

- Channel Path Activity (CHANNEL)
 - Monitor II report
 - command [267](#)
 - how to request [267](#)
 - purpose of [267](#)
 - Monitor III report
 - command abbreviation [19](#)
 - customizing with Monitor III Utility [55](#)
 - purpose of [52](#)
- Channel Path Activity report (Monitor III CHANNEL report) [52](#)
- channel paths field
 - in Monitor III IOQUEUE report [116](#)
- CHPID taken field
 - in Monitor III IOQUEUE report [117](#)
- class (C) field
 - in Monitor III STOR report [160](#)
 - in Monitor III STORF report [168](#)
 - in Monitor III WFEX report options action panel [222](#)
- class field
 - in Monitor III GROUP report [106](#)
 - in Monitor III SYSWKM report [208](#)
- CML and LOCAL lock [456](#)
- CMS/CMSEQDQ/CMSLatch/CMSSMF locks [456](#)
- common storage
 - Monitor III report (STORC)
 - field descriptions - system information [164](#)
- Common Storage (STORC)
 - Monitor III report
 - command [162](#)
 - how to request [162](#)
 - purpose of [162](#)
- common storage remaining
 - Monitor III report (STORCR)
 - field descriptions [167](#)
- Common Storage Remaining (STORCR)
 - Monitor III report
 - command [166](#)
 - how to request [166](#)
 - purpose of [166](#)
- Communication equipment Activity (DEVICE)
 - Postprocessor report
 - contents of [380](#)
 - pending time overflow [380](#)
- Concepts, Monitor III
 - shortened intervals, reason for [15](#)
- connect time percentage field
 - in Monitor III DEV report (CON %) [75](#)
 - in Monitor III DEVJ report [123](#)
- contact
 - z/OS [525](#)
- contention rate field
 - in Monitor III IOQUEUE report [116](#)
- control unit busy field
 - in Monitor III IOQUEUE report (% CU BUSY) [117](#)
- control unit field
 - in Monitor III IOQUEUE report [116](#)
 - in Monitor III STORR report [174](#)
- Coupling facility
 - Postprocessor report
 - subchannel activity section, field description [337](#), [338](#)
- Coupling Facility Activity (CF)
 - Postprocessor report

Coupling Facility Activity (CF) *(continued)*

Postprocessor report *(continued)*

- activity section [332](#)
- CF to CF activity section [339](#)
- command [324](#)
- contents of [324](#)
- how to request [324](#)
- purpose of [324](#)
- Spreadsheet and Overview reference [340](#)
- subchannel activity section [336](#)
- usage summary section [325](#)

Coupling Facility Activity (CFACT)

Monitor III report

- command [39](#)
- command abbreviation [19](#)
- contents of [40](#)
- field descriptions [42](#)
- how to request [39](#)
- purpose of [39](#)
- report options panel [46](#)
- special considerations [39](#)

Coupling Facility Overview (CFOVER)

Monitor III report

- command abbreviation [19](#)
- purpose of [47](#)

Coupling Facility Systems report

CFSYS report (Monitor III)

- command abbreviation [19](#)
- purpose of [49](#)

CPC (CPC Capacity)

Monitor III report

- command [56](#)
- contents of [56](#)
- field descriptions [57](#)
- how to request [56](#)
- purpose of [56](#)

CPC Capacity (CPC)

CPC report (Monitor III)

- contents of [56](#)
- field descriptions [57](#)

CPC Capacity report

CPC report (Monitor III)

- purpose of [56](#)

CPC Capacity report (Monitor III report) [19](#), [56](#)

CPU (CPU Activity)

Postprocessor report

- Blocked Workload Analysis [350](#)
- command [348](#)
- contents of [348](#), [350](#)
- duration report [369](#)
- how to request [348](#)
- purpose of [347](#)
- Spreadsheet and Overview reference [356](#), [363](#), [364](#), [366](#)
- System Address Space [350](#)
- using information in CPU Activity report [350](#)
- using information in Group Capacity report [365](#)
- using information in LPAR Cluster report [363](#)
- using information in Partition Data report [357](#)
- Work Unit Analysis [350](#)

CPU Activity (CPU)

Postprocessor report

- Blocked Workload Analysis [350](#)
- command [348](#)

CPU Activity (CPU) *(continued)*

Postprocessor report *(continued)*

- contents of [348](#), [350](#)
- duration report [369](#)
- how to request [348](#)
- purpose of [347](#)
- Spreadsheet and Overview reference [356](#), [363](#), [364](#), [366](#)
- System Address Space [350](#)
- using information in CPU Activity report [350](#)
- using information in Group Capacity report [365](#)
- using information in LPAR Cluster report [363](#)
- using information in Partition Data report [357](#)
- Work Unit Analysis [350](#)

CPU utilization

LPAR view [253](#), [349](#)

MVS view [253](#), [349](#)

CPU utilization field

in Monitor III SYSINFO report (average CPU UTIL) [183](#)

CPU Utilization field

in Monitor III WFEX report (Average CPU Util) [218](#)

in the Monitor III SYSTREND report (CPU %) [204](#)

criteria set

in Monitor III WFEX report definition and criteria panel [225](#)

cross system coupling facility job delays (XCFJ)

Monitor III report

field descriptions [128](#)

Cross-System Coupling Facility Activity (XCF)

Postprocessor report

- command [509](#)
- contents of [509](#)
- how to request [509](#)
- purpose of [509](#)

Cross-System Coupling Facility Delays (XCF)

Monitor III report

- command [232](#)
- command abbreviations [21](#)
- contents of [232](#)
- field descriptions [233](#)
- how to request [232](#)
- purpose of [232](#)
- report options [233](#)

CRYACC (CRYACC Crypto Accelerator Activity)

Monitor III report

- command [60](#)
- contents of [60](#)
- how to request [60](#)

CRYACC (CRYACC Crypto Accelerator)

Crypto Accelerator Activity report

- purpose of [59](#)

CRYACC (CRYACC Cryptographic accelerator activity)

Monitor III report

- field descriptions [60](#)

CRYACC Crypto Accelerator Activity (CRYACC)

CRYACC report (Monitor III)

- contents of [60](#)

CRYACC Crypto Accelerator report

CRYACC report (Crypto Accelerator Activity)

- purpose of [59](#)

CRYACC Crypto Accelerator report (Crypto Accelerator Activity report) [59](#)

- CRYACC Cryptographic accelerator activity (CRYACC)
 - CRYACC report (Monitor III)
 - field descriptions [60](#)
- CRYOVW (CRYOVW Crypto Hardware Overview)
 - Monitor III report
 - command [61](#)
 - contents of [61](#)
 - field descriptions [63](#)
 - how to request [61](#)
- CRYOVW Crypto (Hardware Overview Report) [61](#)
- CRYOVW Crypto Hardware Overview
 - CRYOVW Crypto Hardware Overview)
 - purpose of [61](#)
- CRYOVW Crypto Hardware Overview (CPC)
 - CRYOVW report (Monitor III)
 - contents of [61](#)
- CRYOVW Crypto Hardware Overview (CRYOVW)
 - CRYOVW report (Monitor III)
 - field descriptions [63](#)
- CRYOVW(CRYOVW Crypto Hardware Overview)
 - Hardware Overview report
 - purpose of [61](#)
- CRYPKC (CRYPKC Capacity)
 - PKCS11 Coprocessor Activity Report
 - purpose of [64](#)
- CRYPKC (CRYPKC Crypto PKCS11 Coprocessor)
 - Monitor III report
 - contents of [65](#)
 - field descriptions [65](#)
- CRYPKC (CRYPKC PKCS11 Coprocessor)
 - Monitor III report
 - command [64](#)
 - how to request [64](#)
- CRYPKC Crypto PKCS11 (Coprocessor Activity Report) [64](#)
- CRYPKC Crypto PKCS11 Coprocessor (CRYPKC)
 - CRYPKC report (Monitor III)
 - contents of [65](#)
 - field descriptions [65](#)
- CRYPTO (Crypto Hardware Activity)
 - Postprocessor report
 - command [370](#)
 - contents of [370](#)
 - how to request [370](#)
 - purpose of [369](#)
 - spreadsheet and overview reference [375](#)
- Crypto Hardware Activity (CRYPTO)
 - Postprocessor report
 - command [370](#)
 - contents of [370](#)
 - how to request [370](#)
 - purpose of [369](#)
 - spreadsheet and overview reference [375](#)
- Crypto Hardware Overview report
 - contents and field descriptions [63](#)
- Crypto PKCS11 Coprocessor report
 - contents and field descriptions [65](#)
- cryptographic accelerator [369](#)
- Cryptographic accelerator activity report
 - contents and field descriptions [60](#)
- cryptographic coprocessor [369](#)
- cryptographic processors
 - adjunct processor [502](#)
 - asynchronous message processor [502](#)
- CU (control unit) field

- CU (control unit) field (*continued*)
 - in the Monitor III DEVT report [83](#)

D

- Data Encryption Standard (DES) [369](#), [370](#), [373](#)
- data field
 - in Monitor III SYSRTD report [193](#)
- Data Index (DI)
 - Monitor III report
 - command [26](#)
 - command abbreviations [19](#)
 - condensed information [28](#)
 - contents of [26](#)
 - cursor-sensitive control [30](#)
 - field descriptions [29](#)
 - how to request [26](#)
 - messages [29](#)
 - options [31](#)
 - preallocated data sets [27](#)
 - purpose of [25](#)
 - sample [26](#)
- data set delays
 - DSNJ (Monitor III report) [88](#)
 - DSNV (Monitor III report) [89](#)
- Data Set Delays - Job (DSNJ)
 - Monitor III report
 - command [88](#)
 - how to request [88](#)
- Data Set Delays - Volume (DSNV)
 - Monitor III report
 - command [89](#)
 - how to request [89](#)
- Data Set Delays (DSND)
 - Monitor III report
 - command [86](#)
 - command abbreviations [19](#)
 - how to request [86](#)
 - purpose of [85](#)
- DCM
 - Monitor II
 - IOQUEUE report [281](#)
 - Monitor III
 - IOQUEUE report [114](#)
- DCM (dynamic channel path management)
 - in Monitor II Channel Path Activity report [267](#)
 - in Monitor III Channel Path Activity report [52](#)
 - in Postprocessor Channel Path Activity report [341](#)
 - in Postprocessor I/O Queuing Activity report [417](#)
- DDNAME
 - in Monitor III DI report [29](#)
- DEL Q LENGTH (delayed queue length) field
 - in Monitor III IOQUEUE report [116](#)
- DELAY
 - Monitor III report
 - command [66](#)
 - command abbreviation [19](#)
 - how to request [66](#)
 - job selection panel [72](#)
 - purpose of [66](#)
 - sample [66](#)
- Delay DB% (device busy delay percentage) field
 - in Monitor III DEVJ report [123](#)
- delay monitoring [4](#)

- delay percentage field
 - in Monitor III DELAY report (DLY %) [67](#)
 - in Monitor III DEV report (DLY %) [74](#)
 - in Monitor III DEVR report (DLY %) [81](#)
 - in Monitor III ENQ report (DLY %) [99](#)
 - in Monitor III ENQR report (Delayed) [104](#)
 - in Monitor III HSM report (DLY %) [113](#)
 - in Monitor III HSMJ report [125](#)
 - in Monitor III JESJ report [125](#)
 - in Monitor III STOR report (DLY %) [161](#)
 - in Monitor III XCF report (DLY %) [233](#)
 - in Monitor III XCFJ report [128](#)
- DELAY Report
 - Monitor III Utility fields [70](#)
- delayed name field
 - in Monitor III ENQR report [104](#)
- DELAYJ (Job Delay)
 - Monitor III report
 - command abbreviations [19](#)
- delta mode [254](#)
- DES (Data Encryption Standard) [369](#), [370](#), [373](#)
- description field
 - in Monitor III GROUP report [106](#)
- DEV (device delays)
 - Monitor III report
 - options panel [75](#)
 - sample [74](#)
- DEV (Device Delays)
 - Monitor III report
 - command [73](#)
 - command abbreviations [19](#)
 - how to request [73](#)
 - purpose of [73](#)
- DEV/CU (device/control unit)
 - field
 - in Monitor III DEVR report [80](#)
- DEV/DEVV (Device Activity)
 - Monitor II report
 - command [270](#)
 - contents of [272](#)
 - field descriptions [273](#)
 - how to request [270](#)
 - purpose of [270](#)
 - report options [275](#)
 - special considerations [271](#)
- DEVICE (Device Activity)
 - Postprocessor report
 - average cannot be calculated [387](#)
 - byte-multiplexor-channel-attached device [378](#)
 - command [377](#)
 - connect/disconnect time overflow [387](#)
 - contents of [378](#), [380](#), [381](#), [386–388](#)
 - data inaccuracies in duration report [389](#)
 - device changed or deleted [386](#)
 - device data incomplete or missing [386](#)
 - device in use [387](#)
 - device not available during entire interval [386](#)
 - Direct Access Device Activity report [378](#)
 - hardware data not available [387](#)
 - how to request [377](#)
 - NMBR and SG suboptions [378](#)
 - overflow condition occurred [387](#)
 - overflow in non-Monitor II reports [387](#)
 - pending time overflow [380](#), [381](#), [388](#)
 - purpose of [377](#)
 - required data not available [387](#)
 - spreadsheet and overview reference [388](#)
 - using information in the report [377](#)
- DEVICE (Device Activity) (*continued*)
 - Postprocessor report (*continued*)
 - purpose of [377](#)
 - required data not available [387](#)
 - spreadsheet and overview reference [388](#)
 - using information in the report [377](#)
- DEVICE (Synchronous I/O Device Activity)
 - Postprocessor report
 - contents of [379](#)
 - Direct Access Device Activity report [379](#)
- Device Activity (DEV/DEVV)
 - Monitor II report
 - command [270](#)
 - contents of [272](#)
 - field descriptions [273](#)
 - how to request [270](#)
 - purpose of [270](#)
 - report options [275](#)
 - special considerations [271](#)
- Device Activity (DEVICE)
 - Postprocessor report
 - average cannot be calculated [387](#)
 - byte-multiplexor-channel-attached device [378](#)
 - command [377](#)
 - contents of [378](#), [381](#), [386–389](#)
 - data inaccuracies in duration report [389](#)
 - device changed or deleted [386](#)
 - device data incomplete or missing [386](#)
 - device in use [387](#)
 - device not available during entire interval [386](#)
 - Direct Access Device Activity report [378](#)
 - hardware data not available [387](#)
 - how to request [377](#)
 - NMBR and SG suboptions [378](#)
 - overflow condition occurred [387](#)
 - overflow in non-Monitor II reports [387](#)
 - pending time overflow [381](#), [388](#)
 - purpose of [377](#)
 - required data not available [387](#)
 - spreadsheet and overview reference [388](#)
 - using information in the report [377](#)
- device activity (DEVN)
 - Monitor III report
 - sample [76](#)
- Device Activity (DEVN)
 - Monitor III report
 - command [76](#)
 - how to request [76](#)
 - purpose of [75](#)
- Device Activity Trend (DEVT)
 - Monitor III report
 - command [82–84](#)
 - description [84](#)
 - how to request [83](#), [84](#)
 - purpose of [82](#), [84](#)
- device busy delay field
 - in Monitor III DEVJ report (Delay DB%) [123](#)
- device delays
 - Monitor III report (DEV)
 - options panel [75](#)
 - sample [74](#)
- Device Delays - Job (DEVJ)
 - Monitor III report
 - command abbreviations [19](#)

- Device Delays - Job (DEVJ) (*continued*)
 - Monitor III report (*continued*)
 - how to request [122](#)
 - variation of JOB report [19](#), [122](#)
- Device Delays (DEV)
 - Monitor III report
 - command abbreviations [19](#)
 - purpose of [73](#)
- device field
 - in Monitor III DEVJ report [122](#)
- device number field
 - in the Monitor III DEVT report [83](#)
- Device Resource Delays (DEVR)
 - Monitor III report
 - command [78](#)
 - command abbreviations [19](#)
 - how to request [78](#)
 - purpose of [78](#)
- device type field
 - in Monitor III STORR report [174](#)
- device using percentage field
 - in Monitor III PROCJ report [126](#)
- device/control unit (DEV/CU)
 - field
 - in Monitor III DEVR report [80](#)
- DEVJ (Device Delays - Job)
 - Monitor III report
 - command abbreviations [19](#)
 - how to request [122](#)
 - variation of JOB report [19](#), [122](#)
- DEVJ (device delays job) [19](#), [122](#)
- DEVN (device activity)
 - Monitor III report
 - sample [76](#)
- DEVN (Device Activity)
 - Monitor III report
 - command [76](#)
 - how to request [76](#)
 - purpose of [75](#)
- DEVR (Device Resource Delays)
 - Monitor III report
 - command [78](#)
 - command abbreviations [19](#)
 - how to request [78](#)
 - purpose of [78](#)
- DEVT (Device Activity Trend)
 - Monitor III report
 - command [82–84](#)
 - how to request [83](#), [84](#)
 - purpose of [82](#)
- DI (Data Index)
 - Monitor III report
 - command [26](#)
 - command abbreviations [19](#)
 - condensed information [28](#)
 - contents of [26](#)
 - cursor-sensitive control [30](#)
 - field descriptions [29](#)
 - how to request [26](#)
 - messages [29](#)
 - options [31](#)
 - preallocated data sets [27](#)
 - purpose of [25](#)
 - sample [26](#)
- direct memory access (DMA) [447](#)
- director port busy field
 - in Monitor III IOQUEUE report (% DP BUSY) [117](#)
- disconnect time field
 - in Monitor III DEVJ report [123](#)
 - in Monitor III STORR report (DSC %) [175](#)
- Disk Space (SPACED)
 - Monitor III report
 - command [157](#), [158](#)
 - how to request [158](#)
 - purpose of [157](#)
- DLY % (delay percentage) field
 - in Monitor III DELAY report [67](#)
 - in Monitor III DEV report [74](#)
 - in Monitor III DEVR report [81](#)
 - in Monitor III ENQ report [99](#)
 - in Monitor III HSM report [113](#)
 - in Monitor III STOR report [161](#)
 - in Monitor III XCF report [233](#)
- DMA (direct memory access) [447](#)
- DMA (direct-memory-access) [138](#)
- DOMINO (Lotus Domino Server)
 - Postprocessor report
 - command [390](#)
 - contents of [390](#)
 - how to request [390](#)
 - purpose of [390](#)
- DSC % (disconnect time percentage) field
 - in Monitor III STORR report [175](#)
- DSD (Detailed Storage Delays)
 - Monitor III report
 - purpose of [84](#)
- DSINDEX
 - Monitor III report
 - command abbreviations [19](#)
 - purpose of [25](#)
- DSND (data set delays)
 - Monitor III report
 - field descriptions [86](#)
 - report options panel [87](#)
 - sample [86](#)
- DSND (Data Set Delays)
 - Monitor III report
 - command [86](#)
 - command abbreviations [19](#)
 - how to request [86](#)
 - purpose of [85](#)
- DSNJ (data set delays - Job report) [19](#)
- DSNJ (Data Set Delays - Job)
 - Monitor III report
 - command [88](#)
 - how to request [88](#)
- DSNV (data set delays - Volume report)
 - Monitor III report
 - command [89](#)
 - field descriptions [90](#)
 - purpose of [89](#)
 - report options panel [91](#)
 - sample [90](#)
- DSNV (Data Set Delays - Volume)
 - Monitor III report
 - command [89](#)
 - how to request [89](#)
- duration field

- duration field (*continued*)
 - in Monitor III SYSSUM report [202](#)
- duration reports [308](#)
- dynamic channel path management
 - Monitor II
 - IOQUEUE report [281](#)
 - Monitor III
 - IOQUEUE report [114](#)
- dynamic channel path management (DCM)
 - in Monitor II Channel Path Activity report [267](#)
 - in Monitor III Channel Path Activity report [52](#)
 - in Postprocessor Channel Path Activity report [341](#)
 - in Postprocessor I/O Queuing Activity report [417](#)

E

- EADM (EADM activity) [19](#)
- EADM (Extended Asynchronous Data Mover Activity Report)
 - Monitor III report
 - command [91](#)
 - contents of [91](#)
 - how to request [91](#)
 - purpose of [91](#)
 - Postprocessor report
 - command [393](#)
 - contents of [394](#)
 - how to request [393](#)
 - purpose of [393](#)
- earliest field
 - in the Monitor III DEVT report [83](#)
 - in the Monitor III SYSTREND report [204](#)
- ECKD [402](#)
- elapsed time (ELAP time) field
 - in Monitor III STORC report [165](#)
- ENCLAVE (enclave activities)
 - Monitor III report
 - command [93, 94](#)
 - field descriptions [94](#)
 - how to request [94](#)
 - sample [93](#)
- ENCLAVE (enclave activity) [19](#)
- enclave activities
 - Monitor III report
 - command [94](#)
 - how to request [94](#)
 - Monitor III report (ENCLAVE)
 - command [93](#)
 - field descriptions [94](#)
 - sample [93](#)
- Enclave Classification Attributes pop-up [95](#)
- Enclave Details pop-up [95](#)
- enclave support
 - DELAY report (Monitor III) [71](#)
 - ENCLAVE report (Monitor III) [93](#)
 - PROC report (Monitor III) [147](#)
 - SYSINFO report (Monitor III) [182](#)
- encryption [373](#)
- ENQ (Enqueue Activity)
 - Postprocessor report
 - command [396](#)
 - contents of [397](#)
 - how to request [396](#)
 - purpose of [395](#)
 - report levels [396](#)

- ENQ (Enqueue Activity) (*continued*)
 - Postprocessor report (*continued*)
 - spreadsheet and overview reference [399](#)
- ENQ (enqueue delays)
 - Monitor III report
 - sample [99](#)
- ENQ (Enqueue Delays)
 - Monitor III report
 - command [98, 99](#)
 - how to request [99](#)
 - purpose of [98](#)
- ENQJ (enqueue delays job) [20](#)
- ENQR (enqueue resource delays)
 - Monitor III report
 - field descriptions [104](#)
 - sample [103](#)
- ENQR (Enqueue Resource Delays)
 - Monitor III report
 - command [103](#)
 - how to request [103](#)
 - purpose of [103](#)
- ENQR (resource-oriented enqueue delays) [20](#)
- Enqueue Activity (ENQ)
 - Postprocessor report
 - command [396](#)
 - contents of [397](#)
 - how to request [396](#)
 - purpose of [395](#)
 - report levels [396](#)
 - spreadsheet and overview reference [399](#)
- enqueue delays
 - Monitor III report (ENQ)
 - sample [99](#)
 - Monitor III report (SYSENQ)
 - field descriptions [180](#)
- Enqueue Delays (ENQ)
 - Monitor III report
 - command [98, 99](#)
 - how to request [99](#)
 - purpose of [98](#)
- enqueue resource delays (ENQR)
 - Monitor III report
 - field descriptions [104](#)
 - sample [103](#)
- Enqueue Resource Delays (ENQR)
 - Monitor III report
 - command [103](#)
 - how to request [103](#)
 - purpose of [103](#)
- enterprise disk system [399](#)
- Enterprise Disk Systems (ESS)
 - Postprocessor report
 - command [401](#)
 - contents of [401](#)
 - ESS statistics [399, 401](#)
 - how to request [401](#)
 - purpose of [399](#)
- ESS (Enterprise Disk Systems)
 - Postprocessor report
 - command [401](#)
 - contents of [401](#)
 - ESS statistics [399, 401](#)
 - how to request [401](#)
 - purpose of [399](#)

- ESS rank statistics
 - solid state drive [405](#)
- EXCEPT (Exception)
 - Postprocessor report
 - command [513](#)
 - contents of [514](#), [515](#)
 - data fields [515](#)
 - heading fields [515](#)
 - how to request [513](#)
 - improved exception reporting [516](#)
 - purpose of [513](#)
- Exception (EXCEPT)
 - Postprocessor report
 - command [513](#)
 - contents of [514](#), [515](#)
 - data fields [515](#)
 - heading fields [515](#)
 - how to request [513](#)
 - improved exception reporting [516](#)
 - purpose of [513](#)
- Exceptions
 - in WFEX (Workflow/Exceptions) Monitor III report [216](#)
- EXCP rate field
 - in Monitor III GROUP report [108](#)
- execution velocity (actual) field
 - in Monitor III SYSSUM report [198](#)
- execution velocity (goal) field
 - in Monitor III SYSSUM report [198](#)
- Extended Asynchronous Data Mover (EADM) Activity Report (EADM)
 - Postprocessor report
 - command [393](#)
 - how to request [393](#)
- Extended Asynchronous Data Mover Activity Report (EADM)
 - Monitor III report
 - command [91](#)
 - contents of [91](#)
 - how to request [91](#)
 - purpose of [91](#)
 - Postprocessor report
 - contents of [394](#)
 - purpose of [393](#)
- extended count key data [402](#)

F

- FCD (FICON Director Activity)
 - Postprocessor report
 - command [407](#)
 - how to request [407](#)
 - purpose of [407](#)
- FICON Director Activity (FCD)
 - Postprocessor report
 - command [407](#)
 - how to request [407](#)
 - purpose of [407](#)
- fixed 1 MB memory objects [442](#)
- frames online field
 - in Monitor III STORR report [174](#)

G

- global resource serialization (GRS) [453](#)

- glossary [531](#)
- goal field
 - in Monitor III SYSWKM report [208](#)
- goal type field
 - in Monitor III SYSSUM report [202](#)
- goals versus actual field
 - in Monitor III SYSSUM report [198](#)
- GROUP (group response time) [20](#)
- GROUP (Group Response Time)
 - Monitor III report
 - command [105](#)
 - how to request [105](#)
 - purpose of [105](#)
- Group Capacity Report [365](#)
- Group Data(XCF)
 - Monitor III report
 - command [233](#)
 - how to request [233](#)
- Group Response Time (GROUP)
 - Monitor III report
 - command [105](#)
 - how to request [105](#)
 - purpose of [105](#)
- GRS (global resource serialization) [453](#)
- GRS enqueue request details [456](#)
- GRS enqueue requests summary [454](#)
- GRS latches [454](#), [457](#)

H

- Hardware Group Report
 - using the information in [366](#)
- help panels [17](#)
- heterogeneous report class period [196](#)
- heterogeneous report class periods [488](#)
- HFS (HFS statistics)
 - Postprocessor report
 - field descriptions [410–412](#)
 - sample [410](#), [411](#)
- HFS (Hierarchical File System Statistics)
 - Monitor II report
 - command [276](#)
 - contents of [277](#)
 - field descriptions [277](#)
 - how to request [276](#)
 - purpose of [276](#)
 - report options [278](#)
 - Postprocessor report
 - command [409](#)
 - how to request [409](#)
 - purpose of [409](#)
- Hierarchical File System Statistics (HFS)
 - Monitor II report
 - command [276](#)
 - contents of [277](#)
 - field descriptions [277](#)
 - how to request [276](#)
 - purpose of [276](#)
 - report options [278](#)
 - Postprocessor report
 - command [409](#)
 - how to request [409](#)
 - purpose of [409](#)
- hierarchical storage manager delays (HSM)

- hierarchical storage manager delays (HSM) *(continued)*
 - Monitor III report
 - report options panel [114](#)
 - sample [113](#)
- Hierarchical Storage Manager Delays (HSM)
 - Monitor III report
 - command [112](#)
 - how to request [112](#)
- High Performance FICON for System z
 - Monitor III [55](#), [270](#), [345](#)
- HiperDispatch
 - parked time [353](#)
- holding field
 - in Monitor III ENQJ report [124](#)
 - in Monitor III PROCJ report [126](#)
- holding name field
 - in Monitor III ENQR report [104](#)
- holding name/SYS field
 - in Monitor III ENQ report [103](#)
- holding percentage field
 - in Monitor III ENQ report [102](#)
- holding status field
 - in Monitor III ENQ report [103](#)
 - in Monitor III ENQR report [104](#)
- homogeneous report class period [195](#)
- homogeneous report class periods [488](#)
- How to read the [191](#)
- How to read the response time distribution table [191](#)
- HSM (hierarchical storage manager delays)
 - Monitor III report
 - report options panel [114](#)
 - sample [113](#)
- HSM (Hierarchical Storage Manager Delays)
 - Monitor III report
 - command [112](#)
 - how to request [112](#)
 - purpose of [112](#)
- HSMJ (hierarchical storage manager delays job) [20](#)
- HTTP (HTTP Server)
 - Postprocessor report
 - command [413](#)
 - how to request [413](#)
 - purpose of [413](#)
- HTTP Server (HTTP)
 - Postprocessor report
 - command [413](#)
 - how to request [413](#)
 - purpose of [413](#)

I

- I/O Queuing Activity (IOQ)
 - Postprocessor report
 - command [418](#)
 - contents of [420](#)
 - data gathering considerations [418](#)
 - how to request [418](#)
 - messages [419](#)
 - missing data in report fields [419](#)
 - purpose of [417](#)
 - spreadsheet and overview reference [427](#)
 - using information in report [418](#)
- I/O Queuing Activity (IOQUEUE)
 - Monitor II report

- I/O Queuing Activity (IOQUEUE) *(continued)*
 - Monitor II report *(continued)*
 - command [282](#)
 - contents of [283](#)
 - evaluating details of cumulative mode output [281](#)
 - field descriptions [283](#)
 - how to request [282](#)
 - purpose of [281](#)
 - report options [285](#)
 - special considerations [282](#)
 - Monitor III report
 - command [114](#), [115](#)
 - how to request [114](#), [115](#)
- IBM System z Application Assist Processor (zAAP)
 - in CPC Capacity Report [58](#)
 - in ENCLAVE Report [95](#)
 - in PROC Report [146](#)
 - in Processor Usage report [148](#)
- IBM System z Integrated Information Processor (zIIP)
 - in CPC Capacity Report [58](#)
 - in ENCLAVE Report [95](#)
 - in PROC Report [146](#)
 - in Processor Usage report [148](#)
 - in SYSINFO Report [183](#)
- IBM Z Application Assist Processor (zAAP)
 - definition [533](#)
 - how referred to [xxxv](#)
- IBM Z Integrated Information Processor (zIIP)
 - definition [533](#)
 - how referred to [xxxv](#)
- ID field
 - in Monitor III STORCR report [167](#)
- idling percentage (IDL %) field
 - in Monitor III DELAY report [68](#)
- IFL (Integrated Facility for Linux)
 - in CPC Capacity Report [58](#)
- ILOCK (IRLM Long Lock Detection)
 - Monitor II report
 - command [279](#)
 - contents of [280](#)
 - field descriptions [280](#)
 - how to request [279](#)
 - purpose of [279](#)
 - special considerations [279](#)
- importance (I) field
 - in Monitor III SYSSUM report [198](#)
- in Monitor III SYSSUM report (Actual) [202](#)
- indicator field
 - in Monitor III WFEX report options action panel [222](#)
- installed at field
 - in Monitor III SYSSUM report [197](#)
- Integrated Facility for Linux (IFL)
 - in CPC Capacity Report [58](#)
- Integrated Information Processor, *See* IBM System z Integrated Information Processor (zIIP)
- interval reports [308](#)
- IOQ (I/O Queuing Activity)
 - Postprocessor report
 - command [418](#)
 - contents of [420](#)
 - data gathering considerations [418](#)
 - how to request [418](#)
 - messages [419](#)
 - missing data in report fields [419](#)

- IOQ (I/O Queuing Activity) (*continued*)
 - Postprocessor report (*continued*)
 - purpose of [417](#)
 - spreadsheet and overview reference [427](#)
 - using information in report [418](#)
- IOQUEUE (I/O queuing activity) [20](#)
- IOQUEUE (I/O Queuing Activity)
 - Monitor II report
 - command [282](#)
 - contents of [283](#)
 - evaluating details of cumulative mode output [281](#)
 - field descriptions [283](#)
 - how to request [282](#)
 - purpose of [281](#)
 - report options [285](#)
 - special considerations [282](#)
 - Monitor III report
 - command [114](#), [115](#)
 - how to request [115](#)
 - purpose of [114](#)
- IOS queue time field
 - in the Monitor III DEVN report (IosQ) [77](#)
- IPL definitions field
 - in Monitor III STORC report [163](#)
- IRLM Long Lock Detection (ILOCK)
 - Monitor II report
 - command [279](#)
 - contents of [280](#)
 - field descriptions [280](#)
 - how to request [279](#)
 - purpose of [279](#)
 - special considerations [279](#)

J

- JES (Job Entry Subsystem Delays)
 - Monitor III report
 - command [117](#), [118](#)
 - how to request [118](#)
 - purpose of [117](#)
- JES delays
 - Monitor III reporter session command [20](#)
- JESJ (JES delays for specified job) [20](#)
- JOB (Job Delay)
 - Monitor III report
 - command [120](#)
 - command abbreviations [20](#)
 - Device Delay variation [122](#)
 - how to request [120](#)
 - purpose of [120](#)
- Job Delay (JOB)
 - Monitor III report
 - command [120](#)
 - command abbreviations [20](#)
 - Device Delay variation [122](#)
 - how to request [120](#)
 - purpose of [120](#)
- Job Delay report
 - Monitor III Utility fields [128](#)
- job ended date field
 - in Monitor III STORCR report [167](#)
- job ended time field
 - in Monitor III STORCR report [167](#)
- Job Entry Subsystem Delays (JES)

- Job Entry Subsystem Delays (JES) (*continued*)
 - Monitor III report
 - command [117](#), [118](#)
 - how to request [118](#)
 - purpose of [117](#)
- job field
 - in Monitor III ENQJ report [124](#)
 - in Monitor III HSMJ report [124](#)
 - in Monitor III JESJ report [124](#)
 - in Monitor III PROCJ report [126](#)
- job report menu [23](#)
- job report options panel [128](#)
- Job Usage (USAGE)
 - Monitor III report
 - command [212](#)
 - contents of [212](#)
 - cursor-sensitive control on [215](#)
 - field descriptions [213](#)
 - how to request [212](#)
 - utility fields [214](#)
- jobname field
 - in Monitor III DEV report [74](#)
 - in Monitor III DEVJ report (JOB) [121](#)
 - in Monitor III DEVR report [80](#)
 - in Monitor III ENQ report [99](#)
 - in Monitor III HSM report [113](#)
 - in Monitor III MSGJ report [125](#)
 - in Monitor III PROC report [146](#)
 - in Monitor III STOR report [160](#)
 - in Monitor III STORCR report [167](#)
 - in Monitor III STORF report [168](#)
 - in Monitor III SYSWKM report [211](#)

K

- keyboard
 - navigation [525](#)
 - PF keys [525](#)
 - shortcut keys [525](#)

L

- label field
 - in Monitor III WFEX report definition and criteria panel [225](#)
 - in Monitor III WFEX report options action panel [222](#)
- latest field
 - in the Monitor III DEVT report [83](#)
 - in the Monitor III SYSTREND report [204](#)
- LCU field
 - in Monitor III IOQUEUE report [116](#)
- Library List (LLI)
 - Monitor II report
 - command [285](#)
 - Contents of the report - APF Library List [287](#)
 - Contents of the report - Link Library List [286](#)
 - Contents of the report - LPA Library List [287](#)
 - field descriptions [287](#)
 - how to request [285](#)
 - purpose of [285](#)
- LLI (Library List)
 - Monitor II report
 - command [285](#)

- LLI (Library List) *(continued)*
 - Monitor II report *(continued)*
 - contents of the report - APF Library List [287](#)
 - contents of the report - Link Library List [286](#)
 - contents of the report - LPA Library List [287](#)
 - field descriptions [287](#)
 - how to request [285](#)
 - purpose of [285](#)
- LLI (link library list)
 - Monitor II report
 - field descriptions [287](#)
- local page data set slots [439](#)
- LOCKSP (Spin Lock)
 - Monitor III report
 - command [129](#)
 - how to request [129](#)
 - purpose of [129](#)
- LOCKSP (Spin locks) [20](#)
- LOCKSU (Suspend Lock)
 - Monitor III report
 - command [130](#), [131](#)
 - contents of [131](#)
 - how to request [131](#)
 - purpose of [130](#)
- LOCKSU (Suspend locks) [20](#)
- long-term overview reporting
 - with the Postprocessor [307](#)
- long-term performance analysis
 - RMF XP
 - long-term performance analysis [4](#)
 - with RMF XP [4](#)
- Lotus Domino Server (DOMINO)
 - Postprocessor report
 - command [390](#)
 - how to request [390](#)
 - Postprocessor reports
 - contents of [390](#)
 - purpose of [390](#)
- LPAR cluster
 - description [363](#)
 - Postprocessor report [363](#)
- LPAR mode
 - CPU utilization [349](#)

M

- magnetic tape Activity (DEVICE)
 - Postprocessor report
 - contents of [380](#)
 - pending time overflow [380](#)
- main delay path field
 - in Monitor III XCF report [233](#)
- main delay volume(s) field
 - in Monitor III DEV report [75](#)
- Memory field
 - in Monitor III SYSSUM report (Actual) [202](#)
- memory objects
 - fixed 1 MB [442](#)
 - shared 1 MB [442](#)
- MNTJ (operator mount job delays) [20](#)
- model number field
 - in Monitor III SYSINFO report [183](#)
- Monitor I
 - Real-time reporting with [305](#)

- Monitor I session
 - system activity measured [3](#)
- Monitor II background session
 - system activity measured [3](#)
- Monitor II display session
 - system activity measured [3](#)
- Monitor III
 - shortened intervals, reason for [15](#)
 - sysplex support in different time zones [16](#)
- Monitor III data gatherer session
 - delay information [4](#)
 - system activity measured [3](#)
- Monitor III menu
 - Overview Report Selection Menu [23](#)
- Monitor III report commands
 - header for single-system reports [21](#)
 - header for sysplex reports [22](#)
- Monitor III reporter session
 - commands [21](#)
 - delay information [4](#)
 - starting [17](#)
 - system activity measured [3](#)
- Monitor III Utility
 - customizing CHANNEL report with [55](#)
- MSGJ (operator message job delays) [20](#)
- MSU/h [56](#), [57](#)
- Multiple Lock Structure [157](#)
- MVS view
 - CPU utilization [253](#)

N

- name field
 - in Monitor III SYSINFO report [184](#)
 - in Monitor III SYSSUM report [197](#)
 - in Monitor III WFEX report [218](#), [231](#)
 - in Monitor III WFEX report definition and criteria panel [226](#)
- navigation
 - keyboard [525](#)
- No. of completed transactions (For nnnn TRX) field
 - in Monitor III SYSWKM report [208](#)
- number field
 - in Monitor III DEVJ report [122](#)

O

- OMVS
 - address spaces [132](#)
 - OPD report [132](#)
- OMVS (OMVS Kernel Activity)
 - Postprocessor report
 - command [428](#)
 - contents of [429](#)
 - field descriptions [429](#)
 - how to request [428](#)
 - purpose of [428](#)
- OMVS Kernel Activity (OMVS)
 - Postprocessor report
 - command [428](#)
 - contents of [429](#)
 - field descriptions [429](#)
 - how to request [428](#)

OMVS Kernel Activity (OMVS) (*continued*)

Postprocessor report (*continued*)

purpose of [428](#)

OMVS Process Data (OPD)

Monitor III report

command [132](#)

contents of [133](#)

how to request [132](#)

purpose of [132](#)

report options [137](#)

OPD (OMVS process data) [20](#)

OPD (OMVS Process Data)

Monitor III report

command [132](#)

contents of [133](#)

how to request [132](#)

purpose of [132](#)

report options [137](#)

OPT (OPT Settings)

Monitor II report

command [288](#)

contents of [289](#)

field descriptions [289](#)

how to request [288](#)

purpose of [288](#)

OPT Settings (OPT)

Monitor II report

command [288](#)

contents of [289](#)

field descriptions [289](#)

how to request [288](#)

purpose of [288](#)

OVERVIEW

Postprocessor report

command [517](#)

creating Overview records [519](#)

description of [517](#)

how to request [517](#)

improved summary reporting [518](#)

Overview data record [521](#)

Overview header record [519](#)

purpose of [517](#)

overview condition names [312](#)

overview conditions [312](#)

overview names [312](#)

Overview Report Selection Menu [23](#)

P

Page Data Set Activity (PAGESP)

Postprocessor report

command [431](#)

contents of [431](#)

how to request [431](#)

overview names in [433](#)

purpose of [431](#)

using information in the report [431](#)

Page Data Set Activity (PGSP)

Monitor II report

command [290](#)

contents of [290](#)

field descriptions [291](#)

how to request [290](#)

purpose of [290](#)

Page Data Set Activity (PGSP) (*continued*)

Monitor II report (*continued*)

special considerations [290](#)

PAGESP (Page Data Set Activity)

Postprocessor report

command [431](#)

contents of [431](#)

how to request [431](#)

overview names in [433](#)

purpose of [431](#)

using information in the report [431](#)

PAGING (Paging Activity)

Postprocessor report

Central Storage Movement Rates [437](#)

central storage paging rates [434](#)

Central Storage Request Rates [437](#)

command [433](#)

fixed 1 MB memory objects [442](#)

frame counts [439](#)

high virtual storage frames [442](#)

how to request [433](#)

local page data set slots [439](#)

memory objects [442](#)

purpose of [433](#)

SCM paging blocks [439](#)

shared 1 MB memory objects [442](#)

shared frames [439](#)

slot counts [439](#)

spreadsheet and overview reference [437](#), [438](#), [441](#), [444](#)

Paging Activity (PAGING)

Postprocessor report

Central Storage Movement Rates [437](#)

central storage paging rates [434](#)

Central Storage Request Rates [437](#)

command [433](#)

Enhanced DAT Facility [442](#)

fixed 1 MB memory objects [442](#)

frame counts [439](#)

high virtual storage frames [442](#)

how to request [433](#)

local page data set slots [439](#)

memory objects [442](#)

purpose of [433](#)

SCM paging blocks [439](#)

shared 1 MB memory objects [442](#)

shared frames [439](#)

slot counts [439](#)

spreadsheet and overview reference [437](#), [438](#), [441](#), [444](#)

Paging Activity (SPAG)

Monitor II report

command [299](#)

contents of [300](#)

field descriptions [300](#)

how to request [299](#)

purpose of [299](#)

Paging Activity report

request rates and frames [438](#)

parked time [353](#)

Partition Data Report

z/VM or alternate VM guest [347](#)

path attributes [116](#), [283](#)

Path Data(XCF)

Path Data(XCF) (*continued*)
 Monitor III report
 command [235](#)
 how to request [235](#)
 path number field
 in Monitor III XCFJ report [128](#)
 PAV (parallel access volume)
 in Monitor I DEVICE report [381](#)
 PCI functions [138](#), [447](#)
 PCIE (PCI Express-based functions) [20](#)
 PCIE (PCIE Activity Report)
 Monitor III report
 command [137](#), [138](#)
 contents of [138](#)
 how to request [138](#)
 purpose of [137](#)
 PCIE (PCIE Activity)
 Postprocessor report
 command [447](#)
 contents of [447](#)
 how to request [447](#)
 purpose of [447](#)
 spreadsheet and overview [452](#)
 PCIE Activity (PCIE)
 Postprocessor report
 command [447](#)
 contents of [447](#)
 how to request [447](#)
 purpose of [447](#)
 spreadsheet and overview [452](#)
 PCIE Activity Report
 cursor sensitive control [139](#)
 PCIE Activity Report
 pop-up panels [139](#)
 PCIE Activity Report (PCIE)
 Monitor III report
 command [137](#), [138](#)
 contents of [138](#)
 how to request [138](#)
 purpose of [137](#)
 peak allocation values field
 in Monitor III STORC report [163](#)
 peer-to-peer remote copy [402](#)
 pend time percentage field
 in Monitor III DEVJ report [123](#)
 in Monitor III DEVR report (PND % Reasons) [80](#)
 in Monitor III STORR report (PND %) [175](#)
 in the Monitor III DEVN report (PND % Reasons) [77](#)
 Pending (pend time percentage) field
 in the Monitor III DEVN report [77](#)
 percent delayed for field
 in Monitor III DELAY report (% Delayed for) [69](#)
 in Monitor III STOR report [161](#)
 percent used field
 in Monitor III STORC report [165](#)
 percentage frames (% Frames) field
 in Monitor III STORR report [174](#)
 performance index [199](#)
 period field
 in Monitor III GROUP report [106](#)
 in Monitor III SYSWKM report [208](#)
 peripheral-component-interconnect (PCI) devices [138](#), [447](#)
 PGSP (Page Data Set Activity)
 Monitor II report

PGSP (Page Data Set Activity) (*continued*)
 Monitor II report (*continued*)
 command [290](#)
 contents of [290](#)
 field descriptions [291](#)
 how to request [290](#)
 purpose of [290](#)
 special considerations [290](#)
 phase (P) field
 in Monitor III SYSWKM report [209](#)
 PND % (pend time percentage) field
 in Monitor III DEVR report [80](#)
 in Monitor III STORR report [175](#)
 possible cause or action field
 in Monitor III WFEX report [221](#)
 Postprocessor
 long-term overview reporting with [307](#)
 PPRC [402](#)
 preallocated datasets
 in DSINDEX report [27](#)
 preferred path [116](#), [283](#)
 primary delay category field
 in Monitor III HSMJ report [125](#)
 in Monitor III JESJ report [125](#)
 primary delay reason field
 in Monitor III HSMJ report [125](#)
 in Monitor III JESJ report [125](#)
 primary reason field
 in Monitor III DELAY report [70](#)
 primary response time component field
 in Monitor III GROUP report [107](#)
 probable cause field
 in Monitor III DEVJ report [122](#)
 PROC (processor delays) [20](#)
 PROC (Processor Delays)
 Monitor III report
 command [145](#)
 contents of [146](#)
 field descriptions [146](#)
 how to request [145](#)
 purpose of [145](#)
 report options [148](#)
 utility fields [147](#)
 Processor Delays (PROC)
 Monitor III report
 command [145](#)
 contents of [146](#)
 field descriptions [146](#)
 how to request [145](#)
 purpose of [145](#)
 report options [148](#)
 utility fields [147](#)
 Processor Usage (PROCU)
 Monitor III report
 command [148](#)
 contents of [149](#)
 field descriptions [149](#)
 how to request [148](#)
 purpose of [148](#)
 report options [151](#)
 utility fields [150](#)
 processor using percentage field
 in Monitor III PROCJ report [126](#)
 in Monitor III SYSWKM report [211](#)

- PROCJ (processor job delays) [20](#)
- PROCU (processor usage) [20](#)
- PROCU (Processor Usage)
 - Monitor III report
 - command [148](#)
 - contents of [149](#)
 - field descriptions [149](#)
 - how to request [148](#)
 - purpose of [148](#)
 - report options [151](#)
 - utility fields [150](#)
- Programming interface information [530](#)
- promotion threshold [356](#)

Q

- QSCJ (quiescent jobs) [20](#)
- qualifier field
 - in Monitor III WFEX report options action panel [222](#)
- quiesce delay (Quies) field
 - in Monitor III SYSWKM report [211](#)

R

- Range field
 - in the Monitor III DEVT report [83, 84](#)
 - in the Monitor III SYSTREND report [204](#)
- Real-time reporting
 - with Monitor I [305](#)
- reason field
 - in Monitor III WFEX report [219](#)
- Reasons field
 - in Monitor III DEVR report [80](#)
 - in the Monitor III DEVN report [77](#)
- refresh field
 - in Monitor III SYSSUM report [197](#)
- report class period
 - heterogeneous [196](#)
 - homogeneous [195](#)
- report class periods
 - heterogeneous [488](#)
 - homogeneous [488](#)
- report mode
 - delta [254](#)
- report options (Monitor III) [16](#)
- Resource Group Activity Report (SYSRG)
 - Monitor III report
 - command [186–189](#)
 - how to request [187](#)
 - purpose of [186–189](#)
- resource group field
 - in Monitor III SYSSUM report [202](#)
- resource name field
 - in Monitor III ENQR report [104](#)
- Resource Report Selection Menu [24](#)
- resource waiting percentage field
 - in Monitor III ENQ report [100](#)
- resource waiting status field
 - in Monitor III ENQ report [100](#)
- response time
 - average [503](#)
 - midpoint [503](#)
- response time (actual) field

- response time (actual) field (*continued*)
 - in Monitor III SYSSUM report [199](#)
- response time (goal) field
 - in Monitor III SYSSUM report [199](#)
- response time breakdown field (delayed by)
 - in Monitor III SYSWKM report [210](#)
- response time distribution (SYSRTD)
 - Monitor III report
 - option panel [193](#)
 - sample [190](#)
- Response Time Distribution (SYSRTD)
 - Monitor III report
 - command [190](#)
 - how to request [190](#)
 - purpose of [190](#)
- response time distribution table [191](#)
- response time field
 - in Monitor III SYSINFO report (RESP Time) [184](#)
 - in Monitor III SYSRTD report [192](#)
 - in the Monitor III DEVN report (RspT) [77](#)
- RLSDS (VSAM RLS activity by data set) [20](#)
- RLSLRU (VSAM LRU Overview)
 - Monitor III report
 - command [151](#)
 - how to request [151](#)
 - purpose of [151](#)
- RLSLRU (VSAM RLS activity LRU statistics) [20](#)
- RLSLRU report
 - buffer counts by pool [152](#)
- RLSSC (VSAM RLS activity by storage class) [20](#)
- RLSSC/RLSDS (VSAM RLS Activity)
 - Monitor III report
 - command [153](#)
 - how to request [153](#)
 - purpose of [153](#)
 - RLSDS command [153](#)
 - RLSSC command [153](#)
- RMF (Resource Measurement Facility)
 - help panels [17](#)
 - tutorial [17](#)
- RMF command [17](#)
- row and position field
 - in Monitor III WFEX report options action panel [222](#)

S

- SAID [402](#)
- SCM
 - in Postprocessor PAGESP report [432](#)
 - Storage Class Memory [432](#)
 - See also* EADM
- SCM (storage class memory) [325](#)
- SCM paging blocks [439](#)
- SCSI [402](#)
- SDELAY (Serialization Delay)
 - XML Postprocessor report
 - command [453](#)
 - contents of [454](#)
 - how to request [453](#)
 - purpose of [453](#)
 - serialization delay details [456](#)
 - serialization delay summary [454](#)
- SDEVICE (Shared Device Activity)
 - Postprocessor report

SDEVICE (Shared Device Activity) *(continued)*

Postprocessor report *(continued)*

- command [461](#)
- field descriptions [462](#)
- how to request [461](#)
- purpose of [461](#)
- Shared Direct Access Device Activity report [461](#)
- Shared Magnetic Tape Device Activity report [462](#)
- spreadsheet and overview reference [467](#)
- using information in the report [461](#)

SDS (Sysplex Data Server)

Monitor II report

- command [292](#)
- contents of [292](#)
- field descriptions [292](#)
- how to request [292](#)
- purpose of [292](#)

selecting a report [17](#)

selection criteria

- in the Monitor III DEVN report [77](#)

SENQ (System Enqueue Contention)

Monitor II report

- command [293](#)
- contents of [295](#)
- field descriptions [296](#)
- how to request [293](#)
- purpose of [293](#)
- report options [297](#)
- types of [294](#)

SENQR (System Enqueue Reserve)

Monitor II report

- command [298](#)
- contents of [298](#)
- field descriptions [298](#)
- how to request [298](#)
- purpose of [297](#)
- report options [299](#)

Serialization Delay (SDELAY)

XML Postprocessor report

- command [453](#)
- contents of [454](#)
- how to request [453](#)
- purpose of [453](#)
- serialization delay details [456](#)
- serialization delay summary [454](#)

server field

- in Monitor III ENQJ report [124](#)

service class field

- in Monitor III SYSWKM report [211](#)

service definition field

- in Monitor III SYSSUM report [197](#)

service field

- in Monitor III SYSWKM report [211](#)

Service Policy page [489](#)

SERVICE TIME [493](#)

session options [xxxv](#)

SHA-1 hash algorithm [374](#)

SHA-256 hash algorithm [374](#)

Shared Device Activity (SDEVICE)

Postprocessor report

- command [461](#)
- field descriptions [462](#)
- how to request [461](#)
- purpose of [461](#)

Shared Device Activity (SDEVICE) *(continued)*

Postprocessor report *(continued)*

- Shared Direct Access Device Activity report [461](#)
- Shared Magnetic Tape Device Activity report [462](#)
- spreadsheet and overview reference [467](#)
- using information in the report [461](#)

shared field

- in Monitor III DEVJ report [122](#)

shared frames, PAGING report [439](#)

shortcut keys [525](#)

shortened intervals, reason for [15](#)

SHR field

- in Monitor III CHANNEL report [54](#)

slot counts

- available slots, PAGING report [441](#)
- bad slots, PAGING report [441](#)
- NON-VIO slots, PAGING report [441](#)
- total slots, PAGING report [441](#)
- VIO slots, PAGING report [441](#)

slots

- PAGING report [439](#)

small computer system interface [402](#)

software pricing

- Postprocessor CPU report [357](#)

solid state drive [405](#)

sort criteria

- in the Monitor III DEVN report [77](#)

SPACED (disk space) [20](#)

SPACED (Disk Space)

Monitor III report

- command [157](#), [158](#)
- how to request [158](#)
- purpose of [157](#)

SPACEG (storage group space) [20](#)

SPACEG (Storage Space)

Monitor III report

- command [158](#), [159](#)
- how to request [159](#)
- purpose of [158](#)

SPAG (Paging Activity)

Monitor II report

- command [299](#)
- contents of [300](#)
- field descriptions [300](#)
- how to request [299](#)
- purpose of [299](#)

Speed (Workflow)

- in WFEX (Workflow/Exceptions) Monitor III report [216](#)

speed field

- in Monitor III WFEX report [218](#)

speedometer

- in Monitor III WFEX report [232](#)

Spin Lock (LOCKSP)

Monitor III report

- command [129](#)
- how to request [129](#)
- purpose of [129](#)

Spin Lock Report

- commands [129](#)

SQA + ESQA (system queue area) overflow field

- in Monitor III STORR report [174](#)

SRB field

- in the Monitor III SYSTREND report [204](#)

- SRCS (Central Storage/Processor/
SRM)
 - Monitor II report
 - command [301](#)
 - contents of [302](#)
 - field descriptions [302](#)
 - how to request [301](#)
 - purpose of [301](#)
- starting and ending an RMF reporter session [17](#)
- status field
 - in Monitor III ENQJ report [124](#)
- STOR (storage delays)
 - Monitor III report
 - report options panel [162](#)
- STOR (Storage Delays)
 - Monitor III report
 - command [159](#)
 - how to request [159](#)
 - purpose of [159](#)
- STOR Report
 - Monitor III Utility fields [161](#)
- storage class memory [91](#)
 - See also* extended asynchronous data mover (EADM)
- storage class memory (SCM) [325](#)
- Storage Class Memory (SCM) [432](#)
- storage delay summary
 - Monitor III report (STORS)
 - field descriptions [178](#)
 - options panel [179](#)
 - sample [177](#)
- Storage Delay Summary (STORS)
 - Monitor III report
 - command [177](#)
 - how to request [176](#), [177](#)
 - purpose of [176](#)
- storage delays
 - Monitor III report (STOR)
 - report options panel [162](#)
- Storage Delays (STOR)
 - Monitor III report
 - command [159](#)
 - how to request [159](#)
 - purpose of [159](#)
- storage frames
 - Monitor III report (STORF)
 - report options panel [169](#)
- Storage Frames (STORF)
 - Monitor III report
 - command [167](#), [168](#)
 - how to request [168](#)
 - purpose of [167](#)
- Storage Memory Objects (STORM)
 - Monitor III report
 - command [170](#)
 - how to request [170](#)
 - purpose of [170](#)
- Storage Resource Delays (STORR)
 - Monitor III report
 - command [172](#)
 - contents of [173](#)
 - field descriptions [174](#)
 - how to request [172](#)
 - purpose of [172](#)
 - utility fields [176](#)
- Storage Space (SPACEG)
 - Monitor III report
 - command [158](#), [159](#)
 - how to request [159](#)
 - purpose of [158](#)
 - storage summary field
 - in Monitor III STORS report [177](#)
- STORC (common storage)
 - Monitor III report
 - field descriptions - system information [164](#)
- STORC (Common Storage)
 - Monitor III report
 - command [162](#)
 - how to request [162](#)
 - purpose of [162](#)
- STORC report
 - unallocated CSA + SQA [164](#)
 - unallocated ERUCSA [164](#)
 - unallocated RUCSA [164](#)
- STORC Report
 - Monitor III Utility fields [165](#)
- STORCR (common storage remaining)
 - Monitor III report
 - field descriptions [167](#)
- STORCR (Common Storage Remaining)
 - Monitor III report
 - command [166](#)
 - how to request [166](#)
 - purpose of [166](#)
- STORF (storage frames)
 - Monitor III report
 - report options panel [169](#)
- STORF (Storage Frames)
 - Monitor III report
 - command [167](#), [168](#)
 - how to request [168](#)
 - purpose of [167](#)
- STORJ (storage job delays) [21](#)
- STORM
 - field descriptions [170](#)
- STORM (storage memory objects) [170](#)
- STORM (Storage Memory Objects)
 - Monitor III report
 - command [170](#)
 - how to request [170](#)
 - purpose of [170](#)
- STORM (storage usage by memory objects) [21](#)
- STORR (resource-oriented storage delays) [21](#)
- STORR (Storage Resource Delays)
 - Monitor III report
 - command [172](#)
 - contents of [173](#)
 - field descriptions [174](#)
 - how to request [172](#)
 - purpose of [172](#)
 - utility fields [176](#)
- STORS (storage delay summary)
 - Monitor III report
 - options panel [179](#)
 - sample [177](#)
- STORS (Storage Delay Summary)
 - Monitor III report
 - command [176](#), [177](#)
 - how to request [177](#)

STORS (Storage Delay Summary) (*continued*)

Monitor III report (*continued*)

purpose of [176](#)

STORS (storage delays summary) [21](#)

STORS Report

Monitor III Utility fields [178](#)

subsystem data field

in Monitor III SYSRTD report [193](#)

subsystem menu [24](#)

subsystem reports [24](#)

subsystem type field

in Monitor III SYSWKM report [209](#)

SUMMARY

Postprocessor report

command [521](#)

generating [521](#)

how to request [521](#)

purpose of [521](#)

special considerations [522](#)

using the Spreadsheet Reporter [523](#)

summary of changes [xxxix](#), [xl](#)

Suspend Lock (LOCKSU)

Monitor III report

command [130](#), [131](#)

contents of [131](#)

how to request [131](#)

purpose of [130](#)

Suspend Lock Report

commands [131](#)

switched time field

in Monitor III SYSWKM report [210](#)

Synchronous I/O Device Activity (DEVICE)

Postprocessor report

contents of [379](#)

Synchronous I/O Device Activity report

[379](#)

SYSENQ (sysplex enqueue delays)

Monitor III report

field descriptions [180](#)

SYSENQ (Sysplex Enqueue Delays)

Monitor III report

command [180](#)

how to request [180](#)

purpose of [180](#)

SYSINFO (system information)

Monitor III report

sample [182](#)

SYSINFO (System Information)

Monitor III report

command [181](#)

field descriptions [183](#)

how to request [181](#)

purpose of [181](#)

utility fields [185](#)

Sysplex Data Server (SDS)

Monitor II report

command [292](#)

contents of [292](#)

field descriptions [292](#)

how to request [292](#)

purpose of [292](#)

Sysplex Enqueue Delays (SYSENQ)

Monitor III report

command [180](#)

Sysplex Enqueue Delays (SYSENQ) (*continued*)

Monitor III report (*continued*)

how to request [180](#)

purpose of [180](#)

Sysplex Report Selection Menu [22](#)

sysplex summary

Monitor III report (SYSSUM)

sample [197](#)

Sysplex Summary (SYSSUM)

Monitor III report

command [193](#), [194](#)

how to request [194](#)

purpose of [193](#)

SYSRG (Resource Group Activity report) [21](#)

SYSRG (Resource Group Activity Report)

Monitor III report

command [186–189](#)

how to request [187](#)

purpose of [186–189](#)

SYSRTD (response time distribution)

Monitor III report

option panel [193](#)

sample [190](#)

SYSRTD (Response Time Distribution)

Monitor III report

command [190](#)

how to request [190](#)

purpose of [190](#)

SYSSUM (sysplex summary)

Monitor III report

sample [197](#)

SYSSUM (Sysplex Summary)

Monitor III report

command [193](#), [194](#)

how to request [194](#)

purpose of [193](#)

SYSSUM Report

field descriptions [197](#)

field descriptions- continuation [200](#)

Monitor III Utility fields [202](#)

system activities measured

Monitor I session [3](#)

Monitor II session [3](#)

Monitor III session [3](#)

Postprocessor [3](#)

System Enqueue Contention (SENQ)

Monitor II report

command [293](#)

contents of [295](#)

field descriptions [296](#)

how to request [293](#)

purpose of [293](#)

report options [297](#)

types of [294](#)

System Enqueue Reserve (SENQR)

Monitor II report

command [298](#)

contents of [298](#)

field descriptions [298](#)

how to request [298](#)

purpose of [297](#)

report options [299](#)

system field

in Monitor III ENQJ report [124](#)

system field (*continued*)
in Monitor III SYSWKM report [192](#)

system information
Monitor III report (SYSINFO)
sample [182](#)
Monitor III report (SYSTREND)
field descriptions [204](#)

System Information (SYSINFO)
Monitor III report
command [181](#)
field descriptions [183](#)
how to request [181](#)
purpose of [181](#)
utility fields [185](#)

system name field
in Monitor III SYSWKM report [211](#)
system queue area (SQA + ESQA) overflow field
in Monitor III STORR report [174](#)

system suspend locks [454](#)

System Trend (SYSTREND)

Monitor III report
command [203](#)
how to request [203](#)
purpose of [203](#)

System z Application Assist Processor, *See* IBM System z
Application Assist Processor (zAAP)

System z Integrated Information Processor, *See* IBM System
z Integrated Information Processor (zIIP)

Systems Data(XCF)
Monitor III report
command [237](#)
how to request [237](#)

Systems Overview(XCF)
Monitor III report
command [234](#)
how to request [234](#)

SYSTREND (system trend)
Monitor III report
field descriptions [204](#)

SYSTREND (System Trend)
Monitor III report
command [203](#)
how to request [203](#)
purpose of [203](#)

SYSWKM (work manager delays)
Monitor III report
report options panel [211](#)
sample [206](#)

SYSWKM (Work Manager Delays)
Monitor III report
command [205](#)
how to request [205](#)
purpose of [205](#)

T

TCB field
in the Monitor III SYSTREND report [204](#)

text field
in Monitor III WFEX report definition and criteria panel
[225](#)

time field
in Monitor III SYSSUM report [199](#)

Time field

Time field (*continued*)
in the Monitor III DEVT report (Time) [84](#)
in the Monitor III SYSTREND report (Time) [204](#)

Time Stamp field
in the Monitor III DEVT report (Time) [84](#)
in the Monitor III SYSTREND report (Time) [204](#)

time zones
Monitor III sysplex support [16](#)

total utilization field
in Monitor II channel path activity (CHANNEL) report [55](#)

TRACE - Trace Activity report
purpose of [467](#)

TRACE (Trace Activity)
Postprocessor report
command [467](#)
contents of [468](#)
how to request [467](#)
standard deviation output [471](#)
using information in the report [468](#)

Trace Activity (TRACE)
Postprocessor report
command [467](#)
contents of [468](#)
how to request [467](#)
standard deviation output [471](#)
using information in the report [468](#)

trademarks [530](#)

TRANS/SEC (transactions per second)
field
in Monitor III SYSINFO report [184](#)

Transaction ended rate field
in Monitor III GROUP report [108](#)

transaction rate (Trx Rate) field
in Monitor III SYSRTD report [193](#)

transaction response time field
in Monitor III SYSINFO report (RESP Time) [184](#)

transactions per second (TRANS/SEC)
field
in Monitor III SYSINFO report [184](#)

tutorial [17](#)

type (T) field
in Monitor III SYSSUM report [198](#)

Type field
in the Monitor III DEVT report [83](#)

U

UKN % (unknown percentage) field
in Monitor III DELAY report [68](#)

unallocated CSA + SQA
in Monitor III STORC report [164](#)

unallocated ERUCSA
in Monitor III STORC report [164](#)

unallocated RUCSA
in Monitor III STORC report [164](#)

unknown percentage (UKN %) field
in Monitor III DELAY report [68](#)

USAGE (Job Usage)
Monitor III report
command [212](#)
contents of [212](#)
cursor-sensitive control on [215](#)
field descriptions [213](#)
how to request [212](#)

- USAGE (Job Usage) *(continued)*
 - Monitor III report *(continued)*
 - utility fields [214](#)
- Usage report
 - purpose of [212](#)
- user interface
 - ISPF [525](#)
 - TSO/E [525](#)
- user menu [25](#)
- user reports [25](#)
- users field
 - in Monitor III SYSINFO report [184](#)
- users/active field
 - in Monitor III WFEX report [218](#)
- using percentage field
 - in Monitor III DELAY report (USG %) [67](#)
 - in Monitor III DEV report (USG %) [74](#)
 - in Monitor III HSMJ report [125](#)
 - in Monitor III JESJ report [125](#)
 - in Monitor III PROC report (USG %) [147](#)

V

- vector utilization field
 - in Monitor III GROUP report (Vector UTIL) [107](#)
- velocity field
 - in Monitor III SYSWKM report [211](#)
- Virtual Storage Activity (VSTOR)
 - High Virtual Memory Usage section [479](#)
 - Postprocessor report
 - allocated storage [472](#)
 - command [472](#)
 - Common Storage Detail section [475](#)
 - Common Storage Summary section [473](#)
 - contents of [473](#)
 - data gathering considerations [473](#)
 - free storage [472](#)
 - how to request [472](#)
 - overview reference [475](#)
 - partial private area data [478](#)
 - Private Area Detail section [479](#)
 - Private Area Summary section [476](#)
 - Private Storage Map [477](#), [478](#)
 - purpose of [472](#)
 - using information in [472](#)
- VMGUEST Monitor I gatherer option [347](#)
- VolSer field
 - in the Monitor III DEVT report [83](#)
- volume serial field
 - in Monitor III STORR report [174](#)
- VSAM LRU Overview (RLSLRU)
 - Monitor III report
 - command [151](#)
 - how to request [151](#)
 - purpose of [151](#)
- VSAM RLS Activity (RLSSC/RLSDS)
 - Monitor III report
 - command [153](#)
 - how to request [153](#)
 - purpose of [153](#)
 - RLSDS command [153](#)
 - RLSSC command [153](#)
- VSTOR (Virtual Storage Activity)
 - Postprocessor report

- VSTOR (Virtual Storage Activity) *(continued)*
 - Postprocessor report *(continued)*
 - allocated storage [472](#)
 - command [472](#)
 - Common Storage Detail section [475](#)
 - Common Storage Summary section [473](#)
 - contents of [473](#)
 - data gathering considerations [473](#)
 - free storage [472](#)
 - High Virtual Memory Usage section [479](#)
 - how to request [472](#)
 - overview reference [475](#)
 - partial private area data [478](#)
 - Private Area Detail section [479](#)
 - Private Area Summary section [476](#)
 - Private Storage Map [477](#), [478](#)
 - purpose of [472](#)
 - using information in [472](#)

W

- WFEX (workflow/exceptions)
 - Monitor III report
 - field descriptions [232](#)
 - graphic report [231](#)
 - sample [218](#)
 - summary [231](#)
- WFEX (Workflow/Exceptions)
 - Monitor III report
 - command [216](#)
 - how to request [216](#)
 - purpose of [216](#)
- WFL % (workflow percentage) field
 - in Monitor III DELAY report [67](#)
 - in Monitor III SYSINFO report [184](#)
 - the Monitor III GROUP report [107](#)
- WLMGL (Workload Activity)
 - Postprocessor report
 - command [480](#)
 - contents of [481](#), [485–488](#)
 - field descriptions [490](#)
 - how to request [480](#)
 - option list [480](#)
 - Policy Summary report [487](#)
 - purpose of [480](#)
 - Report Class Period report [488](#)
 - Report Class report [487](#)
 - Service Class Period report [481](#)
 - Service Class report [485](#)
 - spreadsheet and overview reference [506](#)
 - WLMGL option list [480](#)
 - Workload Group and Service Class Period report [486](#)
 - Workload Group report [486](#)
- work manager delays (SYSWKM)
 - Monitor III report
 - field descriptions [208](#)
 - report options panel [211](#)
 - sample [206](#)
- Work Manager Delays (SYSWKM)
 - Monitor III report
 - command [205](#)
 - how to request [205](#)
 - purpose of [205](#)

- workflow percentage field
 - in Monitor III DELAY report (WFL %) [67](#)
 - in Monitor III HSMJ report [125](#)
 - in Monitor III JESJ report [125](#)
 - in Monitor III SYSINFO report (WFL %) [184](#)
 - the Monitor III GROUP report (WFL %) [107](#)
- workflow/exceptions (WFEX)
 - Monitor III report
 - field descriptions [232](#)
 - graphic report [231](#)
 - sample [218](#)
 - summary [231](#)
- Workflow/Exceptions (WFEX)
 - Monitor III report
 - command [216](#)
 - how to request [216](#)
 - purpose of [216](#)
- Workload Activity (WLMGL)
 - Postprocessor report
 - command [480](#)
 - contents of [481](#), [485–488](#)
 - field descriptions [490](#)
 - how to request [480](#)
 - option list [480](#)
 - Policy Summary report [487](#)
 - purpose of [480](#)
 - Report Class Period report [488](#)
 - Report Class report [487](#)
 - Service Class Period report [481](#)
 - Service Class report [485](#)
 - spreadsheet and overview reference [506](#)
 - WLMGL option list [480](#)
 - Workload Group and Service Class Period report [486](#)
 - Workload Group report [486](#)

X

- XCF (Cross-System Coupling Facility Activity)
 - Postprocessor report
 - command [509](#)
 - contents of [509](#)
 - how to request [509](#)
 - purpose of [509](#)
- XCF (Cross-System Coupling Facility Delays)
 - Monitor III report
 - command [232](#)
 - command abbreviations [21](#)
 - contents of [232](#)
 - field descriptions [233](#)
 - how to request [232](#)
 - purpose of [232](#)
 - report options [233](#)
- XCF (Group Data report)
 - Monitor III report
 - command [233](#)
 - how to request [233](#)
- XCF (Group Data Report)
 - Monitor III report
 - field descriptions [234](#)
- XCF (Group Data)
 - Monitor III report

- XCF (Group Data) (*continued*)
 - Monitor III report (*continued*)
 - contents of [234](#)
- XCF (Path Data report)
 - Monitor III report
 - command [235](#)
 - how to request [235](#)
- XCF (Path Data) [235](#)
- XCF (Path)
 - Monitor III report
 - contents of [236](#)
- XCF (Systems Data report)
 - Monitor III report
 - command [237](#)
 - how to request [237](#)
- XCF (Systems Data Report) [237](#)
- XCF (Systems Data)
 - Monitor III report
 - contents of [237](#)
 - field descriptions [237](#)
- XCF (Systems Overview report)
 - Monitor III report
 - command [234](#)
 - how to request [234](#)
- XCF (Systems Overview Report)
 - Monitor III report
 - field descriptions [235](#)
- XCF (Systems Overview)
 - Monitor III report
 - contents of [235](#)
- XCF (XCF Group Data) [233](#)
- XCFJ (cross system coupling facility job delays)
 - Monitor III report
 - field descriptions [128](#)

Y

- yellow and red (Yel Red) fields
 - in Monitor III WFEX report definition and criteria panel [226](#)

Z

- z/OS Unix
 - address spaces [132](#)
 - OPD report [132](#)
- z/VM or alternate VM guest
 - Partition Data Report [347](#)
- zAAP, *See* IBM System z Application Assist Processor (zAAP)
- zFS File System (ZFSFS)
 - Monitor III report
 - command [238](#), [240](#)
 - contents of the report [240](#)
 - field descriptions [241](#)
 - how to request [238](#)
 - purpose of [238](#)
- zFS Kernel (ZFSKN)
 - Monitor III report
 - command [243](#)
 - contents of the report [244](#)
 - field descriptions [244](#)
 - how to request [243](#)
 - purpose [243](#)

- zFS Overview (ZFSOVW)
 - Monitor III report
 - command [245](#)
 - contents of the report [245](#)
 - field descriptions [245](#)
 - how to request [245](#)
 - navigating to details in [246](#)
 - purpose of [244](#)
- ZFSFS (zFS File System)
 - Monitor III report
 - command [238](#), [240](#)
 - contents of the report [240](#)
 - field descriptions [241](#)
 - how to request [238](#)
 - purpose of [238](#)
- ZFSKN (zFS Kernel)
 - Monitor III report
 - command [243](#)
 - contents of the report [244](#)
 - field descriptions [244](#)
 - how to request [243](#)
 - purpose [243](#)
- ZFSOVW (zFS Overview)
 - Monitor III report
 - command [245](#)
 - contents of the report [245](#)
 - field descriptions [245](#)
 - how to request [245](#)
 - navigating to details in [246](#)
 - purpose of [244](#)
- zHPF
 - Monitor III [55](#), [270](#), [345](#)
- zIIP, *See* IBM System z Integrated Information Processor (zIIP)



Product Number: 5655-ZOS

SC34-2665-70

