

z/OS



Resource Measurement Facility Performance Management Guide

Version 1 Release 10

z/OS



Resource Measurement Facility Performance Management Guide

Version 1 Release 10

Note

Before using this information and the product it supports, be sure to read the general information under “Notices” on page 213.

| This edition applies to Version 1 Release 10 of z/OS (5694-A01) and to all subsequent releases and modifications
| until otherwise indicated in new editions.

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

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

IBM Deutschland Research & Development GmbH
Department 3248
Schönaicher Str. 220
D-71032 Böblingen
Federal Republic of Germany

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

FAX (RMF Development): Your International Access Code +49+7031+16+4240
Internet: rmf@de.ibm.com

Internet

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

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

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

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

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

© Copyright International Business Machines Corporation 1993, 2008. All rights reserved.

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

Contents

Figures v

About this document vii

Who should use this document	vii
When to use this document	vii
How this document is organized	vii
The z/OS RMF library	viii
Related information	viii
Using LookAt to look up message explanations	viii

Summary of changes xi

What's new in z/OS V1R10	xi
Support of HiperDispatch mode	xi
64-Bit Common Storage measurements	xi
New reports on spin locks and suspend locks	xi
Support of WLM enqueue management	xi
Providing DASD volume capacity	xii
History of changes	xii
What's new in z/OS V1R9	xii
What's new in z/OS V1R8	xiii
What's new in z/OS V1R7	xiii
What's new in z/OS V1R6	xiv
What's new in z/OS V1R4	xiv
What's new in z/OS V1R2	xv
What's new in z/OS V1R1	xv

Chapter 1. Performance overview 1

Defining performance management	2
Setting performance goals	2
Planning system capacity	3
Benefits of capacity planning	3
Common mistakes in capacity planning	3
Performance management versus capacity planning	4
Using RMF	5
What sampling cycle and reporting interval should you use?	6
Analyzing transaction response time	7
General formulas	7
GROUP report	8
Analyzing workload characteristics	9
Identifying workloads	9
Measuring resource utilization by workload	10
Using Monitor I reports to analyze workloads	11
Analyzing processor characteristics	13
Analyzing workload I/O characteristics	16
Analyzing processor storage characteristics	18
Processor speed indicators	20
Where do you go from here?	22

Chapter 2. Diagnosing a problem: The first steps 23

What is a performance problem?	24
Getting started: A top-down approach to tuning	25

Where to start sysplex monitoring?	26
Monitor III Sysplex Summary report	26
Monitor III Response Time Distribution report	26
Monitor III Work Manager Delays report	27
Postprocessor Workload Activity report	27
Monitor III indicators	28
Identifying the major delay in a response time problem	35
Using Monitor III reports	36
Monitoring daily performance: System indicators	41
Summary of major system indicators	42
Using Monitor I reports	43
Where do you go from here?	59

Chapter 3. Analyzing processor activity 61

Do you have a processor problem?	62
Monitor III indicators	62
Monitor I indicators	68
Is your processor idle due to other bottlenecks?	72
DELAY for processor: Guidelines for improvement	73
Determine CPU 'holders'	73
Review dispatching priorities	74
Check for CPU delay due to other LPARs	75
Check LPAR balancing using the CPC capacity report	75
Processor USING: Trimming activities	76
Check for loops	76
Check for slip traps	78
Reduce I/O	78
Decrease monitor overhead	79
Tune system address spaces	79
Tracing	79
Calculate your capture ratio	79
Redesign application	79
Install more CPU power	79
Summary	80

Chapter 4. Analyzing I/O activity 81

Do you have an I/O problem?	82
I/O subsystem health check	84
Understand your business priorities	84
Review CPU and processor storage	84
Establish a baseline measurement	84
Improving I/O performance with the Enterprise Storage Server	85
z/OS Parallel Sysplex I/O management	85
ESS performance features	85
Cache performance	89
CACHE	89
DASD	89
Cache management with ESS	90
Cache Subsystem Activity report	92
Using Monitor III cache reporting	102
DASD indicators and guidelines	104
Response time components	104

Spreadsheet reports	105
General DASD guidelines	107
DASD performance in the sysplex	117
Analyzing specific DASD problems	118
Performance analysis on data set level	122
Improving your DASD performance.	124
Tape indicators and guidelines	126
Identifying tape-bound jobs	126
RMF measurements for tape monitoring	126
Response time components.	127
General tape indicators	128
Improving your tape performance	129
Summary.	130

Chapter 5. Analyzing processor storage activity 131

Do you have a processor storage problem?	132
Monitor III indicators.	133
Monitor I indicators	138
General storage recommendations	144
Increase storage	144
Auxiliary storage tuning.	144
Prioritize access to processor storage	144
Tune DASD	144
Summary.	144

Chapter 6. Analyzing sysplex activity 145

Understanding Workload Activity data for IMS or CICS	146
Interpreting the Workload Activity report	146
Problem: Very large response time percentage	150
Problem: Response time is zero	151
Problem: More executed transactions than ended transactions	152
Problem: Execution time is greater than response time	154
Problem: Large SWITCH percentage in CICS execution phase	154
Problem: Decreased number of ended transaction with increased response times	155
Analyzing coupling facility activity	157
Using the Postprocessor Coupling Facility Activity report	157
Spreadsheet report	172
Using the Monitor III online reports.	172
Factors affecting Coupling Facility performance	174

Appendix A. PR/SM LPAR considerations 179

Understanding the Partition Data report	180
---------------------------------------------------	-----

Defining logical processors	182
---------------------------------------	-----

Appendix B. The Intelligent Resource Director 185

Overview.	186
Dynamic channel path management.	186
Value of dynamic channel path management	187
Reporting of dynamic channel path management	187
Channel subsystem priority queuing	189
Value of channel subsystem priority queuing	190
Reporting of channel subsystem priority queuing	190
LPAR CPU management	190
LPAR cluster	190
Value of LPAR CPU management	192
LPAR CPU management decision controls.	192
LPAR CPU management controls.	193
Reporting of LPAR CPU management	193

Appendix C. Planning considerations for HiperDispatch mode 197

Introduction to HiperDispatch mode	197
How RMF supports HiperDispatch mode	198
Processing benefits	199

Appendix D. Other delays 201

Enqueue delays	202
ENQ report	202
Major names	202
HSM delays	204
JES delays	205
OPER delays	205
Unknown delays	205

Accessibility. 207

Using assistive technologies	207
Keyboard navigation of the user interface	207
z/OS information	207

Glossary 209

Notices 213

Policy for unsupported hardware.	214
Programming Interface Information	214
Trademarks	214

Index 217

Figures

1. End-to-End Response Time Components	7	50. Device Queuing in a Parallel Access Volume Environment	88
2. Monitor III Group Response Time Report	8	51. Cache Performance Management: The Key	89
3. System Resource Summary by Workload	10	52. Cache Subsystem Activity report - Summary	92
4. Monitor I CPU Activity Report - Part 1	11	53. Cache Subsystem Activity Report - Top-20 Device Lists	93
5. Workload Activity Report - Service Class	12	54. Cache Subsystem Activity Report - Status and Overview	94
6. Workload Activity Report - Service Policy	12	55. Cache Subsystem Activity Report - Device Overview	97
7. Processor Utilization by Workload	13	56. Cache Subsystem Activity Report - RAID Rank Activity	97
8. I/O Activity by Workload.	16	57. Cache Subsystem Activity Report - Cache Device Activity (device-level reporting)	98
9. Workload Activity Report - I/O Activity	17	58. Overview Report for Cached DASD Device	99
10. Processor Storage by Workload	18	59. Cache Hits Overview Report	100
11. Workload Activity Report - Processor Storage Use	19	60. Cache Trend Report	100
12. Simplified View of Performance Management	25	61. CACHSUM Report.	102
13. Monitor III Sysplex Summary Report	28	62. CACHDET Report	103
14. Monitor III Sysplex Summary Report - GO Mode	29	63. DASD Response Time Components	104
15. Monitor III Response Time Distribution report	32	64. DASD Summary Report	105
16. Monitor III Work Manager Delays Report	34	65. Activity of Top-10 Volumes	106
17. End-to-End Response Time Components	35	66. Response Time of Top-10 Volumes	107
18. Monitor III Group Response Time Report	37	67. Direct Access Device Activity Report	108
19. Monitor III Storage Delays Report	38	68. Shared DASD Activity Report	117
20. Monitor III Delay Report	38	69. Monitor III System Information Report	118
21. Monitor III Job Delays Report	39	70. Monitor III Delay Report.	119
22. Monitor III Workflow/Exceptions report	40	71. Monitor III Device Delays Report	120
23. Summary of Major System Indicators	42	72. Monitor III Device Resource Delays Report	121
24. Monitor I CPU Activity Report	44	73. Monitor III - Device Resource Delays Report	122
25. CPU Contention	46	74. Monitor III - Data Set Delays by Volume	122
26. Monitor I CPU Activity Report - Partition Data Report Section	47	75. Monitor III - Data Set Delays by Job	123
27. Monitor I Channel Path Activity report	48	76. Monitor III - Data Set Delays by Data Set	124
28. Monitor I I/O Queuing Activity Report	49	77. Magnetic Tape Device Activity Report	127
29. Monitor I DASD Activity Report	50	78. Monitor III System Information Report	133
30. DASD Summary	51	79. Monitor III Delay Report.	135
31. Response Time of Top-10 Volumes	52	80. Monitor III Storage Delays Report	136
32. Paging Activity report - Swap Placement Activity	53	81. Paging Activity report - Central Storage Paging Rates.	138
33. Monitor I Page/Swap Data Set Activity Report	54	82. PAGING Report - Central Storage Movement Rates / Frame and Slot Counts	140
34. Workload Activity report	55	83. Monitor I Page/Swap Data Set Activity Report	141
35. Monitor I Virtual Storage Activity Report - Common Storage Summary	57	84. Direct Access Device Activity Report	142
36. Monitor III Group Response Time Report	63	85. Hotel Reservations Service Class	146
37. Monitor III Delay Report	64	86. Response Time Breakdown of CICSHR accessing DBCTL	147
38. Monitor III Workflow/Exceptions Report	65	87. Point-of-Sale Service Class	148
39. Monitor III Processor Delays Report	66	88. Response Time Breakdown of CICSPS accessing DBCTL with IMS V5.	149
40. Monitor I CPU Activity Report	68	89. CICS User Transactions Service Class	149
41. CPU Contention Report	70	90. Response Time Percentages Greater than 100	150
42. Monitor I CPU Activity Report - Partition Data Report Section	71	91. Response Time Percentages all Zero	152
43. Monitor III Job Delays Report	73	92. Executed Transactions greater than Ended Transactions	153
44. Monitor III Address Space State Data Report	74	93. Execution Time Greater than Response Time	154
45. Monitor III CPC Capacity Report	75		
46. Monitor III Processor Delays Report	76		
47. Monitor III Job Delays Report	77		
48. Monitor III Workflow/Exceptions Report	78		
49. z/OS Traditional Device Serialization	86		

94. Large SWITCH Percentage in a CICS Execution Environment	155	102. CFSYS Report	174
95. Coupling Facility Activity Report - Usage Summary	158	103. CPU Activity Report	180
96. Coupling Facility Activity Report - Structure Activity	163	104. Partition Data Report	181
97. Coupling Facility Activity Report - Subchannel Activity	169	105. CHANNEL Report.	188
98. Coupling Facility Activity Report - CF to CF Activity	171	106. IOQUEUE Report	189
99. Coupling Facility Structure Activity	172	107. LPAR Cluster - Example 1	191
100. CFOVER Report	173	108. LPAR Cluster - Example 2	191
101. CFACT Report	174	109. LPAR Cluster Report	193
		110. Partition Data Report	195
		111. CPU Activity report	198
		112. Monitor III Enqueue Delays Report	202
		113. HSM Report	204

About this document

The z/OS Resource Measurement Facility (RMF™) is a performance management tool that is designed to measure selected areas of system activity and present the data collected in the form of *system management facility (SMF)* records, formatted printed reports, or formatted display reports. Use RMF to report on resource usage, evaluate system performance and identify reasons for performance problems.

This document offers a practical, task-oriented approach to analyzing performance issues using RMF. Covering the most frequently asked questions, it offers guidelines and suggestions concerning performance management of a z/OS sysplex.

Who should use this document

Anyone involved in understanding, measuring, and improving the performance of z/OS systems. This document assumes that the reader has a basic understanding of RMF and how z/OS systems work. The detailed description of all RMF functions and reports is available in *z/OS RMF User's Guide* and *z/OS RMF Report Analysis*.

When to use this document

Use this document when you are trying to improve the performance of your z/OS sysplex. Reasons could include:

- Complaints are coming in from users
- Service level objectives are not being met
- Ongoing performance monitoring showing adverse trends
- Resource usage is changing
- Furthering your own understanding

Do not use this document when your efforts have taken you beyond the scope of RMF and MVS™ performance. For example:

- Network tuning
- Detailed internal subsystem tuning for DB2®. Refer to the *DATABASE 2 Client/Server Administration Guide (SC09-1909)* instead.

How this document is organized

This document contains the following chapters:

Chapter 1, "Performance overview"

The chapter describes performance concepts common to all z/OS systems.

Chapter 2, "Diagnosing a problem: The first steps"

This chapter helps you analyze a performance problem to determine which resource is constrained, and thus where you should proceed with further analysis.

Chapter 3, "Analyzing processor activity"

Here, you find information for further analysis of CPU constraints.

Chapter 4, “Analyzing I/O activity”

This chapter concentrates on the analysis of I/O constraints (Cache subsystem, DASD, and tape).

Chapter 5, “Analyzing processor storage activity”

The analysis of processor storage constraints (central storage and paging subsystem) will be explained in this chapter.

Chapter 6, “Analyzing sysplex activity”

This information unit covers key indicators for workload management and coupling facility performance aspects.

Appendix A. PR/SM LPAR considerations

Here, you find performance considerations for systems running in LPAR mode.

Appendix B. The Intelligent Resource Director

This information unit describes the functions of the Intelligence Resource Director and its reporting in RMF.

Appendix C. Planning considerations for HiperDispatch mode

This information unit provides a *CPU Activity* report sample showing metrics for a system running in HiperDispatch mode. These specific metrics are explained and HiperDispatch processing benefits are discussed.

Appendix D. Other delays

This information unit discusses some delays shown in Monitor III that have not been covered before.

The z/OS RMF library

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

Table 1. RMF Library

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

Related information

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

Using LookAt to look up message explanations

LookAt is an online facility that lets you look up explanations for most of the IBM® messages you encounter, as well as for some system abends and codes. Using LookAt to find information is faster than a conventional search because in most cases LookAt goes directly to the message explanation.

You can use LookAt from these locations to find IBM message explanations for z/OS[®] elements and features, z/VM[®], z/VSE[™], and Clusters for AIX[®] and Linux[™]:

- The Internet. You can access IBM message explanations directly from the LookAt Web site at www.ibm.com/servers/eserver/zseries/zos/bkserv/lookat/.
- Your z/OS TSO/E host system. You can install code on your z/OS systems to access IBM message explanations using LookAt from a TSO/E command line (for example: TSO/E prompt, ISPF, or z/OS UNIX[®] System Services).
- Your Microsoft[®] Windows[®] workstation. You can install LookAt directly from the *z/OS Collection* (SK3T-4269) or the *z/OS and Software Products DVD Collection* (SK3T-4271) and use it from the resulting Windows graphical user interface (GUI). The command prompt (also known as the DOS > command line) version can still be used from the directory in which you install the Windows version of LookAt.
- Your wireless handheld device. You can use the LookAt Mobile Edition from www.ibm.com/servers/eserver/zseries/zos/bkserv/lookat/lookatm.html with a handheld device that has wireless access and an Internet browser.

You can obtain code to install LookAt on your host system or Microsoft Windows workstation from:

- A CD in the *z/OS Collection* (SK3T-4269).
- The *z/OS and Software Products DVD Collection* (SK3T-4271).
- The LookAt Web site (click **Download** and then select the platform, release, collection, and location that suit your needs). More information is available in the LOOKAT.ME files available during the download process.

Summary of changes

This document contains information previously presented in *z/OS RMF Performance Management Guide*, SC33-7992-07.

This document includes terminology, maintenance, and editorial changes. The following information describes the enhancements that are being distributed with z/OS Version 1 Release 10. All technical changes or additions to the text are indicated by a vertical line to the left of the change.

What's new in z/OS V1R10

Summary of changes
for SC33-7992-08
z/OS Version 1 Release 10

Support of HiperDispatch mode

“Appendix C. Planning considerations for HiperDispatch mode” on page 197 provides an introduction to HiperDispatch mode with RMF report examples.

64-Bit Common Storage measurements

The new Monitor III **Storage Memory Objects** report (STORM) contains measurement of 64-bit virtual common storage consumption per system and per job. Measurements of the usage of memory objects previously contained in the Monitor III **Storage Resource Delays** (STORR), **Storage Delay Summary** (STORS) and **Storage Frames** (STORF) reports are moved to this new report.

Also, the following Postprocessor reports are enhanced to provide 64-bit virtual common storage usage information:

- In the *Frame and Slot Counts* section of the **Paging Activity** report, a new *Memory Objects and Frames* section is added.
- In the *Private Area Detail* section of the **Virtual Storage Activity** report, the section containing memory allocation values above 2GB is enhanced with new information about memory objects and 1 MB frames.

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

New reports on spin locks and suspend locks

RMF periodically checks all types of system resource locks and provides information about lock holders and lock contention in two new Monitor III reports:

- **Spin Lock Report**
- **Suspend Lock Report**

Support of WLM enqueue management

The Postprocessor **Workload Activity** report is enhanced to show the CPU times for work units that have been promoted by WLM for reasons of enqueue management and chronic resource management.

Providing DASD volume capacity

The Postprocessor *Direct Access Device Activity* report shows the volume capacity for DASD devices specified by the number of cylinders.

In addition, RMF provides a new overview condition for the Postprocessor based on SMF record 74-1.

History of changes

What's new in z/OS V1R9

Summary of changes
for SC33-7992-07
z/OS Version 1 Release 9

Measurements about WLM's blocked workload promotion

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, but does not get CPU service within a certain time interval due to its low dispatch priority, will be temporarily promoted to a higher dispatch priority.

RMF supports this function by reporting relevant measurements in the new *Blocked Workload Promotion* section of the Postprocessor **CPU Activity** report. The Postprocessor **Workload Activity** report provides the CPU time, that transactions of a certain service or report class were running at a promoted dispatching priority.

In addition, RMF provides new overview conditions for the Postprocessor based on SMF record 72-3.

Enhanced monitoring of FICON performance

The average number of concurrently active I/O operations is provided in the LCU summary line of the Postprocessor **I/O Queuing Activity** report if at least one FICON channel is connected to the LCU.

HyperPAV support

In HyperPAV mode, PAV-aliases are no longer statically bound to PAV-bases but are bound only for the duration of a single I/O operation. Thus, the number of aliases required for an LCU is reduced. This enables applications to achieve equal or better performance compared to the original PAV feature.

RMF supports this new feature by extending the following reports:

- The Postprocessor **I/O Queuing Activity** report shows information about HyperPAV activity.
- Several device related and delay related reports show the average number of assigned HyperPAV-aliases.

New Coupling Facility measurements

RMF provides additional data in the Monitor III **Coupling Facility** reports and the Postprocessor **Coupling Facility Activity** report. For example, the reports display CF utilization per structure and whether dynamic CF dispatching is turned on.

In addition, RMF provides new overview conditions for the Postprocessor based on SMF record 74-4.

What's new in z/OS V1R8

Summary of changes
for SC33-7992-06
z/OS Version 1 Release 8

Support of IBM System z9 Integrated Information Processors (zIIP)

RMF provides measurements about zIIP activity in various Postprocessor and Monitor III reports.

Terminology:

RMF provides a consistent naming for zAAPs and zIIPs. The term IFA which was used for the zAAP in all RMF reports, help panels and manuals is now replaced by AAP. That is, AAP is used for the zAAP and IIP is used for the zIIP as field names.

Capacity groups

A new function of z/OS 1.8 allows customers to apply the defined capacity limit not only to a single LPAR, but also to a group of LPARs running on the same CEC, a so-called capacity group.

RMF extends its CPU activity reporting with information about capacity groups. This allows to assess the capacity group definitions and their effect on the system.

- A new *Group Capacity* section is available in the Postprocessor **CPU Activity** report. This section informs about the share each LPAR in the group can take, the guaranteed minimum MSU share of the LPARs and the overall MSU consumption within the group.
- If a partition is a member of a capacity group, the Monitor III **CPC Capacity** report and the *Partition Data Report* section of the **CPU Activity** report include new header information showing the group's name and capacity limit.

What's new in z/OS V1R7

Summary of changes
for SC33-7992-05
z/OS Version 1 Release 7

Support for IBM System z9 Enterprise Class processors

On z900 and z990 processors, special purpose processors IFLs (Integrated Facility for Linux) and IFAs (Integrated Facility for Applications) are reported as ICFs (Internal Coupling Facility). Starting with IBM System z9[®] Enterprise Class (z9 EC) processors, IFLs, IFAs and ICFs are reported separately in the Postprocessor **Partition Data Report** and the Monitor III **CPC Report**.

Queue length distribution in the CPU Activity report

The presentation of the queue length distribution in the *System Address Space Analysis* section of the **CPU Activity** Postprocessor report is changed to reflect the higher number of CPUs available per MVS image.

What's new in z/OS V1R6

Summary of changes
for SC33-7992-03
z/OS Version 1 Release 6

What's new in SC33-7992-04

Obsolete information refreshed or deleted: The former *Appendix C. Data-in-Memory* has been deleted and the *Introduction* section of "Improving I/O performance with the Enterprise Storage Server" on page 85 has been reworked. Also, various other minor changes have been implemented to reflect up-to-date information.

What's new in SC33-7992-03

Support of the @server zSeries® Application Assist Processor: z/OS V1R6 provides the ability to run Java™ applications on a new type of processor called the @server zSeries Application Assist Processor (zAAP). The zSeries Application Assist Processor is also known as an IFA (Integrated Facility for Applications). RMF uses the term IFA in the affected reports, messages, and in this and other publications.

RMF provides measurements of IFA processor activity in the following reports:

- Postprocessor:
 - CPU Activity report and its Partition Data section
 - Workload Activity report
- Monitor III:
 - CPC Capacity report
 - Enclave report
 - System information (SYSINFO) report

This new functionality is available as SPE and needs to be installed as APAR OA05731.

What's new in z/OS V1R4

Note:

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

State samples breakdown in the WLMGL report

Up to this release, state samples have been reported as a percentage of average transaction response time (response time breakdown). The response time is calculated when a transaction completes. This can result in percentages greater than 100 when samples are included for long running transactions which have not completed in the gathering interval (see also "Understanding Workload Activity data for IMS or CICS" on page 146).

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

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

What's new in z/OS V1R2

Software pricing online support

A new Monitor III **CPC Capacity** report is available which provides information about defined and consumed processor capacity of all partitions in a CPC (central processor complex).

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

Reporting of report class periods

Statistics on response times is a key performance metric used for groups of work having response time objectives. WLM report classes allow aggregation of performance data so that you can evaluate performance of applications independent of the definition how they are managed. Now, reporting of report classes is available with the same granularity as for service classes.

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

Enhanced reporting for the coupling facilities

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

What's new in z/OS V1R1

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

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

Intelligent Resource Director (IRD)

The Workload Manager is extended to work with PR/SM™ on zSeries 900 (z900) servers to dynamically expand resources that are available across LPARs.

An **LPAR cluster** is defined as the set of logical partitions in a single CEC that belong to the same parallel sysplex. Based on business goals, WLM can direct PR/SM to enable or disable CP capacity for an LPAR, without human intervention. This combination of WLM working with PR/SM on a z900 server is called IRD.

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

- **LPAR CPU Management**

Based on workload resource demand, the Workload Manager is able to dynamically adjust the number of logical processors and the weight of a logical partition. This allows the system to distribute the CPU resource in an LPAR cluster to partitions where the CPU demand is high.

The dynamic adjustment of processor resources within the partitions is reflected in the Postprocessor **CPU Activity (Partition Data)** report which provides LPAR views as well as aggregated views on LPAR cluster level.

Summary of changes

- **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 managed channels that are not related to a specific control unit. With the help of the Workload Manager, channels can float between control units to best service the work according to their goals and their importance.

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

- **Channel Subsystem Priority Queuing**

This topic is not reflected directly in any RMF report.

FICON® director

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

Goal mode enhancements

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

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

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

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

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

Chapter 1. Performance overview

Let's Start with an Overview

This chapter contains an overview of key performance topics. It explains the concepts and terms we will be using throughout the book:

- Performance Management Definition
- Service Level Agreements
- Capacity Planning Concepts
- Using RMF
- Components of Transaction Response Time
- Analyzing Workload Characteristics
 - Processor
 - Processor Storage
 - I/O Activity
- Processor Speed Indicators

Defining performance management

Performance management means monitoring and allocating data processing resources to applications according to *Service Level Agreements* (SLA) or informal objectives. It involves an ongoing cycle of measuring, planning, and modifying.

The goal of performance management is to make the best use of your current resources, to meet your current objectives without excessive tuning effort.

Setting performance goals

The human view of the performance of a system is often subjective, emotional and difficult to manage to. However, meeting the business needs of the users is the reason the system exists.

To match business needs with subjective perception, the concept of SLA was introduced.

The SLA is a contract that objectively describes such measurables as:

- Average transaction response time for network, I/O, CPU, or total
- The distribution of these response times (for example, 90% TSO trivial at less than 0.2 of a second)
- Transaction volumes
- System availability

A *transaction* is a business unit of work and can be a CICS® end user interaction, a batch job, etc. Ideally, the definition of a transaction is from a user's point of view.

Planning system capacity

We can define capacity planning as:

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

Capacity planning involves asking the following questions:

- How much of your computer resources are being used?
 - CPU
 - Processor storage
 - I/O
 - Network
- Which workloads are consuming the resources (workload distribution)?
- What are the expected growth rates?
- When will the demands on current resources impact service levels?

Benefits of capacity planning

An effective capacity planning process provides:

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

Common mistakes in capacity planning

Things to avoid include:

- **Forgetting the concept of balanced systems.** Remember that it takes a combination of resources to process work: CPU, processor storage, and I/O. Increasing the capacity of one resource by itself, may not allow more work to be processed.
- **Ignoring latent demand.** Latent demand is work that is 'hidden' and waiting to be unleashed. One way to begin quantifying latent demand is to note your peak-to-average ratio (e.g. peak hour CPU busy compared to prime-shift average CPU busy). This ratio will drop as latent demand builds up so track it and compare it to your current system (see *Balanced Systems and Capacity Planning* for more details). Relieving *any* bottleneck (for example, upgrading constrained CPU, increasing network bandwidth, or improving I/O subsystem) will release latent demand.
- **Being out of touch with business forecasts.** This one is obvious, but can be difficult to avoid. The capacity planner needs to stay informed of business plans that will impact IT resource requirements (for example, departmental growth, new applications, and mergers).
- **Ignoring complexity.** Factors other than transaction growth will increase demand on IT resources. Enhancements to existing applications, new government regulations, growing databases, etc. - all can result in an increase in resource consumption. This means that "the same" transaction running today may require 10-15% more resource than it did last year.

Performance management

- **Striving for total precision.** Many variables will impact your capacity plan (growth projections, resource capacity estimates, measurement precision, politics). A plan which achieves an accuracy of $\pm 10\%$ is considered good. Your time is probably better spent on reasonability checks on your input than on trying to achieve total accuracy.

Performance management versus capacity planning

These two disciplines have much in common. Performance management concentrates on allocating existing resources to meet service objectives. Capacity planning is a means of predicting resources needed to meet future service objectives.

Similar approaches can be used for both; here are some examples:

- Rules of thumb
- Comparison with other systems
- Parametric methods
 - Transaction profile (10 read calls, 2 update calls, 8 physical I/Os)
 - Cost of function (CICS: 15ms per physical I/O in a 3390 non-cached)
- Analytic (queuing theory) models: For example the IBM capacity planning tool CP90.
- Simulation, using a computer program that has the essential logic of the system: for example the Snap/Shot modelling system from IBM.
- Benchmarks, *Teleprocessing Network Simulator (TPNS)*.

In addition, performance and capacity work tend to require similar input data from the system. In particular, both require resource consumption by workload data (see “Analyzing workload characteristics” on page 9).

Using RMF

RMF issues reports about performance problems as they occur, so that your installation can take action before the problems become critical.

Your installation can use RMF to:

- Determine that your system is running smoothly
- Detect system bottlenecks caused by contention for resources
- Evaluate the service your installation provides to different groups of users
- Identify the workload delayed and the reason for the delay
- Monitor system failures, system stalls, and failures of selected applications

RMF comes with three monitors, Monitor I, II and III. Monitor III with its ability to determine the 'cause of delay' is where we start:

Monitor III provides short term data collection and online reports for continuous monitoring of system status and solving performance problems.

Monitor III is a good place to begin system tuning. It allows the system tuner to distinguish between delays for important jobs and delays for jobs that are not as important to overall system performance.

Monitor I provides long term data collection for system workload and resource utilization. The Monitor I session is continuous, and measures various areas of system activity over a long period of time.

You can get Monitor I reports directly as real-time reports for each completed interval (single-system reports only), or you let run the **Postprocessor** to create the reports, either as single-system or as sysplex reports. Many installations produce daily reports of RMF data for ongoing performance management. In this publication, sometimes a report is called a Monitor I report (for example, the Workload Activity report) although it can be created only by the Postprocessor.

Monitor II provides online measurements on demand for use in solving immediate problems.

A Monitor II session can be regarded as a snapshot session. Unlike the continuous Monitor I session, a Monitor II session generates a requested report from a single data sample. Since Monitor II is an ISPF application, you might use Monitor II and Monitor III simultaneously in split-screen mode to get different views of the performance of your system.

In addition, you can use the **Spreadsheet Reporter** for further processing the measurement data on a workstation by help of spreadsheet applications. The following chapters provide sample reports including the name of the corresponding macro.

You find a detailed description on how to create the reports and records and on how to use the macros in the *z/OS RMF User's Guide*.

There is another function in RMF to exploit the workstation for monitoring and analyzing your system, called **RMF Performance Monitoring (RMF PM)**, formerly known as PM of OS/390. You find a detailed description on how to use this function in the *z/OS RMF User's Guide*.

Performance management

This book will discuss key RMF indicators from the different monitors that can be used in a daily report and in problem diagnosis.

What sampling cycle and reporting interval should you use?

For Monitor III: Use the default sampling cycle of 1 second, with the default of 100 seconds for the reporting interval. Adjust this interval if needed to match a problem occurrence that you are investigating.

For Monitor I: Again use the default sampling cycle of 1 second. For the reporting interval, the default of 30 minutes is fine to start with. Adjust this if you prefer, or if you need to home in on a problem.

Analyzing transaction response time

To characterize the performance of a transaction you need to understand its different response time components.

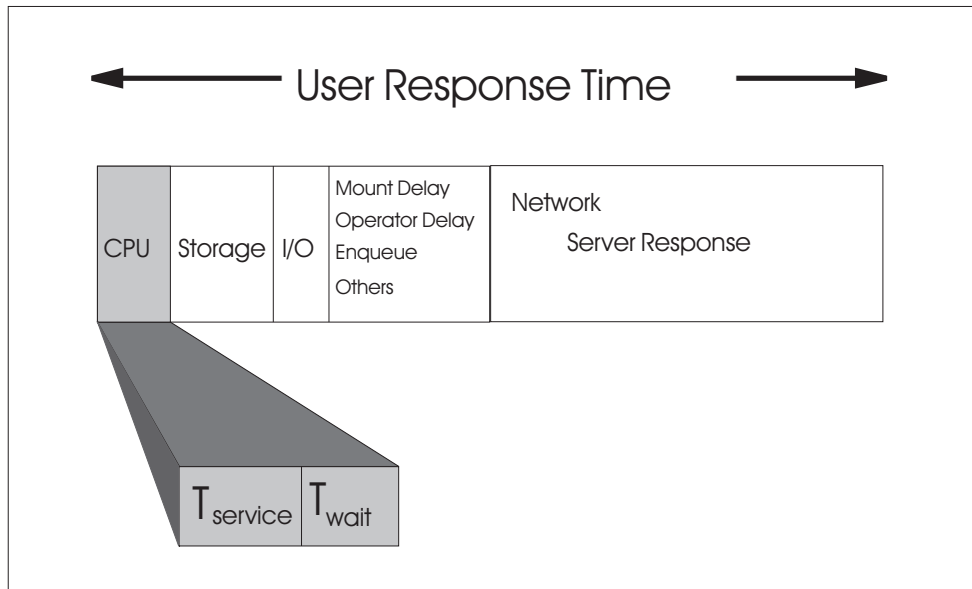


Figure 1. End-to-End Response Time Components

The reasons you should care about this are:

- For capacity planning you need to know resource consumption.
- For performance management you need to break down response time into components, to see where tuning can be done.

General formulas

Response time is made up of service time (the time actual work is done) and waiting time (the time waiting for resource): $T_r = T_s + T_w$

- T_r is Response Time
- T_s is Service Time
- T_w is Waiting Time (that is Queue time or Delay in RMF Monitor III).

Similarly total transaction service and wait times are made up of the individual resource service and wait times.

$$T_s = T_s(\text{CPU}) + T_s(\text{I/O}) + T_s(\text{TP}) + T_s(\text{Other})$$

$$T_w = T_w(\text{CPU}) + T_w(\text{I/O}) + T_w(\text{TP}) + T_w(\text{Storage}) + T_w(\text{Other})$$

The Monitor I Workload Activity report shows some of this data for a group of transactions (service or report class). This applies for TSO and batch work (and potentially CICS, with interval recording).

The Monitor III Group Response Time report also shows this data, with a more useful breakdown of response time components (see Figure 2 on page 8).

GROUP report

```

RMF V1R10 Group Response Time
Command ==>
Samples: 100      System: PRD1 Date: 04/07/08 Time: 10.32.00 Range: 100  Sec
Class: TSOPROD   Period: 1   Description: TSO Production
Primary Response Time Component: Storage delay for local paging

WFL  Users      Frames  Vector  EXCP  PgIn  TRANS  --- Response Time ---
%    TOT  ACT    %ACT    UTIL  Rate  Rate  Rate  -- Ended TRANS-(Sec) -
50   2    1     5      0     3.2  4.1   0.2   WAIT  EXECUT  ACTUAL
                                           0.000  3.432  3.432

                                           -AVG USG-  -----Average Delay-----
Average Users      Total  PROC  DEV  PROC  DEV  STOR  SUBS  OPER  ENQ  OTHER
Response Time ACT  3.432  0.57  0.97  0.11  0.17  1.25  0.00  0.00  0.00  0.34

                                           ---STOR Delay---  ---OUTR Swap Reason---  ---SUBS Delay---
                                           Page Swap  OUTR  TI  TO  LW  XS  JES  HSM  XCF
Average Users      0.19  0.03  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00
Response Time ACT  1.08  0.17  0.00  0.00  0.00  0.00  0.00  0.00  0.00  0.00
    
```

Figure 2. Monitor III Group Response Time Report

Report Analysis

See Response Time ACT field for a breakdown of the components of response time. Think of USING as service time (T_s) and DELAY as wait time (T_w)

- AVG USG: $0.57 + 0.97 = 1.54$ sec
- Average Delay: $0.11 + 0.17 + 1.25 + 0.34 = 1.87$ sec

So here, if you were investigating a response time problem, STOR would be the best starting point.

Monitor III provides insight into the components of response. With an emphasis on causes of *delay*, it shows you the resources for which work is being delayed and the address spaces which are holding the resources.

Analyzing workload characteristics

Much of performance work is done, not at the individual transaction level, but at a larger *workload* level. Understanding resource requirements by workload is key to effective performance management.

The reasons you should analyze your workload are:

- To understand your system's behavior
- For setting your SLA
- For tuning
 - Where is the pain?
 - Interactions with other workloads
- For input to the capacity planning process
 - Workload growth projections
 - Processor requirements
 - Storage requirements
 - DASD requirements
 - Network requirements

Identifying workloads

You will need to identify the different types of work in your system. Usually this is done at a service class level. To assign work to service classes you must classify the various workloads by their unique characteristics and requirements, that is:

- Response time needs
- Resource consumption (CPU, Storage, I/O)
- Priority
- Anticipated growth

Examples of workload identification at the service class level include:

- Trivial TSO (first period)
- Non-trivial TSO
- Batch
- Production CICS
- Development CICS
- IMS™
- Graphics application

Ideally, you may want to take workload differentiation one level further, matching workloads to true business functions (for example, claims processing and order fulfilment). This may require more detailed data from SMF records.

RMF reports data about different workloads grouped into categories which you have defined as *workloads* and *service classes* in your service policy.

The appropriate grouping of workloads is important.

- If you have different applications which should be managed according to the same goals, you should define the same *service class* for them. Applications with different goals need to be assigned to different service classes.

Workload characteristics

- If you want to get separate reporting for different applications in the same service class, you can define separate *report classes* for each of them. Reporting for report classes is possible with the same level of detail (report class period) as for service classes.

Measuring resource utilization by workload

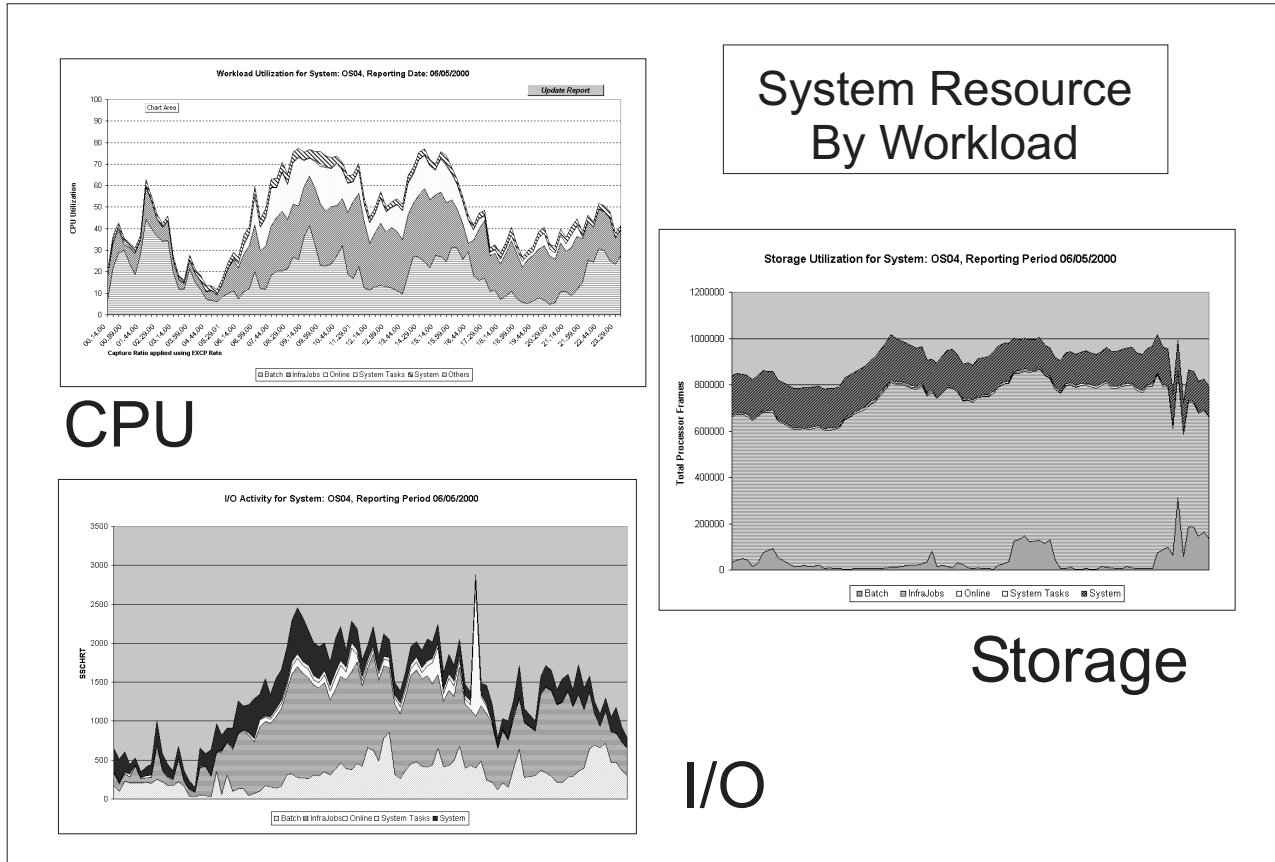


Figure 3. System Resource Summary by Workload

This figure illustrates one way to report your resource utilization by workload.

You can use the Spreadsheet Reporter to create spreadsheet files from Postprocessor data and to display them graphically. The following chapters show how to use the sample macros for creating the graphics that help you in understanding and analyzing the performance of your system.

In a similar fashion (could be graphical or numeric) you will need to record such data to:

- Measure resource consumption:
 - CPU utilization
 - Processor storage usage (*central storage (CS)* and *expanded storage (ES)*)
 - I/O rates
- Understand what makes up response time:
 - Waiting for and using the resources
 - Use Monitor III Group Response Time report
- Understand the factors that influence the above

Virtual storage is another type of resource. While not as dynamic as I/O, CPU, or processor storage, virtual storage is one of the basic resources that you need to plan, even though you cannot buy it. Virtual storage can have a critical impact on delivery of service to end users: an *initial program load (IPL)* is required when not enough virtual storage is available.

RMF can measure virtual storage. It should be monitored, and the Virtual Storage report helps you verify the size of the different MVS components, as explained in “Virtual Storage Activity report” on page 57.

Using Monitor I reports to analyze workloads

The CPU Activity report and the Workload Activity report provide the data you can use for a workload analysis.

CPU Activity report - part 1

CPU ACTIVITY									
z/OS VIRI0		SYSTEM ID OS04		DATE 03/14/2008		INTERVAL 14.59.996		PAGE 1	
RPT VERSION VIRI0 RMF		TIME 09.29.00		CYCLE 1.000 SECONDS					
CPU 2097	MODEL 720	H/W MODEL E26	SEQUENCE CODE 00000000000473BF	HIPERDISPATCH=NO					
---CPU---	ONLINE	LPAR BUSY	TIME %	MVS BUSY	PARKED	LOG PROC SHARE %	--I/O RATE	INTERRUPTS--	% VIA TPI
0	CP	100.00	76.93	99.15	0.00	35.0	320.9	0.33	
1	CP	100.00	76.85	99.13	0.00	35.0	324.0	0.22	
2	CP	100.00	76.86	99.09	0.00	35.0	321.6	0.22	
3	CP	100.00	76.87	98.99	0.00	35.0	329.5	0.25	
4	CP	100.00	76.89	99.01	0.00	35.0	325.9	0.33	
5	CP	100.00	76.87	98.88	0.00	35.0	328.7	0.36	
6	CP	100.00	76.88	98.86	0.00	35.0	338.4	0.46	
7	CP	100.00	76.85	98.73	0.00	35.0	341.1	0.51	
8	CP	100.00	76.82	98.59	0.00	35.0	335.9	0.67	
CP TOTAL/AVERAGE			76.87	98.94		315.0	2966		

Figure 4. Monitor I CPU Activity Report - Part 1

The first part of this report provides an overview on all processors belonging to this system.

Workload characteristics

Workload Activity report

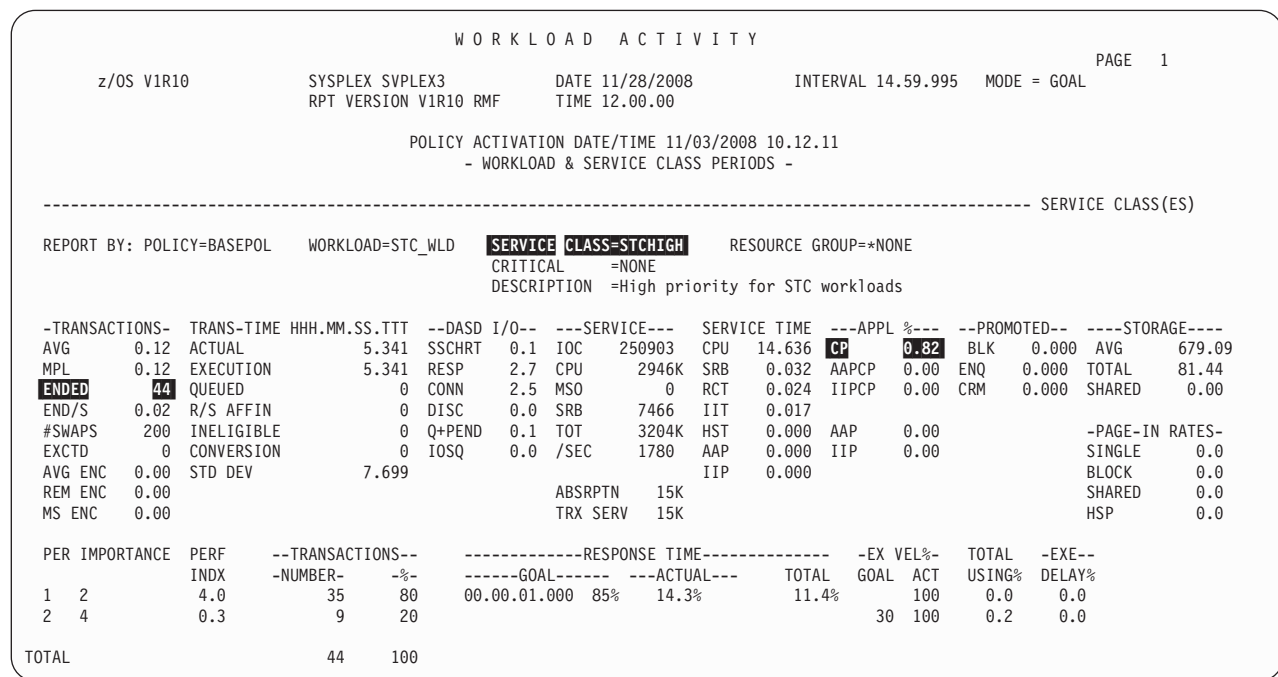


Figure 5. Workload Activity Report - Service Class

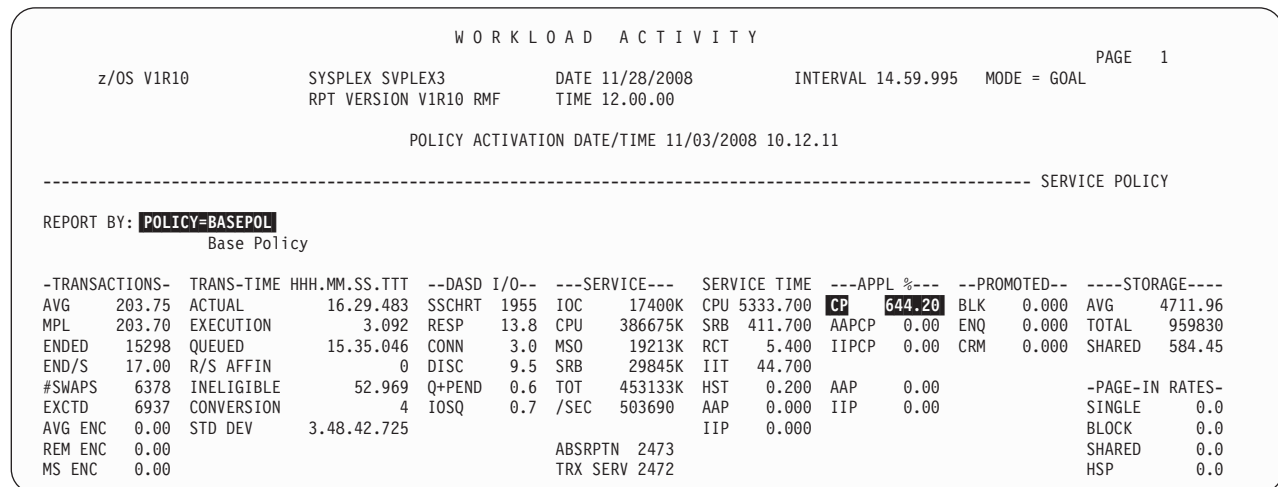


Figure 6. Workload Activity Report - Service Policy

The Workload Activity report shows workload-related data on resource consumption with different levels of detail. The above examples summarize data for one service class (STCHIGH) and for the total system (service policy).

Analyzing processor characteristics

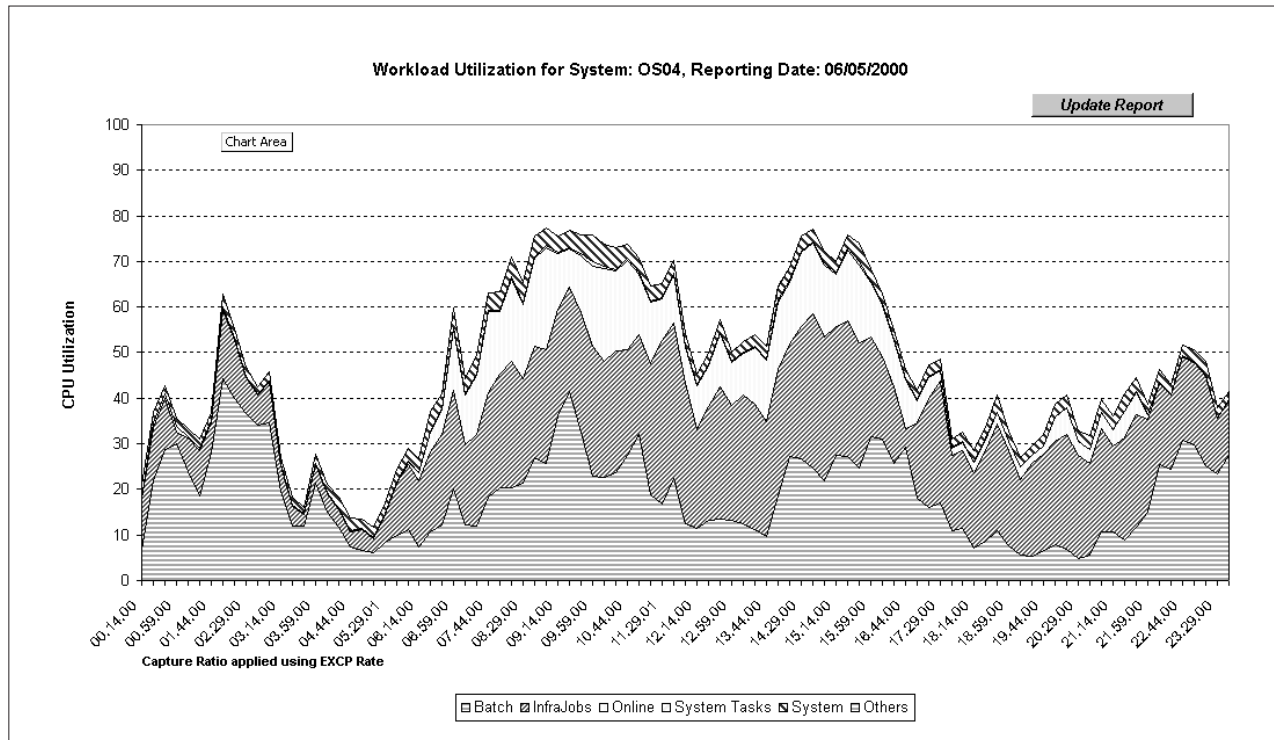


Figure 7. Processor Utilization by Workload. Spreadsheet Reporter macro RMFY9WKL.XLS (Workload Trend Report) - Day Utilization

This information unit discusses how to measure CPU utilization by workload. Our approach is to convert *application time* for each service class into overall percent of CPU used for that workload. The total of all these numbers is the base for the APPL% CP value in the report.

The concept of *capture ratio (CR)* must be understood in order to do this. Simply put, the CPU time reported for the sum of all your workloads never adds up to the total CPU time used by the system. The total CPU time reported for all your workloads will typically account for 85-90% of the total CPU time used. This is not an error. It is the best level of accuracy that the reporting tools can achieve, *in a consistent, repeatable manner*.

The *uncaptured* time is sometimes, misleadingly, called system overhead. This is incorrect, because most of it is genuine user work — it is either a political or philosophical view whether activities like paging are seen as genuine work or just overhead.

Thus the capture ratio is:

$$CR = \frac{\text{Captured CPU Time}}{\text{Total CPU Time}}$$

The question now is: how do you account for the uncaptured, though real, CPU time? The answer is to distribute the uncaptured time among your workloads, so that the total CPU time is accounted for. Several approaches how to do this are discussed.

How to compute CPU time per workload type

Following are the basic steps for computing CPU time by workload. A graphical view of CPU consumption over time (for example, by hour or by day) can help you to understand your workload flow, as shown in Figure 7 on page 13 as an example of a Spreadsheet Reporter macro.

1. Calculate total CPU time in seconds:

$$\text{Total CPU time} = \text{MVS Busy \%} * \text{Interval} * \text{\#CPs}$$

If your system is running in an PR/SM environment, the calculation has to be performed using LPAR BUSY and the number of logical processors.

2. Get the captured time for all address spaces.

The Workload Activity report shows the percentage of application time. It is the sum of *task control block and preemptible-class SRB (CPU) time, non-preemptible service request block (SRB) time, region control task (RCT) time, I/O interrupt (IIT) time, and hiperspace service (HST) time*. Subtract zAAP and zIIP time, because it is contained in CPU time.

$$\text{APPL\% CP} = \frac{(\text{CPU} + \text{SRB} + \text{RCT} + \text{IIT} + \text{HST} - \text{AAP} - \text{IIP})}{\text{Interval}} * 100$$

$$\text{Captured time} = (\text{APPL\% CP} / 100) * \text{Interval}$$

This results in the uncaptured CPU time:

$$\text{Uncaptured CPU time} = \text{Total CPU time} - \text{Captured time}$$

3. Calculate the capture ratio:

$$\text{CR} = \frac{\text{Captured CPU Time}}{\text{Total CPU Time}}$$

4. Get application (captured) percentage for one workload you are interested in:

This APPL% CP value represents the percentage based on one processor. You get the percentage based on the total capacity by:

$$\text{Total Captured \%} = \frac{\text{APPL\% CP}}{\text{\#CPs}}$$

5. Distribute uncaptured time based on the capture ratio:

$$\text{Total \%} = \frac{\text{Total Captured \%}}{\text{CR}}$$

There are more ways to distribute the uncaptured time, see “Methods of distributing uncaptured time” on page 15 for a discussion of this topic.

Report Analysis

Figure 4, Figure 5, and Figure 6 give us the following values for standard processors:

- Total CPU time = $0.7687 \times 900 \times 9 = 6226.47$ sec
- Captured time (POLICY) = $6.442 \times 900 = 5797.8$ sec
- Uncaptured time = $6226.47 - 5797.8 = 428.67$ sec
- Capture ratio = $5797.8 / 6226.47 = 0.93$
- DB2 (Service Class SCONL006) = 64.40 % of one processor

Figure 7 is a diagram being created with the Spreadsheet Reporter, it shows the workload utilization for different service classes during one day. The macro provides different ways to distribute the uncaptured time.

Methods of distributing uncaptured time

To get a full picture of CPU consumption, the uncaptured time needs to be distributed across the workloads. There are several methods of doing this:

1. Not bad: By workload APPL time - as done in the previous calculation.
2. Better: By workload I/O rate (% of system total)
3. Best:
 - a. Distribute uncaptured time by I/O rate, AND...
 - b. Distribute system address space (AS) time to relevant users, e.g.
 - VTAM[®] time: to TSO and CICS
 - JES time: to TSO and BATCH
 - etc.
4. Probably not worth the effort: Regression analysis, etc.

How to compute CPU time per transaction

There are times when additional detail of CPU utilization is needed. For capacity planning especially, you may need to know the average CPU time consumed per transaction. This is a simple calculation, dividing application time by number of transactions.

Report Analysis

Figure 5 on page 12 gives us the following values for DB2 (Service Class SCONL006):

- ENDED TRANSACTIONS: 7288
- APPL% CP * INTERVAL : 581.4 sec (Application time)

This results in

$$\text{CPU time per transaction} = \frac{581400}{7288} = 80 \text{ msec}$$

Workload characteristics

Analyzing workload I/O characteristics

It can also be useful to know the I/O rate each different workload generates. This is especially useful in capacity planning, but also for performance management - a change in workload I/O rates (or relative I/O content) can significantly change its performance.

Base for the following graphic are these OVW statements:

```
OVW(IOBATCH(SSCHRT(W.WLBATCH)))  
OVW(IOINFRA(SSCHRT(W.WLINFRA)))  
OVW(IOONLIN(SSCHRT(W.WLONLINE)))  
OVW(IOSYSTK(SSCHRT(W.WLSYSTEM)))  
OVW(IOSYSTEM(SSCHRT(W.SYSTEM)))
```

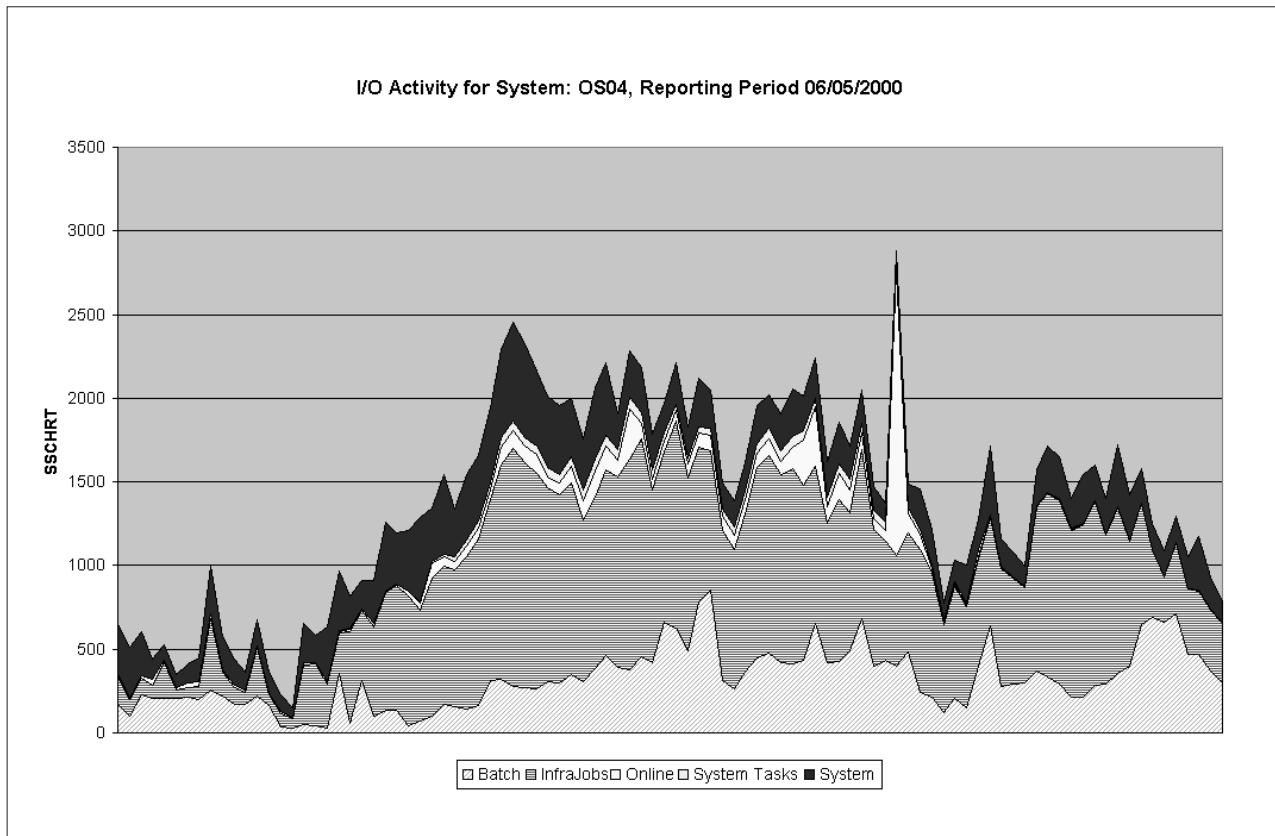


Figure 8. I/O Activity by Workload

```

REPORT BY: POLICY=SP000010  WORKLOAD=WONLINE  SERVICE CLASS=SCONL006  RESOURCE GROUP=*NONE
CRITICAL =NONE
DESCRIPTION =DB2 Production

TRANSACTIONS  TRANS.-TIME  HHH.MM.SS.TTT  --DASD I/O--  ----SERVICE----  --SERVICE TIMES--  ---APPL %---  ----STORAGE-----
AVG 28.23  ACTUAL 769  SSCHRT 41.4  IOC 0  CPU 581.600  CP 64.40  AVG 0.00
MPL 28.23  EXECUTION 768  RESP 2.9  CPU 42164K  SRB 0.000  AAPCP 0.00  TOTAL 0.00
ENDED 7288  QUEUED 0  CONN 2.4  MSO 0  RCT 0.000  IIPCP 0.00  SHARED 0.00
END/SEC 8.10  R/S AFFIN 0  DISC 0.0  SRB 0  IIT 0.000
#SWAPS 0  INELIGIBLE 0  Q+PEND 0.5  TOT 42164K  HST 0.000  AAP 0.00  --PAGE-IN RATES--
EXCTD 0  CONVERSION 0  IOSQ 0.0  /SEC 46869  AAP 0.000  IIP 0.00  SINGLE 0.0
AVG ENC 0.00  STD DEV 1.918  ABSRPTN 1660  IIP 0.000  BLOCK 0.0
REM ENC 0.00  TRX SERV 1660  PROMOTED 0.000  SHARED 0.0
MS ENC 0.00  HSP 0.0

REPORT BY: POLICY=SP000010
Default Policy

TRANSACTIONS  TRANS.-TIME  HHH.MM.SS.TTT  --DASD I/O--  ----SERVICE----  --SERVICE TIMES--  ---APPL %---  ----STORAGE-----
AVG 203.75  ACTUAL 16.29.483  SSCHRT 1955  IOC 17400K  CPU 5333.700  CP 644.20  AVG 4711.96
MPL 203.70  EXECUTION 3.092  RESP 13.8  CPU 386675K  SRB 411.700  AAPCP 0.00  TOTAL 959830
ENDED 15298  QUEUED 15.35.046  CONN 3.0  MSO 19213K  RCT 5.400  IIPCP 0.00  SHARED 584.45
END/SEC 17.00  R/S AFFIN 0  DISC 9.5  SRB 29845K  IIT 44.700
#SWAPS 6378  INELIGIBLE 52.969  Q+PEND 0.6  TOT 453133K  HST 0.200  AAP 0.00  --PAGE-IN RATES--
EXCTD 6937  CONVERSION 4  IOSQ 0.7  /SEC 503690  AAP 0.000  IIP 0.00  SINGLE 0.0
AVG ENC 0.00  STD DEV 3.48.42.725  ABSRPTN 2473  IIP 0.000  BLOCK 0.0
REM ENC 0.00  TRX SERV 2472  PROMOTED 0.000  SHARED 0.0
MS ENC 0.00  HSP 0.0
    
```

Figure 9. Workload Activity Report - I/O Activity

The Workload Activity report directly shows the I/O activities by workload in the value SSCHRT. This is the number of start subchannels (SSCH) per second and gives the number of DASD non-paging I/Os.

Report Analysis

We see the following numbers in the sample report:

Workload	SSCH/second
DB2 Production	41.4
Policy - Total System	1955

Analyzing processor storage characteristics

As with CPU and I/O, it is useful to build a clear picture of processor storage use on your system. The Workload Activity report can give you a good start on the processor storage use by workload, at the workload level.

The Spreadsheet Reporter does not contain a macro for displaying storage utilization by workload. But you can create very easily a graphic as shown in the following example. Assuming that you have the workloads WLBATCH, WLINFRA, WLONLINE, and WLSYSTEM, you can use the following OVW statements to create the required spreadsheet data:

```
OVW(TOTBATCH(STOTOT(W.WLBATCH)))  
OVW(TOTINFRA(STOTOT(W.WLINFRA)))  
OVW(TOTONLIN(STOTOT(W.WLONLINE)))  
OVW(TOTSYSK(STOTOT(W.WLSYSTEM)))  
OVW(TOTSYSTEM(STOTOT(W.SYSTEM)))
```

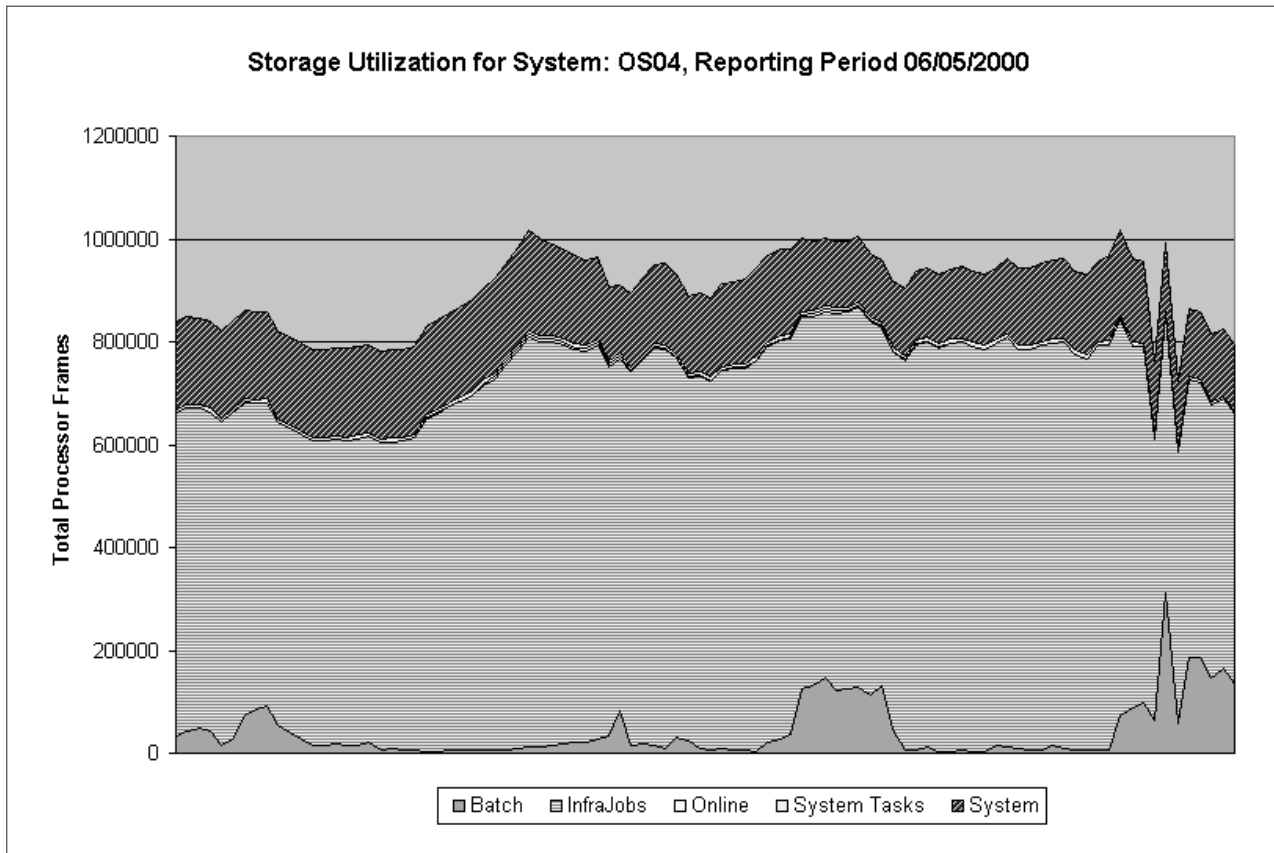


Figure 10. Processor Storage by Workload

Of course, you can display this in more detail by using other OVW conditions, for example STOCEN and STOEXP instead of STOTOT. You find a detailed description of all OVW conditions in chapter "Overview and Exception Conditions" in the z/OS RMF User's Guide.

Workload Activity report

REPORT BY: POLICY=SP000010 WORKLOAD=OMVS_WLD SERVICE CLASS=OMVLOW RESOURCE GROUP=*NONE															
CRITICAL =NONE															
DESCRIPTION =Low priority for OMVS workloads															
-TRANSACTIONS-	TRANS-TIME	HHH.MM.SS.TTT	--DASD	I/O--	---	SERVICE----	---	SERVICE	TIMES--	---	APPL	%---	---	STORAGE-----	
AVG	3.47	ACTUAL	4.912	SSCHRT	1.4	IOC	345125K	CPU	3.736	CP	0.44			AVG	2068.48
MPL	3.47	EXECUTION	4.911	RESP	70.0	CPU	6951K	SRB	0.086	AAPCP	0.00			TOTAL	7185.56
ENDED	135	QUEUED	1	CONN	4.3	MSO	6470	RCT	0.119	IIPCP	0.00			SHARED	1.27
END/S	0.15	R/S AFFIN	0	DISC	0.9	SRB	163	IIT	0.023						
#SWAPS	306	INELIGIBLE	0	Q+PEND	64.9	TOT	352083K	HST	0.000	AAP	0.00			--PAGE-IN	RATES--
EXCTD	0	CONVERSION	0	IOSQ	0.0	/SEC	391203	AAP	0.000	IIP	0.00			SINGLE	0.0
AVG ENC	0.00	STD DEV	11.041					IIP	0.003					BLOCK	0.0
REM ENC	0.00							ABSRPTN	113K					SHARED	0.0
MS ENC	0.00							TRX SERV	113K	PROMOTED	0.000			HSP	0.0

Figure 11. Workload Activity Report - Processor Storage Use

Processor Storage Use is given in the fields:

AVG # of processor frames on average per swapped-in AS in the group.

TOTAL # of processor frames on average for all AS in the group.

TOTAL = AVG * MPL

If your system is running on a zSeries processor, all fields in the RMF reports which are related to expanded storage will be empty. Expanded storage is no longer supported in the zArchitecture.

This workload use of processor storage will account for most of the processor storage on your system. To account for the rest, you should also look at:

- System area use of storage (see Monitor I Paging Activity report)
- Available frames (see Monitor I Paging Activity report)
- Swapped-out users (both in CS & ES - see Monitor II ASD report)

It can also be useful to know how much paging your individual workloads are doing. From the Workload Activity report, you also get paging data from the fields:

SINGLE Page-ins per second from DASD

BLOCK Page-ins per second from DASD, for blocked pages

Use this paging information to help decide whether a workload's current processor storage is sufficient (see Chapter 5, "Analyzing processor storage activity" for more detail on paging).

Monitor II and III also provide processor storage use (and paging) data, down to the AS level.

Processor speed indicators

This information unit presents some of the terms and formulas used in discussing CPU performance and capacity.

- **Cycle Time**

$\text{CPU time} = \text{Path length} * (\text{Cycles} / \text{Instruction}) * \text{Cycle time}$

- Path length depends on the instruction set and the compiler
- Cycles/instruction depends on the design (e.g. microcoding of instructions)
- Cycle time depends on the technology (for example, CPU model)

Thus, cycle time by itself is not a complete indicator of CPU speed.

- **MIPS**

This term depends on cycle time and number of cycles per instructions, it is only valid for comparison if the instruction sets are the same.

In common usage today, MIPS are used as a single number to reflect relative processor capacity. As with any single-number capacity indicator, your experience may vary considerably. These numbers are often based on vendor announcements.

By the way, *MIPS* means *Millions of Instructions Per Second* (some people say *Misleading Information about Processor Speed*). There does not exist a term *MIP* which often is used erroneously. A small processor has a speed of 1 MIPS, not 1 MIP.

- **SRM Constant — MSU**

The SRM constant is a number derived by product engineering to normalize the speed of a CPU, so that a given amount of work would report the same service unit consumption on different processors.

Service Units (SU) are typically used by the *System Resource Manager* (SRM) as a basis for resource management.

$\text{Service Units} = \text{CPU seconds} * \text{SRM constant}$

Many installations who have to charge the users do so on the basis of SUs. This is a reasonable approximation of relative capacity. But of course, *any* single-number metric for processor capacity comparison can be misleading, and LSPR numbers, based on your own workload mix, are the best to use. You get more details about this in the following section that describes ITR and LSPR.

The SRM constants are contained in internal tables within MVS and are published in the *z/OS MVS Initialization and Tuning Guide*. The SUs are reported in many Monitor I and Monitor III reports.

The power of a system can be characterized also in terms of MSUs (millions of service units per hour). These numbers are published in the LSPR documentation.

- **Internal Throughput Rate (ITR)**

ITR numbers are measurements by workload type which determine the capability of a machine in terms of the number of transactions per CPU second. This is important, since processor capacity varies according to the workload mix.

$$\text{ITR} = \frac{\text{Number of transactions}}{\text{CPU time}}$$

ITR is a function of:

- CPU speed
- Operating system
- Transaction characteristics (workloads differ)

ITRs are derived from the LSPR methodology. See *Large Systems Performance Reference* for details. ITR provides a reliable basis for measuring capacity and performance.

- **External Throughput Rate (ETR)**

$$\text{ETR} = \frac{\text{Number of transactions}}{\text{Elapsed time}} = \text{ITR} * \text{CPU utilization}$$

Any bottleneck, internal or external to the system, will effect the ETR. Examples include: I/O constraints, tape mount delay and paging delay. Thus it is more difficult to get a repeatable measure than with ITR.

- **Relative Processor Power (RPP)**

RPP is the ratio of ITRs for a specific workload mix (usually 25% Batch, 25% TSO, 25% CICS, and 25% IMS) for different machines. RPP is normalized to a base machine.

So, if a given function takes a certain amount of RPPs, you can estimate how much it would consume on a different machine by using the formula:

$$\text{Your utilization} = \frac{\text{Used RPPs} * 100}{\text{Your machine RPPs}}$$

- **MFLOPS**

MFLOPS = millions of floating point instructions per second. Used only in numerically intensive computing (scientific/technical).

In summary, various numbers may commonly be used as rough indicators of CPU capacity. Remember that most of these are very rough approximations only, and your actual capacity may vary significantly.

To get the most accurate assessment of CPU capacity for your workload, differences in workload processing characteristics must be taken into account using a methodology such as that described in the *Large Systems Performance Reference*. You find the most current version via

<http://www.ibm.com/servers/eserver/zseries/lspr/zSeries.html>

Where do you go from here?

Having reviewed some of the main MVS performance concepts, the following chapters will discuss how to recognize and resolve specific performance problems.

One thought to bear in mind on your journey:

You can only do so much as a system programmer; you need help from your application designers and data base designers. A poorly written application can nullify a lot of system tuning effort. Participate in the design reviews, where you can stress that you need reasonable call structures to data bases, practical data base design etc.

Chapter 2. Diagnosing a problem: The first steps

Let's Start the Diagnosis

This chapter explains the first steps on your way to performance management:

- How to recognize a performance problem
- How to find that system in the sysplex that contributes most to the problem
- How to find the major cause of the problem (CPU, processor storage, I/O, etc.).

Generally, there are two different approaches:

- One for response time problems (primarily using Monitor III delay reports)
- And one for system indicators showing stress (primarily using Monitor I/II resource-usage reports)
- How to continue by reading the other chapters in this book for further problem analysis and resolution

What is a performance problem?

There are many views on what constitutes a performance problem. Most of them revolve around unacceptably high response times or resource usage, which we can collectively refer to as "pain." The need for performance investigation and analysis is for example indicated by:

- Bad or erratic response time
 - Service level objectives are being exceeded
 - Users complaining about slow response
 - Unexpected changes in response times or resource utilizations
- Other indicators showing stress
 - Monitor III Workflow/Exceptions
 - System resource indicators (for example, paging rates, DASD response)
 - Expected throughput on the system not being attained

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

Generally, a performance problem is the result of some workload not getting the resource(s) it needs to complete in a timely manner. Or, less common, the resource is obtained but is not fast enough to provide the desired response time.

There are three ways to gain the resource needed:

- Buy it.
- Create the illusion you bought it. This is known as tuning. Capability to do this implies you have been previously wasting resources. Like purchasing, there is a cost. The cost is people; while it may be higher, it is always less visible than purchase.
- Steal it (take it from a less important application). Again there is a cost. Here the cost is lower service to the application from which the resources were stolen.

If none of these options are technically or financially possible, it will be necessary to change users' (and management's) expectations.

You may have experienced the situation where you complete an extensive performance analysis only to conclude that no further tuning or stealing of resources can be done. One of the goals of this book is to assist you in determining whether you have reached this point.

Getting started: A top-down approach to tuning

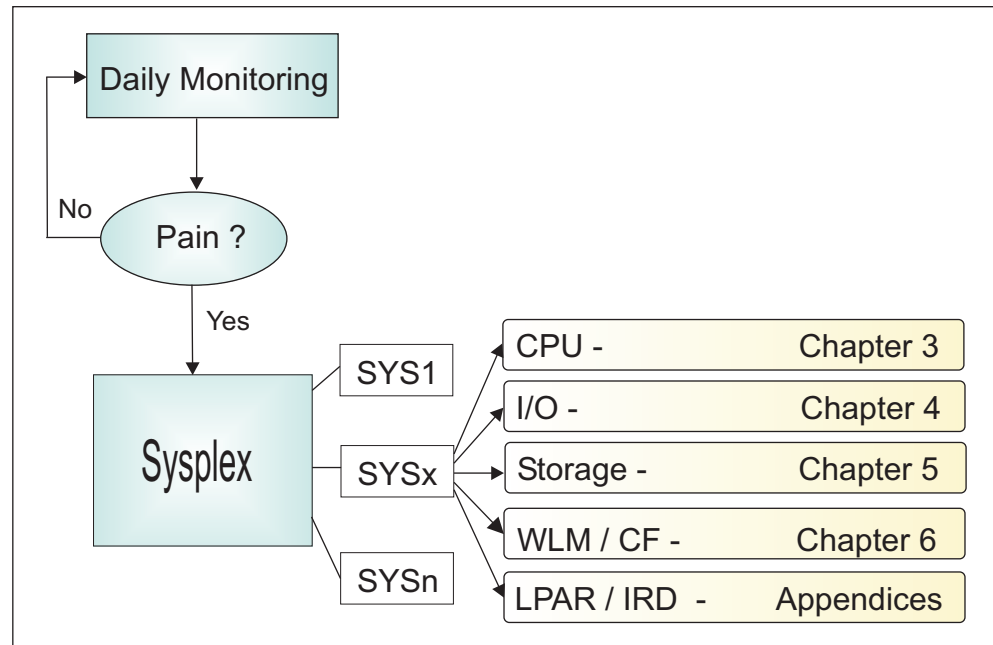


Figure 12. Simplified View of Performance Management

This book will take the following top-down approach to investigating and resolving performance problems. Figure 12 is a simplified view:

1. Through your daily performance monitoring activities recognize any of the "pain" indicators as being a problem on your sysplex.
2. Discover the system where you can see the most important performance problems.
3. Analyze the data to determine which component of the response time holds the most promise for relief.
4. Go to the appropriate chapter or appendix for more detailed advice.

Daily performance monitoring — what does this mean? If you ask ten people, probably you will get eleven answers. Or vice versa, there is not just one answer to the question: *How should I do daily performance monitoring?* Having in mind customers who want to be satisfied with the service that they receive from the system you are responsible for, the focus should be on observing the service level agreements. If the goals being described there can be attained — then you are providing the service that your customers are expecting. In the past, it was not a simple task of monitoring the response time and throughput objectives that are described in the service level agreements. And furthermore, it was not a simple task to adjust all parameters in the Parmlib members IEAIPSxx and IEAICSxx to reach the optimal performance of your system.

But now, it's easy and simple: you just define your goals in the service policy and let the system run. The workload manager is doing its best in managing the workload in your system for achieving all goals — your job is done automatically, be happy. STOP — if this would be true we could stop at this point of time. Of course, it's not true because there are many aspects which have to be mentioned with regard to performance management.

Performance analysis

Starting with the sysplex view, you get a high-level overview about the performance indicators of your workloads. This will lead you to the most relevant systems in the sysplex. There, you may start with the largest cause of delay, do what you can to address it, go back and address the next largest cause of delay, and so on. Continue to do this until the service objectives are met, or until you reach a point of diminishing returns, where further effort does not pay off.

In short ...

... use the data available from RMF to ensure that your efforts are focused where they can do the most good.

Of course, with any performance problem there are many different approaches you can take to analyze the data. We have selected these two paths:

- **Response time problems.** Follow this path if you have a user or group of users complaining about response time.
- **System indicators.** Follow this path if your ongoing performance monitoring shows signs of a problem, or for a general health check. This path starts on page 41.

Where to start sysplex monitoring?

When using RMF, you are interested in getting a performance overview at a glance. You have defined a service policy that contains several workloads, each of them with service classes and report classes that have goals of different types.

Monitor III and the Postprocessor offer you some reports that show how your sysplex is running at different levels of detail.

To start with:

Monitor III Sysplex Summary report

This report shows (based on your option selection) all or only selected workloads, service and report classes, and periods with their goals and actual performance values.

When you run Monitor III in GO mode, you see the *Performance Status Line* that summarizes the key performance indicators for up to 80 time ranges, showing you the performance history of your sysplex.

To continue with:

Monitor III Response Time Distribution report

You can use this report to see how several systems of the sysplex contribute in servicing one specific service class. This is the entry point from the sysplex to single-system reports, for example:

GROUP	Group Response Time Report
STOR	Storage Delays Report
DELAY	Delay Report
JOB	Job Delays Report
WFEX	Workflow/Exceptions Report

Monitor III Work Manager Delays report

Here, RMF provides detailed performance data about the CICS and IMS subsystems. The Work Manager Delays report shows you for CICS and IMS several kinds of delay data about your transactions, being in the begin-to-end or in the execution phase. You get an overview of how the different CICS address spaces (for example, the AOR, TOR, or FOR) provide service to the service classes your transactions belong to.

Postprocessor Workload Activity report

You can use this report to get performance and resource data for your service classes and workloads in the sysplex. You will find a detailed discussion including report samples in “Understanding Workload Activity data for IMS or CICS” on page 146.

Monitor III indicators

Use the following reports for monitoring performance of your sysplex:

- “Sysplex Summary report”
- “Response Time Distribution report” on page 32
- “Work Manager Delays report” on page 34

Sysplex Summary report

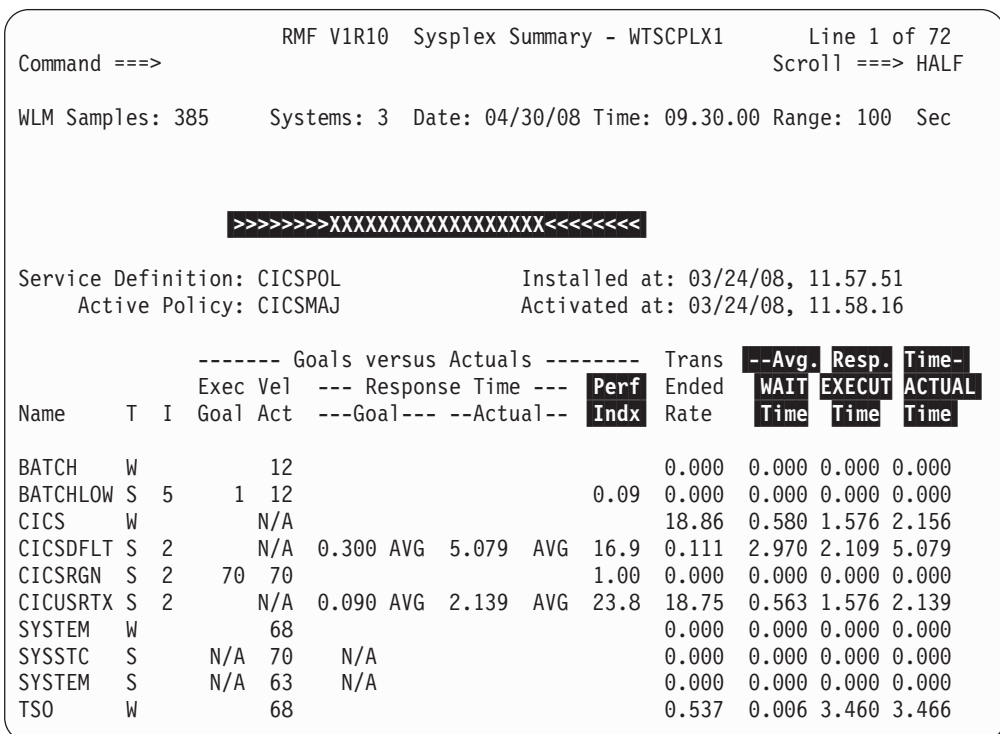


Figure 13. Monitor III Sysplex Summary Report

You can select what should be part of this report by using the report options:

- *What types should be reported?* These can be everything from workload to service class period, as well as report classes.
- *What performance index should be used as threshold?* This allows to show only the data that have an index high enough for you to want to be informed about it.
- *Is any importance important?* If you want to see only high or medium important work - you can select this.
- *Are you interested in all workload groups you have defined?* Then you can display them - otherwise the report shows only active workload groups and their details.

Performance Status Line in the Sysplex Summary Report




Figure 14. Monitor III Sysplex Summary Report - GO Mode

Indicator

Field: `-----|||---X--XX-XXXX-XX-XX-|---XXXXX---|---|||-----XX-`

Description: The performance status line gives a performance indication, by its symbols and colors, for each range when the Monitor III reporter session was in GO mode:

- Green** All goals have been attained for this range
- Yellow** Service class periods with low or medium importance have not attained the goal.
- Red** Service class periods with high importance have not attained the goal.

If you switch from GO mode to STOP mode, the reporter builds no reports for this time, and therefore, no indicator for the status line will be created. When continuing with GO mode, you will see one or several blank fields in the status line.

The status line is updated with each refresh of the report in GO mode. Assuming a range and refresh value of 100 seconds, you will see the *full* status line after 8000 seconds with each interval indicator shifted to the left by one position.

You can take this status line to see at a glance the performance status of your sysplex during the previous minutes or hours.

If green is the dominant color, your sysplex seems to be in good shape, as more yellow or red indicators appear, you might start investigating the sources of possible problems.

Indicator

Field: Perf Indx

Description: The performance index shows how well a performance goal could be achieved and is calculated from the goal and actual performance data.

Guideline: A value of 1 means that the actual value met the goal exactly, a lesser value indicates that the actual value is better than the goal, a higher value could be seen as an indication of a performance problem.

If the goal could not be attained, the lines for the service class period, service class and workload are displayed in red (for high importance) and yellow (for medium and low importance). The same color is also given to the corresponding field in the performance status line.

Problem Area: The red and yellow lines can be indicators of performance problems in the sysplex. If it is a one-time event, you might ignore it. If some lines show red continuously, further investigation is recommended.

Potential Solution: Depending on the type of the workload you want to study, you can proceed with the Response Time Distribution report to get an impression of how this service class performs on the different systems in the sysplex, or you might choose the Work Manager Delays report that shows you more about CICS and IMS delays.

Note: The calculation of the performance index is done with actual values that are averages for the reporting range (could be, for example, 100 seconds for Monitor III, or 30 minutes in the Workload Activity report).

You might distinguish between this value and the performance index that is used by the workload manager internally to manage all service classes. This index is calculated every 10 seconds and you find the values in SMF records type 99. Due to the different ranges it is obvious that a difference will be seen.

Indicator**Field:** Avg. Resp. Time**Description:** The response time information for all workloads and service classes is given in three fields:

1. ACTUAL Time

The average response time for all ended transactions.

Note that 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, you can't see the problem on this report until the locked out transactions end.

2. EXECUT Time

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.

Note: In the Postprocessor Workload Activity report, you see this field as EXECUTION TIME.

3. WAIT Time

This time is calculated as the difference between ACTUAL and EXECUT time, as long as ACTUAL time is the bigger value. However, for subsystem data, it can happen that EXECUT time is more than the ACTUAL time.

For CICS transactions, this includes not only queuing in the TOR and AOR, but also processing time within the TOR.

For IMS transactions, this includes not only queuing for the MPR, but also processing time within the CTL region.

Otherwise, this is the average time that transactions spent waiting on a JES or APPC queue. Note that WAIT 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.

The "server" service classes are blank in the Avg. Resp. Time and Trans Ended Rate columns, because their "transactions" are address spaces, and response times are available only for ended transactions.

Problem Area / Potential Solution: See chapter "Understanding Workload Activity data for IMS or CICS" on page 146.

Response Time Distribution report

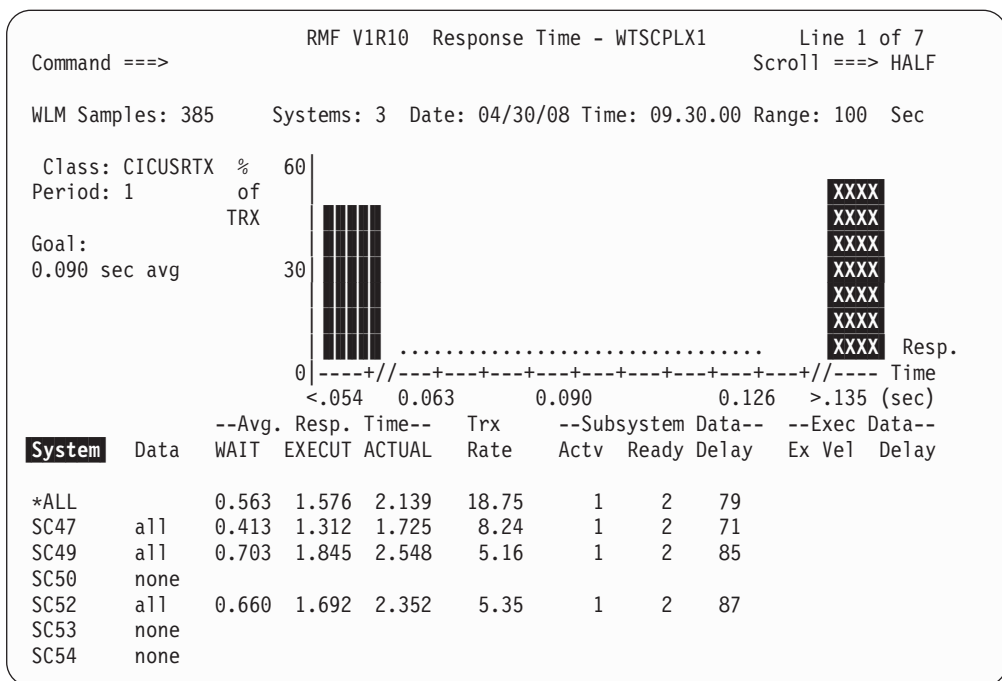


Figure 15. Monitor III Response Time Distribution report

This report enables you to analyze the distribution of a response time to see whether a response time goal was met and, if not, how close it came to failing. The report shows how the response time for a specific service or report class period is distributed. Two levels of detail are shown:

- A character graphic shows the distribution of response time for all systems in a sysplex which have data available in this period.
This graphic is shown only for a periods with a response time goal.
- The details of how each system contributed to the overall response time.

Using cursor-sensitive control, you can navigate from this report to the GROUP report, which provides a detailed analysis of the response time, or to the SYSINFO report.

Indicator

Field: Response Time Distribution Graphic

Description: The horizontal axis shows response time (in seconds) with the response time goal in the middle. The middle section of the graph surrounding the goal shows the distribution of transactions that met between 60% and 150% of the goal.

Guideline: Easy — the more *green* fields you see the better your system is performing.

Problem Area: A big number of transactions on the right of the graphic in *red* indicates a problem.

Potential Solution: There can be different reasons for getting a distribution graphic like the one shown in the sample report:

- If you get this graphic more or less permanently, then probably different types of transactions are classified for the same service class. You can analyze the response time of the transactions by classifying them to different report classes, and then — as the second step — to different service classes to get more meaningful reports.
- The problem can also be either an untuned sysplex, a temporary performance problem, or an unrealistic performance goal. Then you need to investigate further.

Indicator

Field: System

Description: The bottom part of the report shows each system in the sysplex that provides service to the chosen service class.

Guideline: You can use this report to evaluate possible anomalies among the different systems that provide service.

Problem Area: This part of the report shows whether the any system has problems achieving the required goal.

Potential Solution: With cursor-sensitive control you can navigate from this report into the SYSINFO report for the system you are interested in. Then you can continue performance analysis as described in “Using Monitor III reports” on page 36.

Work Manager Delays report

```

RMF V1R10 Work Manager Delays - WTSCPLX1 Line 1 of 7
Command ==> Scroll ==> HALF

WLM Samples: 385 Systems: 3 Date: 04/30/08 Time: 09.30.00 Range: 100 Sec

Class: CICUSRTX Period: 1 Avg. Resp. time: 2.139 sec for 1805 TRX.
Goal: 0.090 sec average Avg. Exec. time: 1.576 sec for 1558 TRX.
Actual: 2.139 sec average Abnormally ended: 0 TRX.

Sub P ----- Response time breakdown (in %) ----- -Switched--
Type Tot Act Rdy Idle -----Delayed by----- Time (%)
LOCK I/O CONV DIST SESS TIME PROD MISC LOC SYS REM

CICS B 79.0 0.5 1.8 0 2.3 1.3 67.0 0 6.1 0 0 0 67 0 0
CICS X 66.4 0.5 0.4 0 2.8 21.6 0 0 40.8 0 0 0.3 22 0 0

----- Address Spaces Serving this Service Class CICUSRTX -----
Jobname M ASID System Serv-Class Service Proc-Usg I/O-Usg Veloc Capp Quies
CICSAOR1 46 SC47 SYSSTC 100 .61 4 55 0 0
CICSAOR2 49 SC47 SYSSTC 100 3.7 3 74 0 0
CICSAOR3 50 SC49 SYSSTC 100 4.2 1 43 0 0
CICSAOR4 52 SC52 SYSSTC 100 2.1 2 79 0 0
CICSCMAS 54 SC47 SYSSTC 50 11 8 54 0 0
CICSFOR 39 SC47 CICSRGN 100 14 6 77 0 0
CICSTOR 51 SC47 CICSRGN 100 14 8 63 0 0
    
```

Figure 16. Monitor III Work Manager Delays Report

Indicator

Field: Response time breakdown (in %)

Description: This part of the report provides performance information for the begin-to-end phase (CICS only) and the execution phase of CICS and IMS.

Guideline: This information is the same as in the Postprocessor Workload Activity report with just the difference in the interval length.

For details, refer to "Understanding Workload Activity data for IMS or CICS" on page 146.

Indicator

Field: Address Spaces Serving this Service Class

Description: You can see in this part of the report what address spaces provide service on different systems in the sysplex.

Guideline: The velocity (Veloc) and the processor service (Proc-Usg) for each address space can be seen as initial indicators to know how well the address spaces are providing service to the service classes of your transactions.

Identifying the major delay in a response time problem

You have a response time problem? Want to know where to go next? This section will take you through several methods of analyzing a response time problem to find the major cause. It then points you to the relevant chapter for further diagnosis and problem resolution suggestions.

For a response time problem this means understanding the end-to-end response time components as completely as possible. As shown in Figure 17, most transactions can be viewed as having two major response time components:

- *internal* (for example, CPU time, paging, I/O)
- *external* (for example, network delay, server delay from other platforms)

RMF can help you analyze the internal response time components, and that is where this book will focus. The external response time components are beyond the scope of RMF, and will require other monitoring tools such as NetView^{®*}.




Figure 17. End-to-End Response Time Components

Understanding the relative proportions of these components allows you to see where you have the most leverage to improve response times. For example, if 90% of your response time is in I/O, why start with the CPU?

For response time problems, generally the best place to start is with Monitor III. It reports on contention for resources (both hardware and software) and any associated *delays*.

Start by setting the *range* (or interval) of the data you're looking at to match at least a portion of the time in which the problem occurred. This could be current data in GO mode, or prior data. The choice of the range value must not be too high, which would smooth the peaks, or too low, to have enough samples to provide useful data.

The range should contain at least about 100 cycles. If the cycle is one-second (which is a good value, and is the default in RMF) a 1-5 minute time for the interval will generally fit. 100 seconds is the recommended value. Displayed

Response time analysis

percentages are then easily converted into elapsed time. If you are unfamiliar with altering these values, see the *z/OS RMF User's Guide* for more details.

If your response problem is sporadic, or due to brief spikes, you may need to adjust your range, moving forward or backward in time to find the problem. If you're having trouble, set the range to 100 seconds and move through time to pinpoint a spike.

Do not base your analysis on a single occurrence, as it could be an anomaly. Check over multiple days or times of day, to verify that the problem is real.

Using Monitor III reports

This section will discuss several reports, showing you how to identify the major cause of delay to a workload. Your starting point might be the GROUP report when you used the cursor sensitivity in the SYSRTD report.

GROUP Report

If you have a response time problem for TSO or batch, go to the Group Response Time report. Look at the Primary Response Time Component field to see the largest component of response time.

```

RMF V1R10 Group Response Time
Command ==>

Samples: 100      System: PRD1 Date: 04/07/08 Time: 10.32.00 Range: 100 Sec
Class: TSOPROD   Period: 1   Description: TSO Production
Primary Response Time Component: Storage delay for local paging

WFL  Users      Frames  Vector  EXCP  PgIn  TRANS  --- Response Time ---
%    TOT  ACT    %ACT    UTIL  Rate  Rate   -- Ended TRANS-(Sec) -
50   2    1      5       0     3.2   4.1    0.2    0.000  3.432  3.432

                                -AVG USG-  -----Average Delay-----
                                Total  PROC  DEV  PROC  DEV  STOR  SUBS  OPER  ENQ  OTHER
Average Users                    0.600  0.10 0.17  0.02 0.03  0.22  0.00  0.00  0.00  0.06
Response Time ACT                 3.432  0.57 0.97  0.11 0.17  1.25 0.00  0.00  0.00  0.34

                                ---STOR Delay---  ---OUTR Swap Reason---  ---SUBS Delay---
                                Page Swap  OUTR   TI   TO   LW   XS   JES  HSM  XCF
Average Users                    0.19 0.03 0.00  0.00 0.00  0.00  0.00  0.00  0.00  0.00
Response Time ACT                 1.08 0.17 0.00  0.00 0.00  0.00  0.00  0.00  0.00  0.00
    
```

Figure 18. Monitor III Group Response Time Report

Report Analysis

This report has two indicators for storage problems:

- Primary Response Time Component: Storage delay for local paging
- Response Time ACT: Average storage delay of 1.25 seconds

For more information, position the cursor on the field you are interested in and press ENTER. For this example, the reporting link will take you to the Storage Delay report, shown in figure 19.

STOR Report

```

RMF V1R10 Storage Delays                               Line 1 of 2
Command ==>>>                                         Scroll ==>> HALF

Samples: 100      System: PRD1 Date: 04/07/08 Time: 10.32.00 Range: 100 Sec

Jobname  C  Class  DLY  ----- % Delayed for ----- -- Working Set --
          C  Class  %    COMM LOCL  VIO  SWAP  OTR   Central  Expanded
BAJU     T  TSO    21    0   18   0    3    0    273     5
BHOL     T  TSO     1    1    0    0    0    0    1111    4
Report is for service class TSO only.
    
```

Figure 19. Monitor III Storage Delays Report

This report shows the address spaces which are affected by storage delays, in this example just for service class TSO.

Report Analysis

Address space BAJU is shown with the largest delay.

DELAY Report

If you do not know which group a delayed user belongs to, use the Delay report.

```

RMF V1R10 Delay Report                               Line 1 of 43
Command ==>>>                                         Scroll ==>> HALF

Samples: 100      System: PRD1 Date: 04/07/08 Time: 10.32.00 Range: 100 Sec

Name      Service  WFL USG DLY IDL UKN ----- % Delayed for ---- Primary
          CX  Class  Cr  %  %  %  %  %  PRC  DEV STR SUB OPR  ENQ  Reason
BHOLEQB  B  BATCH    0  0  51  0  49  0  0  0  0  0  0  51  SYSDSN
JES2     S  SYSSTC   0  0  2  0  98  0  2  0  0  0  0  0  SYSPAG
SMF      S  SYSSTC   0  0  1  0  99  0  1  0  0  0  0  0  SYSPAG
BAJU     T  TSO    29  9  22  67  2  1  0  21  0  0  0  0  LOCL
RMFGAT   S  SYSSTC  50  1  1  0  98  1  0  0  0  0  0  0  BAJU
BHOL     T  TSO    78 18  5  73  4  1  3  1  0  0  0  0  RMFUSR
BHOLPRO2 B  BATCH   93 93  7  0  0  7  0  0  0  0  0  0  BHOLPRO1
    
```

Figure 20. Monitor III Delay Report

The report lists all delayed address spaces sorted in ascending order by workflow percentage.

Report Analysis

There are two jobs with significant delays:

- BHOLEQB with 51% enqueue delay
- BAJU with 21% storage delay

We will continue analyzing the storage problem by positioning the cursor at the BAJU address space name for further processing. This leads to the Job Delay report with initial information about possible causes, as shown in figure 21.

JOB Report

```

RMF V1R10 Job Delays                               Line 1 of 4
Command ==>                                         Scroll ==> HALF

Samples: 100    System: PRD1 Date: 04/07/08 Time: 10.32.00 Range: 100 Sec
Job: BAJU      Primary delay: Paging of private area storage.

Probable causes: 1) Job may be using excessive central storage.
                  2) Paging configuration may need tuning.
Help panels contain more possible causes.

----- Job Storage Usage Data -----
Average Frames: 300    Working Set: 278    Fixed Frames: 43
Active Frames: 92     Aux Slots: 978    DIV Frames: 0
Idle Frames: 209     Page In Rate: 14.1    ES Move Rate: 0.0

----- Job Performance Summary -----
Service  P WFL -Using%- DLY IDL UKN ----- % Delayed for ----- Primary
CX ASID Class P Cr % PRC DEV % % % PRC DEV STR SUB OPR ENQ Reason
T 0027 TSO * 29 4 5 22 67 2 1 0 21 0 0 0 0 LOCL
      TSO 1 33 1 1 4 67 2 0 0 4 0 0 0 0 SWAP
      TSO 2 35 2 4 11 0 0 1 0 10 0 0 0 0 LOCL
      TSO 3 13 1 0 7 0 0 0 0 7 0 0 0 0 LOCL
    
```

Figure 21. Monitor III Job Delays Report

Report Analysis

Monitor III has done basic analysis of the status and provides some possible causes. Use these as a starting point, and investigate further to confirm or rule-out these possibilities. In this example, swap trim may also be worth investigating.

In this example you would need to investigate processor storage further. See Chapter 5, "Analyzing processor storage activity" for details.

WFEX report

If your help desk uses the Workflow/Exceptions report as a continual system monitor, they would be alerted to exception conditions. These could be your first indicators to a problem. This depends on your having customized the screen to your particular needs and thresholds. See the *z/OS RMF Report Analysis* for tailoring information.

Or, you can use this screen as your first diagnosis aid if investigating a user problem.

```

RMF V1R10 Workflow/Exceptions                               Line 1 of 4
Command ==>                                               Scroll ==> HALF

Samples: 100      System: PRD1 Date: 04/07/08 Time: 10.32.00 Range: 100 Sec

----- Speed (Workflow) -----
Speed of 100 = Maximum, 0 = Stopped      Average CPU Util: 100 %
Name    Users Active   Speed      Name    Users Active   Speed
*SYSTEM    43     3       72      TSOPROD     2     1       50
ALL TSO     2     1       50      BATCHPRD    6     3       75
ALL STC    35     0       84
ALL BATCH   6     3       75
ALL ASCH                    Not avail
*PROC      17     2       93
*DEV        6     0       21

----- Exceptions -----
Name    Reason          Critical val. Possible cause or action
*ECSA*  SECS% > 85      90.6 %      System ECSA usage 91 %.
*STOR   TSQA0 > 0       784K bytes  SQA overflow into CSA 784K.
BHOLEQB ENQ -SYSDSN     51.0 % delay ENQ.TEST.DAT
HSM     Not avail                      Job HSM is not running.

```

Figure 22. Monitor III Workflow/Exceptions report

Speed, or workflow, can be used as a performance indicator either for the total system or for groups of users or resources. It is an indicator of how well a workload is able to get the system resources it wants. 100% means it gets whatever it wants whenever it wants it; 0% means it never gets what it wants. Speed is also used on a resource level (processors or devices) to reflect contention for that resource.

Speed can also be a useful indicator for exception reporting for your workloads. Make note of the workflow values reported when your system is running well. Use these as your own guideline values to set up exception reporting.

Report Analysis

The Speed of 50% for all TSO users might suggest a response time problem, if users are complaining, or if your experience has shown that a higher workflow is required to meet service objectives on your system.

However, your system may function well at a 50% speed. Use your judgment and review related indicators to determine whether or not a low speed is a problem.

DELAY versus USING

So far, we have been concentrating on workload delays. Your workload's dominant state could also be USING. Both cases will be discussed further in the relevant chapters.

Performance problem due to DELAY: If the largest single component of response or delay for the workload you're interested in is:

- PROC (or processor delays), go to Chapter 3, "Analyzing processor activity" for further analysis
- DEV (or device delays), go to Chapter 4, "Analyzing I/O activity" for further analysis
- STOR (or storage delays), go to Chapter 5, "Analyzing processor storage activity" for further analysis

If some other component (ENQ, IDL, SUBS, OPER, UKN) indicates a problem, see Appendix D. Other delays.

Performance problem due to USING: If there are no delays of any significance (at least 5-10%), look at the USING% values for PROC and DEV in the Job Delay report (other reports also have this information). Select the larger, and go to the corresponding CPU or I/O chapter for some thoughts on reducing the workload's need for that resource or speeding up the processing.

Monitoring daily performance: System indicators

As discussed at the beginning of this chapter, there are different kinds of performance problems. Response time complaints or SLA exceptions are two examples. Potential problems may also be indicated by your daily performance monitoring. This section will discuss some RMF indicators which will be of use in your daily monitoring.

Here are some of the most frequently used global indicators of MVS system health, along with some *rules-of-thumb* (ROTs), on what constitutes potential trouble.

As with any ROTs, these should be looked at as a starting point only. There is no magic single number for any indicator that suggests you have a problem.

The most important thing is that your service levels are being met.

We will highlight the indicators by using RMF reports, with pointers to the appropriate chapter of this book for more details.

As you read these, remember that Monitor III will show you the impact of most of these indicators:

- *WHO* is being delayed?
- For *HOW LONG*?
- And *BY WHOM*?

Use this delay information to help you decide whether a given indicator really means trouble or not.

Summary of major system indicators

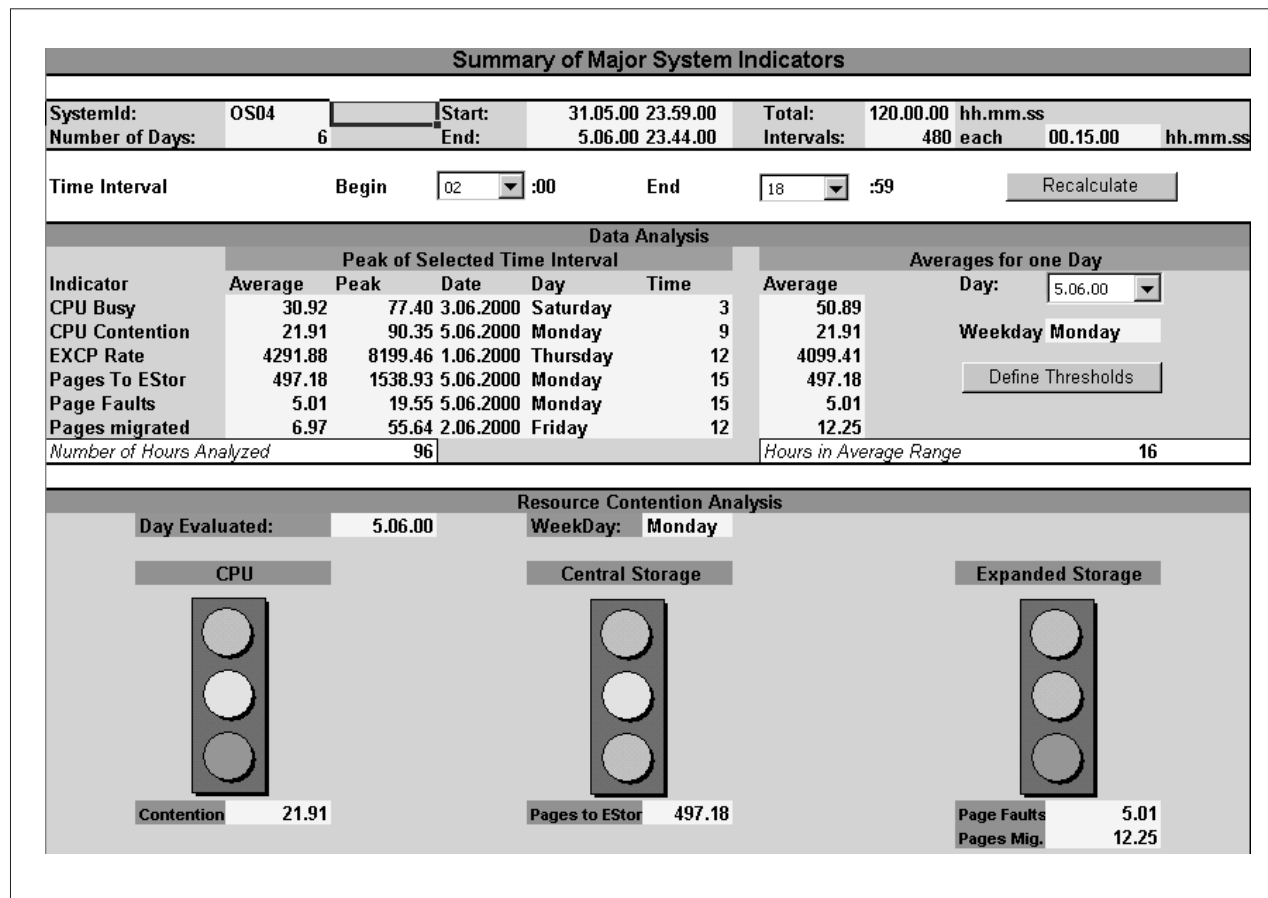


Figure 23. Summary of Major System Indicators. Spreadsheet Reporter macro RMFY9OVW.XLS (System Overview Report - Summary)

The Summary report displays key indicators representing the performance of your system. At a glance, you see whether your system experiences processor or storage contention.

A resource contention analysis is done for three indicators:

- CPU Contention** The percent of time with at least one job waiting for the CPU.
- CS Contention** The page movement rate from CS to ES is used as indicator for CS contention.
- ES Contention** The page fault rate and the migration rate are used as indicators for ES contention. However, expanded storage is not used on systems running in z/Architecture mode.

You can specify a time range and a day to perform the contention evaluation. The calculated values will be compared against a set of defined thresholds:

- CPU Contention** 20 - 80
- CS Contention** 200 - 800

ES Contention	10 - 50 (page fault rate)
	30 - 80 (migration rate)

You can specify the low and high thresholds which are used to set the traffic light to red or green according to your own experiences. Especially, the ES contention threshold depends on the key workload you are running in your system. If it is an online application (for example CICS or DB2), you should decrease the values to 10 - 30 (page fault rate) and to 20 - 50 (migration rate). On the other hand, a TSO system can run with much higher values, here you might specify 20 - 70 and 50 - 150.

Please keep in mind that the contention analysis should be done with data for a longer interval, the sample report covers prime shift data (8:00am to 3:59pm), therefore, the thresholds are different to other values in this publication, because those will be taken for interval reports of 15 or 30 minutes.

Using Monitor I reports

This section discusses the use of the following Monitor I reports:

- "CPU Activity report" on page 44
- "Partition Data report" on page 47
- "Channel Path Activity report" on page 48
- "I/O Queuing Activity report" on page 49
- "DASD Activity report" on page 50
- "Paging Activity - Swap Placement Activity" on page 53
- "Page/Swap Data Set Activity report" on page 54
- "Workload Activity report" on page 55
- "Virtual Storage Activity report" on page 57

Rules-of-Thumb**BUSY TIME PERCENTAGE**

A value of over 85%, along with the value for the NUMBER OF ADDRESS SPACES in the distribution of the IN READY QUEUE for line $\leq N$ below 60% implies contention for the CPU.

If your CPU is more than 85% busy, look at the values for the DISTRIBUTION OF QUEUE LENGTHS (%). If the number of address spaces in line $\leq N$ for example is **58.8**, this means that in 58.8% of the time, you had enough CPs to handle all the work on the IN READY queue. In the remaining time you had ready work waiting for CPs.

This does not mean that you cannot run a z/OS system at higher utilizations, up to 100% busy. You can. Just be aware that the busier the CPU, the longer the CPU delay for lower-priority workloads will be. The more low-priority work you have, and the fewer non-CPU bottlenecks you have (for example, responsive I/O), the busier you can run your CPU and still maintain good response for your high-priority workloads.

If your goal is to run your system to 100% busy, minimize your non-CPU bottlenecks and use the information from the Monitor III Processor Delay report to keep an eye on CPU delay for your low-priority work.

BUSY TIME PERCENTAGE

Conversely, a low CPU busy could indicate that other bottlenecks in the system are preventing work from being dispatched.

A peak-to-average ratio (for example, peak hour busy/prime shift average busy) of less than 1.3-1.4 for a commercial workload could also indicate bottlenecks preventing work from using the CPU.

This ratio is inversely correlated to IN READY users: the numbers of IN READY users go up as the peak-to-average ratio comes down. This is due to latent demand and the inability of the system to dispatch all the ready work.

OUT READY

Number of address spaces swapped out (probably TSO and batch), but ready to execute. If it is greater than 1, and TSO or batch response is an issue, look into this. This could reflect processor storage constraints, and probably you may need to update the WLM service policy.

Report Analysis

- BUSY TIME PERCENTAGE is **97.31** %
- Distribution of address spaces in IN READY QUEUE $\leq N$: **58.8** %

This might point to a problem, because at least 41.2% of the time one or more address spaces are delayed for CPU. The Monitor III Processor Delay report will tell you which AS were delayed, and for how long.

System indicators

See Chapter 3, “Analyzing processor activity” and Chapter 5, “Analyzing processor storage activity” for further analysis.

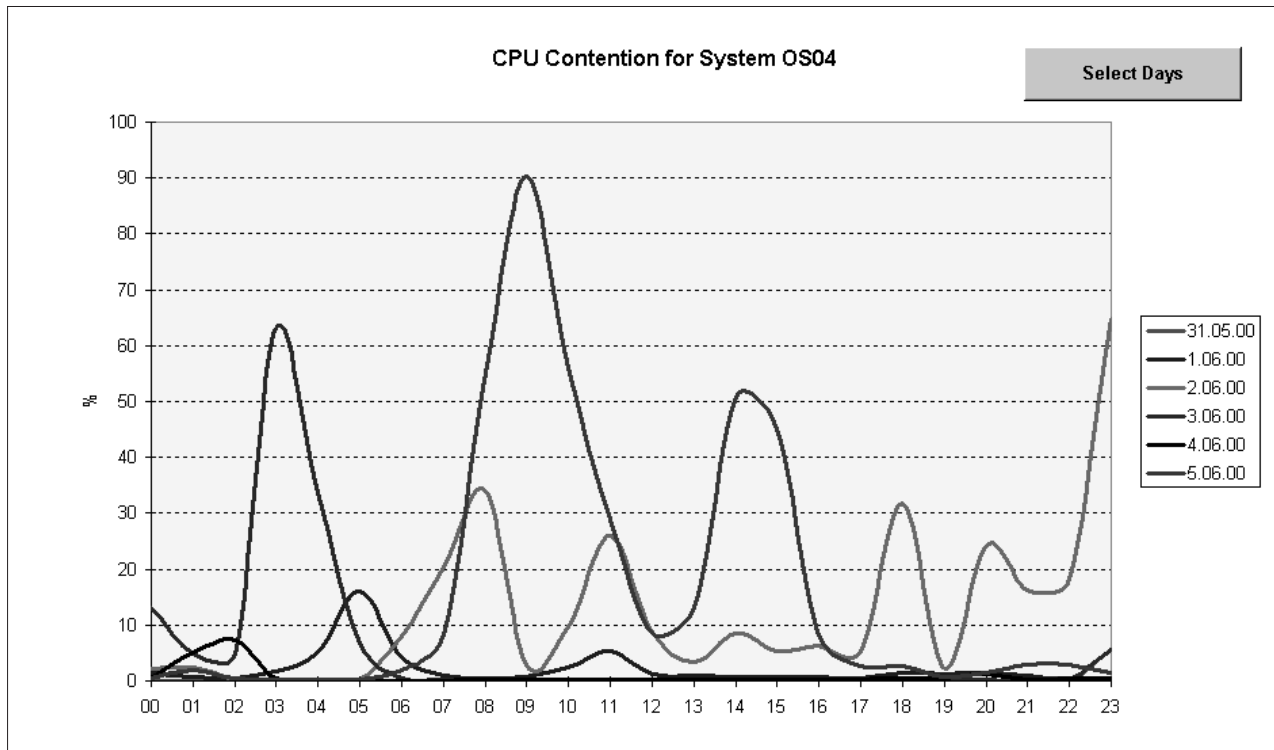


Figure 25. CPU Contention. Spreadsheet Reporter macro RMFY9OVW.XLS (System Overview Report - DaysCont)

The CPU Contention report helps you in understanding your processor capacity. It extends the *long-term* indicator in the Summary report (Figure 23 on page 42) and informs you about the contention of your system in some more detail.

Partition Data report

PARTITION DATA REPORT															
z/OS V1R10		SYSTEM ID NP1				DATE 11/28/08				INTERVAL 14.59.678				PAGE 2	
		RPT VERSION V1R10 RMF				TIME 09.30.00				CYCLE 1.000 SECONDS					
MVS PARTITION NAME		NP1				NUMBER OF PHYSICAL PROCESSORS				9		GROUP NAME		N/A	
IMAGE CAPACITY		100				CP				7		LIMIT		N/A	
NUMBER OF CONFIGURED PARTITIONS		9				ICF				2					
WAIT COMPLETION		NO													
DISPATCH INTERVAL		DYNAMIC													
----- PARTITION DATA ----- -- LOGICAL PARTITION PROCESSOR DATA -- -- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --															
NAME	S	---MSU---			-CAPPING--		PROCESSOR- NUM TYPE	---DISPATCH TIME DATA---		LOGICAL PROCESSORS		--- PHYSICAL PROCESSORS ---			
		WGT	DEF	ACT	DEF	WLM%		EFFECTIVE	TOTAL	EFFECTIVE	TOTAL	LPAR MGMT	EFFECTIVE	TOTAL	
NP1	A	20	100	10	NO	62.2	1.2	CP	00.04.27.302	00.04.27.519	24.86	24.92	0.01	4.24	4.25
NP2	A	1	0	1	YES	0.0	4	CP	00.00.21.680	00.00.22.083	0.60	0.61	0.01	0.34	0.35
NP3	A	10	5	8	NO	3.3	1.0	CP	00.03.35.761	00.03.35.859	23.97	23.98	0.02	3.41	3.43
NP4	A	300	95	155	NO	0.0	6.0	CP	01.12.05.405	01.12.06.405	80.10	80.15	0.01	68.68	68.69
NP5	A	200	50	52	NO	0.0	4	CP	00.24.11.147	00.24.11.311	40.30	40.31	0.01	23.02	23.03
PHYSICAL								00.00.03.103				0.05		0.05	
TOTAL								01.44.41.295 01.44.46.280				0.11		99.69 99.80	
CFC1	A	DED					1	ICF	00.14.59.611	00.14.59.625	99.98	99.99	0.01	99.95	99.96
CFC2	A	DED					1	ICF	00.00.00.000	00.00.00.000	0.00	0.00	0.00	0.00	0.00
PHYSICAL								00.00.00.065				0.03		0.03	
TOTAL								-----				0.04		99.95 99.99	
CB88	D														
CB89	D														

Figure 26. Monitor I CPU Activity Report - Partition Data Report Section

If your system is showing signs of CPU constraint (see the CPU Activity Report from Figure 24 on page 44), check the Partition Data report to be sure that your LP is correctly configured. You would do this to allocate more CPU resource to a particular LP.

See Appendix A. PR/SM LPAR considerations for further analysis.

Channel Path Activity report

CHANNEL PATH ACTIVITY									
z/OS V1R10			SYSTEM ID OS04		DATE 06/05/2008		INTERVAL 14.59.996		
			RPT VERSION V1R10 RMF		TIME 09.29.00		CYCLE 1.000 SECONDS		
IODF = 01	CR-DATE: 07/10/08	CR-TIME: 21.00.01	ACT: POR	MODE: LPAR	CPMF: EXTENDED MODE				
B5 CNC_S Y	25.22	26.16		B0 CNC_S Y	22.65	23.48			
B6 CNC_P	OFFLINE			B1 CNC_S Y	0.00	0.01			
B7 CNC_S Y	0.01	0.01		B2 CTC_S Y	0.40	1.48			
B8 CNC_S Y	0.00	0.00		C0 CNC_S Y	2.11	2.77			
B9 CNC_S Y	0.00	0.00		C1 CNC_S Y	0.00	0.00			
BA CNC_S Y	1.11	1.16		C2 CNC_S Y	0.00	0.00			
BB CNC_S Y	2.55	2.56		C3 CNC_P	OFFLINE				
BC CNC_S Y	4.63	9.21		C4 CNC_S Y	0.00	0.00			
C5 CNC_S Y	0.00	0.00		CD CFS Y	-----	0.00			
C6 CNC_S Y	13.70	14.72		D0 CNC_S Y	2.03	2.68			
C7 CNC_S Y	33.26	33.26		D1 CNC_S Y	0.00	0.00			
C8 CNC_S Y	0.00	0.00		D2 CVC	0.00	0.00			
C9 CNC_S Y	4.40	8.84		D3 CNC_S Y	32.96	32.97			
CA CNC_S Y	0.00	0.00		D4 CNC_S Y	37.13	39.54			
CB CNC_? Y	0.00	0.00		D5 CNC_S Y	1.24	1.27			
CC CFS Y	-----	0.00		D6 CNC_S Y	0.00	0.00			

Figure 27. Monitor I Channel Path Activity report

Tip:

Using URL <http://www.ibm.com/servers/eserver/zseries/library/techpapers/> in section **Networking & Connectivity**, you can find a document called *FICON and FICON Express Channel Performance* that helps you with understanding the performance characteristics of the different versions of FICON and FICON Express channels. This paper also explains in detail the benefits of FICON and FICON Express channels in both FC and FCV mode, discusses several channel configurations, compares ESCON and FICON technology and, last not least, recommends how to use FICON RMF information from various reports for I/O configuration performance measurement.

Report Analysis

Channel 54 has a partition utilization of 37.13% and a total utilization of 39.54%. This may not be a problem.

To find out which *logical control unit (LCU)* is using this channel, look in the I/O Queuing Activity report (see Figure 28 on page 49). From there, you can go on to check device response times.

See Chapter 4, “Analyzing I/O activity” for further analysis.

I/O Queuing Activity report

I/O QUEUING ACTIVITY															PAGE 1
z/OS VIR10			SYSTEM ID SYS1		DATE 11/28/08		INTERVAL 15.00.036								
			RPT VERSION VIR10 RMF		TIME 16.30.00		CYCLE 1.000 SECONDS								
TOTAL SAMPLES = 900			IODF = 26		CR-DATE: 12/08/2006		CR-TIME: 16.09.06		ACT: ACTIVATE						
IOP	- INITIATIVE QUEUE -		IOP UTILIZATION			-- % I/O REQUESTS RETRIED --					----- RETRIES / SSCH -----				
	ACTIVITY RATE	AVG Q LNGLTH	% IOP BUSY	I/O START RATE	INTERRUPT RATE	ALL	CP BUSY	DP BUSY	CU BUSY	DV BUSY	ALL	CP BUSY	DP BUSY	CU BUSY	DV BUSY
00	1456.345	0.01	28.06	1456.345	1910.201	3.9	1.7	2.2	0.0	0.0	0.04	0.02	0.02	0.00	0.00
01	81.124	0.00	3.45	81.124	213.122	0.1	0.1	0.0	0.0	0.0	0.00	0.00	0.00	0.00	0.00
02	1305.816	0.00	2.25	1305.816	1083.198	0.9	0.7	0.0	0.2	0.0	0.01	0.01	0.00	0.00	0.00
03	5469.766	0.34	16.33	5469.766	5061.730	59.8	59.8	0.0	0.0	0.0	1.49	1.49	0.00	0.00	0.00
04	2057.740	0.01	4.21	2057.739	2079.958	0.2	0.0	0.0	0.2	0.0	0.00	0.00	0.00	0.00	0.00
05	253.301	0.00	0.83	253.301	413.870	0.4	0.2	0.2	0.0	0.0	0.00	0.00	0.00	0.00	0.00
06	5154.934	0.35	18.62	5154.934	5498.320	50.6	50.5	0.0	0.1	0.0	1.02	1.02	0.00	0.00	0.00
07	112.071	0.00	1.47	112.071	136.763	24.7	4.4	20.2	0.0	0.0	0.33	0.06	0.27	0.00	0.00
SYS	15891.10	0.23	9.40	15891.10	16397.16	46.0	45.7	0.2	0.1	0.0	0.85	0.85	0.00	0.00	0.00
LCU	CU	DCM MIN	GROUP MAX	CHAN DEF	CHPID TAKEN	% DP BUSY	% CU BUSY	AVG CUB DLY	AVG CMR DLY	CONTENTION RATE	DELAY Q LNGLTH	AVG CSS DLY	HPAV WAIT	AVG OPEN EXCH	DATA XFER CONC
0011	1400			90	64.501	0.00	0.05	0.0	3.1						
				B0	83.215	0.00	0.03	0.0	2.6						
				D7	63.353	0.00	0.11	0.0	3.2						
				E3	30.523	0.00	0.42	0.0	3.5						
				*	241.59	0.00	0.11	0.0	3.0	0.008	3.00	22.2	0.000	2	0.052 0.001
0012	1500			10	0.034	0.00	0.00	0.0	0.0						
				11	0.038	0.00	0.00	0.0	0.0						
		2	2	6	0.074	0.00	0.00	0.0	0.0						
				*	0.147	0.00	0.00	0.0	0.0	0.000	0.00	0.1	0.000	0	0.052 0.001
0013	1600			90	64.088	0.00	0.04	0.0	3.0						
				B0	82.921	0.00	0.03	0.0	2.6						
				D7	63.831	0.00	0.12	0.0	3.2						
				E3	30.712	0.00	0.31	0.0	3.5						
				*	241.55	0.00	0.09	0.0	3.0	0.000	0.00	22.2	0.000	3	0.052 0.001
0014	1700			EC	0.020	0.00	0.00	0.0	0.0						
				88	0.052	0.00	0.00	0.0	0.1						
				D9	0.032	0.00	0.00	0.0	0.1						
				96	0.036	0.00	0.00	0.0	0.0						
				*	0.140	0.00	0.00	0.0	0.1	0.000	0.00	0.2	0.000	0	0.052 0.001

Figure 28. Monitor I I/O Queuing Activity Report

Report Analysis

Channel path **90** is connected to LCUs **0011** and **0013**. The corresponding devices can be seen in the Device Activity report.

DASD Activity 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																					
z/OS V1R9				SYSTEM ID UIG1				DATE 02/10/2007				INTERVAL 14.59.996									
				RPT VERSION V1R9 RMF				TIME 09.29.00				CYCLE 1.000 SECONDS									
TOTAL SAMPLES = 1,807 IODF = 12 CR-DATE: 03/01/08 CR-TIME: 10.23.17																					
STORAGE GROUP	DEV NUM	DEVICE TYPE	NUMBER OF CYL	VOLUME SERIAL	PAV	LCU	DEVICE ACTIVITY RATE	AVG RESP TIME	AVG IOSQ	AVG CMR	AVG DB	AVG PEND TIME	AVG DISC TIME	AVG CONN TIME	% DEV CONN	% DEV UTIL	% DEV RESV	AVG NUMBER ALLOC	% ANY ALLOC	% MT PEND	
SGGDG	0340	33903	200	IDP386		0056	0.139	18.4	0.0	0.0	0.0	0.6	0.7	17.1	0.24	0.25	0.0	1.3	100.0	0.0	
SGGDG	0341	33903	200	IDP387		0056	0.001	0.9	0.0	0.0	0.0	0.3	0.0	0.6	0.00	0.00	0.0	0.0	100.0	0.0	
SGGDG	0342	33903	200	IDP388		0056	3.311	22.4	0.0	0.0	0.2	0.6	0.2	21.5	7.12	7.18	0.0	1.0	100.0	0.0	
SGGDG	0343	33903	200	IDP389		0056	0.001	0.9	0.0	0.0	0.0	0.3	0.0	0.6	0.00	0.00	0.0	0.0	100.0	0.0	
SGCCBS	0344	33903	200	CCB178		0056	0.150	12.8	0.0	0.0	0.0	0.5	9.9	2.4	0.04	0.18	0.0	2.0	100.0	0.0	
SGCCBS	0345	33903	200	CCB179		0056	3.703	23.2	11.0	0.0	0.0	0.3	4.8	7.1	2.63	4.40	0.0	3.7	100.0	0.0	
SGDBSWN	0346	33903	200	DBSW11		0056	0.959	2.3	0.0	0.0	0.2	0.5	0.0	1.8	0.17	0.17	0.0	11.0	100.0	0.0	
SGDBSWN	0347	33903	200	DBSW12		0056	0.003	0.9	0.0	0.0	0.0	0.3	0.0	0.6	0.00	0.00	0.0	2.0	100.0	0.0	
SGSW	0348	33903	200	IDP302		0056	0.003	0.9	0.0	0.0	0.0	0.3	0.0	0.6	0.00	0.00	0.0	0.0	100.0	0.0	
		0349	33903	200	PLEXP1		0056	0.002	0.7	0.0	0.0	0.2	0.1	0.4	0.00	0.00	0.0	0.0	100.0	0.0	
:																					
SGSW	037E	33903	200	IDP303		0056	0.003	0.8	0.0	0.0	0.0	0.3	0.0	0.5	0.00	0.00	0.0	0.0	100.0	0.0	
SGCRITC	037F	33903	200	IDP076		0056	0.003	0.8	0.0	0.0	0.0	0.2	0.0	0.6	0.00	0.00	0.0	1.0	100.0	0.0	
						LCU	0056	44.580	13.0	0.5	0.0	0.2	0.7	3.6	8.2	0.57	0.82	0.0	137	100.0	0.0

Figure 29. Monitor I DASD Activity Report

Rule-of-Thumb

Tape LCU AVG RESP TIME

Look at the ratio of LCU DISC + PEND time to CONN time. When the DISC+PEND exceeds CONN for extended periods of time, this is an indication of channel and/or control unit contention.

Report Analysis

Several devices 340, 342, 344, or 345 have extraordinary long response times and need further investigation. (Which workloads are using these devices? How much are they being delayed? Are these response times typical for my system?).

See Chapter 4, "Analyzing I/O activity," on page 81 for further analysis.

RMF DASD System Summary						
System Id: UIG1		Operating System: z/OS V1R9 -		Report Range(seconds): 3600		
Reporting Date: 02.10.2007		Reporting Time: 09.05.11		Report Range(hh.mm.ss): 001.00.00		
Overall	Total Number of Logical Control Units	29	Installed Capacity Estimation (*)			
	Total Number of DASDs	1669	Types	GB (*)	Number	Capacity
	Installed Capacity Estimation (*)	4711,11 GB	3380D	0,63	0	0
	Total Activity Rate	1270,41 I/O per sec	3380E	1,26	0	0
	Access Density	0,27	3380J	0,63	0	0
Times	Connect Time	2,83 ms	3380K	1,89	0	0
	Disconnect Time	1,53 ms	33901	0,94	0	0
	Service Time	4,36 ms	33902	1,89	0	0
	Pending Time	0,24 ms	33903	(2.83)	1660	4711,11
	IOS Queue Time	0,03 ms	33906	8,44	0	0
	Response Time	4,62 ms	33909	(8.44)	0	0,00
			9345	1,50	0	0
			Unknown	NA	9	NA
			Additional	NA	NA	0
		Response Time / Service Time	1,06			
Intensities	I/O Intensity	5875,37 ms / s	Total (*)			
	Service Time Intensity	5536,31 ms / s			1669	4711,11
	Path Intensity	3597,10 ms / s	(*) several large disk models (e.g.27,54,...) are handled as device model 09. This means that the number and capacity and total summary information may be incorrect.			
DASD Skew	Average Service Time Intensity	3,32 ms / s				
	Highest Service Time Intensity	329,61 ms / s				
	Ratio: Highest/Average (DASD Skew)	99,36				
	Typical DASD Skew	26,17				
	Typical Range (low / high)	14,46	47,39			

Figure 30. DASD Summary. Spreadsheet Reporter macro RMFR9DAS.XLS (DASD Activity Report - System)

In the DASD Summary report, you see at a glance the key data of your DASD subsystem: performance values as well as capacity data. If you feel that you might see some more details, you can have a look on the ten most busy DASD volumes in your system:

System indicators




Figure 31. Response Time of Top-10 Volumes. Spreadsheet Reporter macro RMFR9DAS.XLS (DASD Activity Report - Top10RT)

Paging Activity - Swap Placement Activity

PAGING ACTIVITY												
z/OS V1R10		SYSTEM ID OS04			DATE 06/05/2008			INTERVAL 14.59.996				
		RPT VERSION V1R10 RMF			TIME 09.29.00			CYCLE 1.000 SECONDS				
OPT = IEAOPT02 MODE=ESAME		S W A P P L A C E M E N T A C T I V I T Y										
----- AUX STORAGE ----- *---LOGICAL SWAP---* *----- EXPANDED STORAGE -----*												
		TOTAL	AUX STOR TOTAL	AUX STOR DIRECT	AUX STOR VIA TRANSITION	LOG SWAP	LOG SWAP EFFECTIVE	EXP STOR DIRECT	EXP STOR TOTAL	MIGRATED FROM EXP STOR	EXP STOR EFFECTIVE	LOG SWAP /EXP STOR EFFECTIVE
TERMINAL	CT	1,558	0	0	0	1,558	1,535	0	23	0	23	1,558
INPUT/OUTPUT	RT	1.73	0.00	0.00	0.00	1.73	1.71	0.00	0.03	0.00	0.03	1.73
WAIT	%	23.2%	0.0%	0.0%	0.0%	100.0%	98.5%	0.0%	1.5%	0.0%	100.0%	100.0%
LONG	CT	2,081	0	0	0	2,081	2,060	0	21	0	21	2,081
WAIT	RT	2.31	0.00	0.00	0.00	2.31	2.29	0.00	0.02	0.00	0.02	2.31
	%	30.9%	0.0%	0.0%	0.0%	100.0%	99.0%	0.0%	1.0%	0.0%	100.0%	100.0%
DETECTED	CT	2,918	1	0	1	2,918	2,894	0	24	1	23	2,917
WAIT	RT	3.24	0.00	0.00	0.00	3.24	3.22	0.00	0.03	0.00	0.03	3.24
	%	43.4%	0.0%	0.0%	100.0%	100.0%	99.2%	0.0%	0.8%	4.2%	95.8%	100.0%
UNILATERAL	CT	6	0	0	0	6	6	0	0	0	0	6
	RT	0.01	0.00	0.00	0.00	0.01	0.01	0.00	0.00	0.00	0.00	0.01
	%	0.1%	0.0%	0.0%	0.0%	100.0%	100.0%	0.0%	0.0%	0.0%	0.0%	100.0%
EXCHANGE ON	CT	3	0	0	0	3	3	0	0	0	0	3
RECOMMENDA-	RT	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
TION VALUE	%	0.0%	0.0%	0.0%	0.0%	100.0%	100.0%	0.0%	0.0%	0.0%	0.0%	100.0%
TRANSITION	CT	8	0	0	0	0	0	8	8	0	8	8
TO NON-	RT	0.01	0.00	0.00	0.00	0.00	0.00	0.01	0.01	0.00	0.01	0.01
SWAPPABLE	%	0.1%	0.0%	0.0%	0.0%	0.0%	0.0%	100.0%	100.0%	0.0%	100.0%	100.0%
OMVS	CT	153	0	0	0	153	150	0	3	0	3	153
INPUT	RT	0.17	0.00	0.00	0.00	0.17	0.17	0.00	0.00	0.00	0.00	0.17
WAIT	%	2.3%	0.0%	0.0%	0.0%	100.0%	98.0%	0.0%	2.0%	0.0%	100.0%	100.0%
TOTAL	CT	6,727	1	0	1	6,719	6,648	8	79	1	78	6,726
	RT	7.48	0.00	0.00	0.00	7.47	7.39	0.01	0.09	0.00	0.09	7.48
	%	100.0%	0.0%	0.0%	100.0%	99.9%	98.9%	0.1%	1.2%	1.3%	98.7%	100.0%
AUXILIARY STORAGE - AVERAGE PAGES PER SWAP OUT -				1		AVERAGE PAGES PER SWAP IN -				257		
OCCURRENCES OF TERMINAL OUTPUT WAIT -				87								

Figure 32. Paging Activity report - Swap Placement Activity

Rule-of-Thumb

LOG SWAP/EXP STOR EFFECTIVE

This is the percentage of all swaps that are satisfied from processor storage. You want the TOTAL value under this column to be a high number (e.g. over 95 %) to keep swap delay times low.

The percentage under TERMINAL INPUT/OUTPUT WAIT should also be high, as this is where most TSO swaps are reflected.

Report Analysis

The LOG SWAP/EXP STOR EFFECTIVE values for

- TOTAL: 100%
- TERMINAL INPUT/OUTPUT WAIT: 100%

are perfect.

See Chapter 5, “Analyzing processor storage activity” for further analysis.

Page/Swap Data Set Activity report

PAGE / SWAP DATASET ACTIVITY													
z/OS VIRI0			SYSTEM ID OS04			DATE 06/05/2008			INTERVAL 14.59.996				
			RPT VERSION VIRI0 RMF			TIME 09.29.00			CYCLE 1.000 SECONDS				
NUMBER OF SAMPLES = 898			PAGE DATA SET USAGE										
PAGE SPACE TYPE	VOLUME SERIAL	DEV NUM	DEVICE TYPE	SLOTS ALLOC	---- SLOTS USED MIN	--- SLOTS USED MAX	--- SLOTS USED AVG	BAD SLOTS	% IN USE	PAGE TRANS TIME	NUMBER IO REQ	PAGES XFER'D	I DATA SET NAME
PLPA	PAGP11	0378	33903	45000	16466	16466	16466	0	0.11	0.200	6	5	PAGE.OS04.APLPA
COMMON	PAGP10	0377	33903	21240	5345	5378	5355	0	0.11	0.027	38	37	PAGE.OS04.ACOMMON
LOCAL	PAGP12	0379	33903	566820	67670	70965	69140	0	2.67	0.004	711	5,697	N PAGE.OS04.LOCAL01
LOCAL	PAGP13	037A	33903	566820	65887	69007	67190	0	3.01	0.005	771	5,711	N PAGE.OS04.LOCAL02
LOCAL	PAGP14	037B	33903	566820	66611	69723	67928	0	3.01	0.005	821	5,625	Y PAGE.OS04.LOCAL03
LOCAL	PAGP15	038D	33903	566820	66296	69563	67618	0	3.23	0.005	726	5,523	Y PAGE.OS04.LOCAL04
LOCAL	PAGP16	039F	33903	566820	66390	69311	67560	0	3.12	0.005	766	5,548	Y PAGE.OS04.LOCAL05

Figure 33. Monitor I Page/Swap Data Set Activity Report

Rules-of-Thumb

% IN USE

This is really the % busy for the data set. Above 30% you may start to see response times increase. However, you also need to consider other active data sets on the same volume; check the total device utilization on the device activity report.

PAGES XFER'D

The total number of pages transferred to and from a given page/swap dataset. Divide this number by the number of INTERVAL seconds to get the rate in pages/second. Above 30 pages/second for a single data set may mean it is time to add another page pack, dedicate the packs to paging (or better yet, look into more processor storage).

Report Analysis

If we check the two highest values in the report, we get:

- % IN USE: 3.23% - no problem seen
- PAGES XFER'D: 5711 - this is equivalent to 6 pages/second.

See Chapter 5, “Analyzing processor storage activity” for further analysis.

Workload Activity report

REPORT BY: POLICY=SPECIAL															
WORKLOAD=ONLINE			SERVICE CLASS=TSO			RESOURCE GROUP=*NONE			PERIOD=1 IMPORTANCE=2						
TRANSACTIONS	TRANS.-TIME	HHH.MM.SS.TTT	--DASD	I/O--	----	SERVICE----	SERVICE TIME	---	APPL %---	--PROMOTED--	-----STORAGE-----				
AVG	1.25	ACTUAL	96	SSCHRT	10.0	IOC	30100	CPU	32.300	CP	15.00	BLK	0.000	AVG	124.31
MPL	0.86	EXECUTION	96	RESP	15.4	CPU	357.3K	SRB	5.100	AAPCP	0.00	ENQ	0.000	TOTAL	107.44
ENDED	4,945	QUEUED	0	CONN	2.8	MSO	203.7K	RCT	0.000	IIPCP	0.00	CRM	0.000	SHARED	0.00
END/SEC	8.21	R/S AFFIN	0	DISC	2.5	SRB	56334	IIT	0.100						
#SWAPS	4,952	INELIGIBLE	0	Q+PEND	10.2	TOT	647.4K	HST	0.000	AAP	0.00			--PAGE-IN RATES--	
EXCTD	0	CONVERSION	0	IOSQ	0.0	/SEC	1,075	AAP	0.000	IIP	0.00			SINGLE	16.7
AVG ENC	0.00	STD DEV	119					AAP	0.000					BLOCK	6.2
REM ENC	0.00					ABSRPTN	1,244							SHARED	0.0
MS ENC	0.00					TRX SERV	856							HSP	0.0
...															
REPORT BY: POLICY=SPECIAL															
TRANSACTIONS	TRANS.-TIME	HHH.MM.SS.TTT	--DASD	I/O--	----	SERVICE----	SERVICE TIME	---	APPL %---	--PROMOTED--	-----STORAGE-----				
AVG	60.79	ACTUAL	749	SSCHRT	368.7	IOC	1106K	CPU	509.900	CP	105.40	BLK	0.000	AVG	152.18
MPL	55.66	EXECUTION	652	RESP	12.4	CPU	5630K	SRB	124.900	AAPCP	0.00	ENQ	0.000	TOTAL	8471.98
ENDED	14,418	QUEUED	96	CONN	2.4	MSO	9228K	RCT	0.200	IIPCP	0.00	CRM	0.000	SHARED	0.00
END/SEC	23.95	R/S AFFIN	0	DISC	2.1	SRB	1380K	IIT	1.300						
#SWAPS	14.8K	INELIGIBLE	0	Q+PEND	7.9	TOT	17343K	HST	0.000	AAP	0.00			--PAGE-IN RATES--	
EXCTD	0	CONVERSION	0	IOSQ	0.0	/SEC	28,819	AAP	0.000	IIP	0.00			SINGLE	0.9
AVG ENC	0.00	STD DEV	5.769					AAP	0.000					BLOCK	0.5
REM ENC	0.00					ABSRPTN	517							SHARED	0.0
MS ENC	0.00					TRX SERV	474							HSP	0.0

Figure 34. Workload Activity report

Rules-of-Thumb

APPL% CP

This value tells you how much CPU time was captured for the workload in a given service class. You can use this information for several purposes:

- Compute a *capture ratio (CR)* for your system
- Input to capacity planning
- Check for a single-CP constraint

Remember that the APPL% CP is given as percent of a single engine, so you will need to divide by the number of engines if your system has more than one.

Most MVS systems today have capture ratios greater than 80%. If yours is below that, you may want to investigate further.

PAGE-IN RATES

For any workload that is sensitive to paging (for example CICS), check the SINGLE and BLOCK paging values to be sure that no paging is taking place. The amount of paging that these workloads can sustain is up to you, but the best answer is probably zero.

Report Analysis

- Calculation of the capture ratio:

$$\text{Capture ratio} = \frac{\text{APPL\% CP}}{\text{MVS BUSY \% * \#CP}} = \frac{105.4}{73.88 * 2} = 0.71$$

This value is somewhat low.

If your system is running in an PR/SM environment, the calculation has to be performed using LPAR BUSY and the number of logical processors.

- The paging rates for TSO are

SINGLE 16.7 page-ins per second

BLOCK 6.2 page-ins per second

This may be worth further investigation.

See Chapter 3, “Analyzing processor activity” and Chapter 5, “Analyzing processor storage activity” for further analysis.

Virtual Storage Activity report

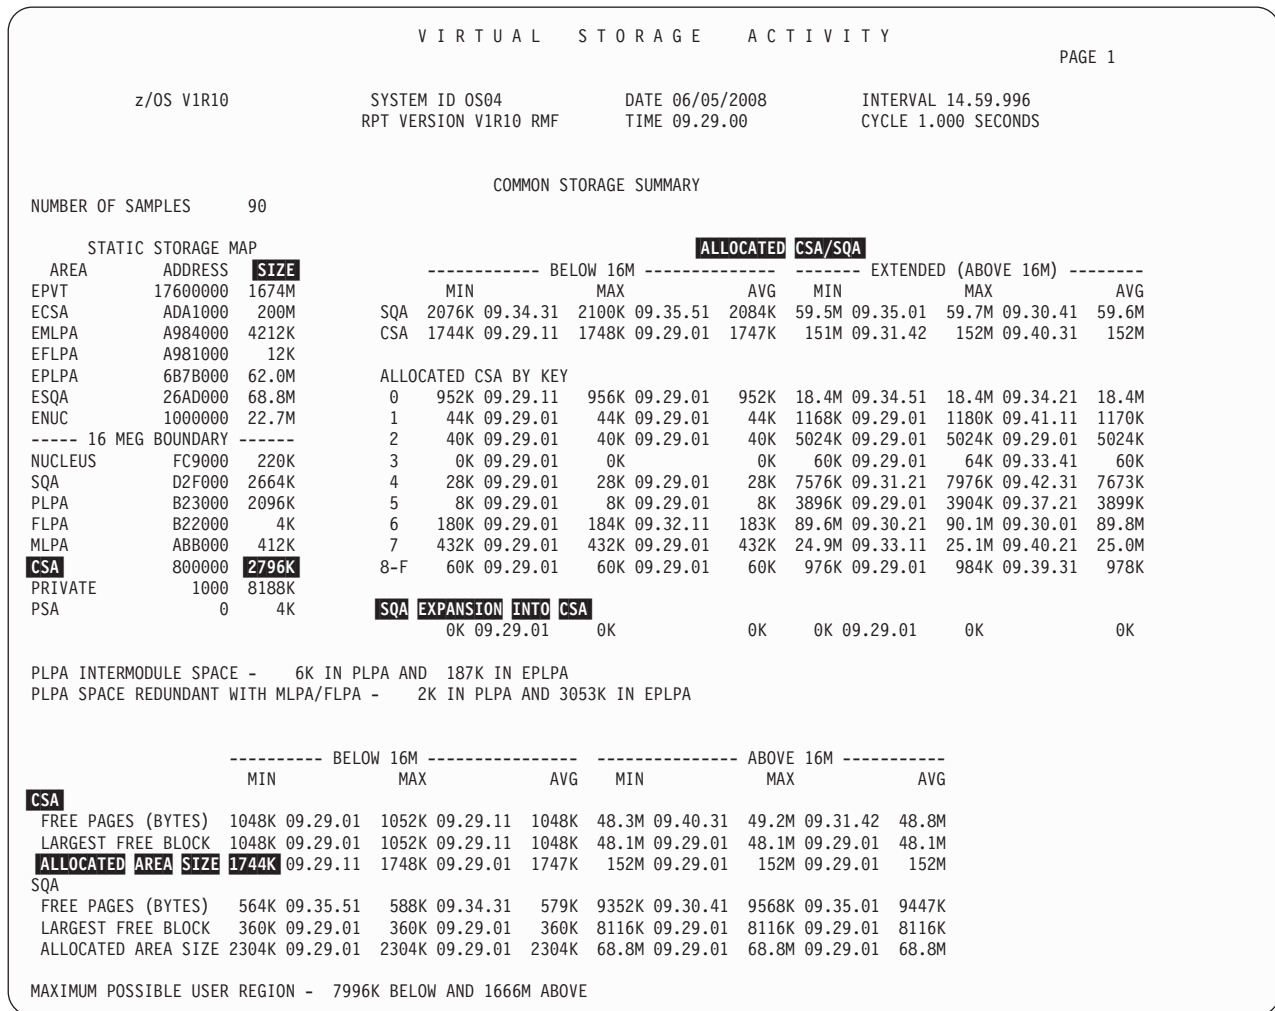


Figure 35. Monitor I Virtual Storage Activity Report - Common Storage Summary

Rules-of-Thumb

Maximum Allocated Percent

Calculate this value for the CSA, the ECSA, the SQA, and the ESQA. For example, for the CSA:

$$\text{Maximum Allocated Percent(CSA)} = \frac{\text{CSA ALLOCATED AREA SIZE}}{\text{CSA SIZE}} * 100$$

For CSA and ECSA, this value should be less than 65%.

The virtual storage that is most volatile is the SQA and CSA, both above and below the 16MB line. This is a number that is readily tracked on the Virtual Storage report on either a daily or weekly basis.

RMF calculates the size of this area as the difference between the highest and lowest address occupied by allocated storage, and includes all free blocks that lie between allocated blocks. Significant segmentation causes this number to be much larger than the amount of storage actually used.

If the size allocated for ESQA is too small or is used up, the system attempts to steal pages from the ECSA. When both areas are full, the system allocates space from the SQA and CSA below the 16MB line. SRM will attempt to reduce the demand for resources. If no storage is available, the result is a system failure.

If the size of SQA below 16M is too small, additional data will spill into the CSA below the 16MB line. If no storage is available, the result is a system failure.

- Is CSA/ECSA allocated percent high?
 - Review ECSA, CSA, ESQA, and SQA sizes
 - Increase CSA size: cleanup LPA, split address spaces (MRO for example), use Version/Release products that exploit ESA, or convert applications to exploit ESA (for example to Cobol II)
- Is SQA/ESQA allocated percent low?
 - ESQA or SQA may be too large.
- Is SQA EXPANSION INTO CSA low?
 - SQA size may be too large; you may be able to decrease SQA size and increase CSA size.

Using the Monitor III Common Storage report is a very good way to understand which AS has the storage.

Report Analysis

The maximum allocated percent for the CSA is:

$$\frac{1744K}{2796K} * 100 = 62.4\%$$

That concludes there is no problem in this area.

The SQA EXPANSION INTO CSA is zero, that means the SQA is large enough.

Where do you go from here?

By now you will have found the major cause of your performance problem. The following chapters will go into more detail to help you further.

As you know, performance tuning is an iterative process and bear in mind:

- Take a base measurement as your starting point
- Only change one thing at a time
- Predict how your key measurements will change
- Measure your change against your expectation

System indicators

Chapter 3. Analyzing processor activity

Let's Analyze Your Processors

This chapter discusses how to analyze a processor problem. It takes the following approach:

- What RMF values can you look at in a performance problem situation?
- What do these fields mean?
- What can you do to solve the problem?

We make recommendations for two cases:

- An application being *delayed* by higher priority work
- An application *using* excessive CPU, or overall CPU under stress

As you continue through this chapter, bear in mind that tuning the processor is generally a high-effort/low-return activity. Quick fixes are rare. However, we will discuss the tuning activities that can apply.

Do you have a processor problem?

There are many views on what constitutes a processor performance problem. You should conclude that you have a processor performance problem because your:

- Service level objectives are being exceeded.
- Users are complaining about slow response.
- AND, the CPU indicators (discussed in the following section) show signs of stress. Or better still, you see a processor delay directly in Monitor III.

We suggest there are two main ways to approach this problem analysis, depending on how you are alerted:

- If a user or group of users complains about response times then start with Monitor III.
- If your normal daily or weekly performance tracking procedures show a trend or problem developing, then start with Monitor I.

Table 2 shows for each monitor, which fields in which RMF reports can detect and resolve a processor performance problem.

Table 2. Processor Indicators in Monitor III and Monitor I

	REPORTS	DELAY Fields	USING Fields
Monitor III	GROUP	Avg Delay PROC	AVG USG PROC
	DELAY	%Delayed for PRC	
	WFEX	Speed *PROC	
	PROC	DLY%	USG% App1% CP
	CPC	WLM Capping %	
Monitor I	CPU Activity		BUSY TIME PERCENTAGE
	Partition Data		PHYSICAL PROCESSOR TOTAL

Monitor III indicators

This section lists the Monitor III indicators that may be used to alert you to a potential processor problem.

GROUP report

```

RMF V1R10 Group Response Time
Command ===>                               Scroll ===> HALF
Samples: 100      System: PRD1 Date: 04/07/08 Time: 10.32.00 Range: 100 Sec
Class: BATCH      Period: 1   Description: Primary Batch
Primary Response Time Component: Using the processor

WFL  Users      Frames  Vector  EXCP  PgIn  TRANS  --- Response Time ---
%    TOT  ACT    %ACT    UTIL  Rate  Rate   -- Ended TRANS-(Sec) -
60   30   3     25     0     2.1   0.5    WAIT  EXECUT  ACTUAL
      830.5 152.7 983.3

                        -AVG USG-  -----Average Delay-----
Average Users      Total  PROC  DEV  PROC  DEV  STOR  SUBS  OPER  ENQ  OTHER
Response Time ACT 152.7  88.3 4.67  34.6 2.80 0.47 0.00 0.00 23.8 0.00

                        ---STOR Delay---  ---OUTR Swap Reason---  ---SUBS Delay---
Page  Swap  OUTR  TI  TO  LW  XS  JES  HSM  XCF
Average Users      0.01 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00
Response Time ACT  0.47 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00 0.00
    
```

Figure 36. Monitor III Group Response Time Report

This report presents information for a specific service class period.

DELAY Indicator

Field: Average Delay PROC for Response Time ACT

Description: This is a component of average response time, showing the average number of seconds, transactions active during the range period in the specific service class period, were delayed by the lack of CPU. It is the CPU queue time.

Guideline: Look at the different Average Delay values, in the line Response Time ACT. If the PROC value is the largest, continue your processor investigation. Look at the different service classes (coming back to the Primary Panel and choosing option 3) and start with the service class showing the largest value.

Problem Area: This points out processor delays for the service class listed.

Potential Solution:

- See "DELAY for processor: Guidelines for improvement" on page 73
- See "Processor USING: Trimming activities" on page 76

USING Indicator

Field: AVG USG PROC for Response Time ACT

Description: This is a component of average response time, namely the average number of seconds, transactions belonging to the service class were using the CPU during the range period.

Guideline: Check all the service classes and start with that one showing the average largest processor consumption.

Problem Area: This points out processor consumption.

Potential Solution:

- See "Processor USING: Trimming activities" on page 76

DELAY report

RMF V1R10 Delay Report											Line 1 of 44		
Command ==>											Scroll ==> HALF		
Samples: 100 System: PRD1 Date: 04/07/08 Time: 10.32.00 Range: 100 Sec													
Name	Service CX Class	WFL Cr	USG %	DLY %	IDL %	UKN %	PRC	----- % Delayed for -----				Primary Reason	
								DEV	STR	SUB	OPR	ENQ	
BHOLEQB	B BATCH	0	51	0	49	0	0	0	0	0	51	0	SYSDSN
CATALOG	S SYSSTC	0	0	3	0	97	0	0	3	0	0	0	LOCL
BAJU	T TSO	14	2	12	85	1	0	0	12	0	0	0	LOCL
BHOLNAM4	B BATCH	62	31	19	0	32	18	0	1	0	0	0	BHOLPRO2
JES2	S SYSSTC	67	13	7	0	83	1	6	0	0	0	0	SYSPAG
BHOLPRO1	B BATCH	72	72	28	0	0	28	0	0	0	0	0	BHOLPRO2
BHOLPRO2	B BATCH	74	74	26	0	0	26	0	0	0	0	0	BHOLPRO1
RMF	S SYSSTC	75	3	1	0	96	1	0	0	0	0	0	BHOLPRO2
MASTER	S SYSTEM	92	23	2	0	76	0	2	0	0	0	0	SYSPAG
BHOL	T TSO	97	29	1	67	3	1	0	0	0	0	0	CONSOLE
CONSOLE	S SYSSTC	100	2	0	0	98	0	0	0	0	0	0	
BHOLST09	B BATCH	100	1	0	99	0	0	0	0	0	0	0	
VTMLCL	S SYSSTC	100	1	0	0	99	0	0	0	0	0	0	

Figure 37. Monitor III Delay Report

This is a good report to look at for a quick snapshot of 'Who is delayed' for CPU and for how long. The report is sorted by workflow percentage (WFL % column) in ascending order. Check this report to see if any priority work is being delayed.

DELAY Indicator

Field: % Delayed for PRC

Description: This gives for each address space (AS), the percentage of the measured time that a transaction running in that AS was delayed by the lack of CPU.

Guideline: Look at the largest %Delayed for PRC value. Compare it with the DLY % value of the same job. If the PRC value is the main part of the delay, look at this job.

You can put the cursor on the line describing the job that may be a problem and press ENTER: a new panel is displayed (see Figure 43 on page 73) that lists the main cause of the delay and suggests some possible actions to reduce the delay.

Problem Area: This points out a delay for processing.

Potential Solution:

- See "DELAY for processor: Guidelines for improvement" on page 73
- See "Processor USING: Trimming activities" on page 76

WFEX report

```

RMF V1R10 Workflow/Exceptions                               Line 1 of 4
Command ==>>>                                           Scroll ==>> HALF
Samples: 100      System: PRD1 Date: 04/07/08 Time: 10.32.00 Range: 100 Sec

----- Speed (Workflow) -----
Speed of 100 = Maximum, 0 = Stopped      Average CPU Util: 100 %
Name      Users Active      Speed      Name      Users Active      Speed
*SYSTEM    41      4      63
ALL TSO     2      0      70
ALL STC     32     1      79
ALL BATCH   7      3      59
*PROC      24     3      73
*DEV        6      1      33
*MASTER*    1      0      92

----- Exceptions -----
Name      Reason      Critical val. Possible cause or action
*ECSA*    SECS% > 85    90.6 %      System ECSA usage 91 %.
*STOR     TSQA0 > 0     784K bytes   SQA overflow into CSA 784K.
BHOLEQB   ENQ -SYSDSN   51.0 % delay ENQ.TEST.DAT
    
```

Figure 38. Monitor III Workflow/Exceptions Report

DELAY Indicator

Field: *PROC Speed (Workflow)

Description: This is a good global indicator for CPU performance. If equal to 100%, no work has been delayed by the CPU. If equal to 50%, this means that half of the attempts to use CPU were denied because the CPU was busy (that is, five out of ten ready AS were being delayed). This field does not deal with how busy the CPU is, but rather the amount of delay that exists.

Guideline: If this value is low (less than 40%), it may be due to CPU delay. If this value is greater than 80%, there is no CPU problem. Between these two values, it depends on the characteristics of the environment.

Problem Area: This points out a delay for processing.

Potential Solution:

- See "DELAY for processor: Guidelines for improvement" on page 73
- See "Processor USING: Trimming activities" on page 76

PROC report

RMF V1R10 Processor Delays									
Command ==>>								Line 1 of 138	
								Scroll ==> HALF	
Samples: 60		System: MVS1		Date: 11/28/08		Time: 09.10.00		Range: 60 Sec	
Jobname	Service CX Class	CPU Type	DLY %	USG %	EApp1 %	----- % Name	----- % Name	----- % Name	----- % Name
WSWS7	0 OMVS	CP	11	46	59.4	9 *ENCLAVE	7 DBS3DIST	7 WSP1S2F	
WSP1S2FS	SO WASCR	CP	4	4	42.5	2 DBS3DIST	2 WSWS7	2 VTAM44	
		AAP	6	0	98.4	6 *ENCLAVE			
WSP1S6FS	SO WASCR	CP	0	0	5.3				
		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 WSWS7	2 DBS3DIST	2 U078069	
WSP1S4F	SO WASCR	CP	0	0	0.1				
		AAP	2	0	0.4	2 WSP1S6F			
U078068	0 OMVS	CP	2	0	0.2	2 XCFAS	2 WSWS7	2 *ENCLAVE	
DBS3DIST	SO DB2HIGH	CP	0	78	111.0				
		IIP	0	2	21.3				

Figure 39. Monitor III Processor Delays Report

This report displays all AS waiting for or using the processor during the range period. The report is sorted by descending overall delay percentages: the first line is the job you need to look at first.

DELAY Indicator

Field: DLY %

Description: The percentage of time an AS is delayed because of contention for the processor during the range period. For the multitask AS, RMF reports only one delay when several tasks are delayed at the same time.

Guideline: Use the cursor sensitivity: put the cursor on the name of the AS you want to analyze (usually we start with the most delayed job). Hence the Job Delay report (as in Figure 43 on page 73) is displayed; it allows further investigation as explained later in this chapter.

Problem Area: Points out a processor delay.

Potential Solution:

- Start by identifying the holding jobs that are causing the most delay
- See "Determine CPU 'holders'" on page 73

USING Indicator

Field: USG % and EApp1%

Description:

- USG %: The percentage of time an AS is using the processor during the range period. This is a sampled single state value.
- EApp1%: Percentage of CPU time as sum of TCB time, global and local SRB time, preemptable or client SRB time, and enclave CPU time consumed within this address space. This is a measured multi-state value: if the address space is using more than one processor, this value can exceed 100%.

Guideline: Start with the AS with the largest EApp1% value, then look at its USG% value: if the USG% value is large, you have a heavy processor consumer.

Problem Area: Points out processor consumption.

Potential Solution:

- See "Processor USING: Trimming activities" on page 76

USING Indicator**Field:** BUSY TIME PERCENTAGE**Description:** The percentage of the interval time that the processor was busy executing instructions; it includes the non-captured CPU-Time.**Guideline:** If this value is greater than 85% AND the IN READY QUEUE value for $\leq N$ is below 80%, as is the case in the sample report from Figure 40 on page 68, then you have some work with CPU delay.**Note:** Depending on your environment, the guidelines may be very different:

- In a batch or scientific environment, a 100% CPU busy is usually not a problem.
- In a transaction environment, user response times usually increase when CPU busy goes over 80%. The type of the application and the number of CPs available to it can influence how high this CPU busy value can go before response time suffers.

Problem Area: Points out processor consumption.**Potential Solution:**

- See "Processor USING: Trimming activities" on page 76

Note: BUSY TIME PERCENTAGE by itself is not enough as a CPU indicator. You must remember that the distribution of the arrival rate of work demanding CPU is not constant. So even at a low utilization level, a queue may build up. The IN READY value $\leq N$ gives the percentage of time when no contention for CPU was detected, that is, the number of AS with a ready process (TCB or SRB) was equal to or less than the number of CPs. In our example, as the total is 58.8%, then 41.2% of the time (100%-58.8%) you have work delayed.

This is not to say that you cannot run a z/OS system at higher utilizations, on up to 100% busy. You can. Just be aware that the busier the CPU, the longer the CPU delay for lower-priority workloads will be. The more low-priority work you have, and the fewer non-CPU bottlenecks you have (e.g. responsive I/O), the busier you will be able to drive your CPU and still maintain good response for your high-priority workloads.

Mon I - Processor analysis




Figure 41. CPU Contention Report. Spreadsheet Reporter macro RMFY9SUM.XLS (System Overview Report - OneCpuCont)

This report provides some more detailed data on CPU contention. In addition to the contention value which is given by the percentage of time when at least one job was ready and waiting for the processor, you see here the percentages when at least two or three jobs were waiting.

Partition Data report

PARTITION DATA REPORT																
z/OS V1R10			SYSTEM ID NP1			DATE 11/28/08			INTERVAL 14.59.678			PAGE 2				
			RPT VERSION V1R10 RMF			TIME 09.30.00			CYCLE 1.000 SECONDS							
MVS PARTITION NAME			NP1			NUMBER OF PHYSICAL PROCESSORS			9		GROUP NAME		N/A			
IMAGE CAPACITY			100			CP			7		LIMIT		N/A			
NUMBER OF CONFIGURED PARTITIONS			9			ICF			2							
WAIT COMPLETION			NO													
DISPATCH INTERVAL			DYNAMIC													
----- PARTITION DATA ----- -- LOGICAL PARTITION PROCESSOR DATA -- -- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --																
NAME	S	----MSU----			-CAPPING--		PROCESSOR- NUM	TYPE	----DISPATCH TIME DATA----		LOGICAL PROCESSORS		--- PHYSICAL PROCESSORS ---			
		WGT	DEF	ACT	DEF	WLM%			EFFECTIVE	TOTAL	EFFECTIVE	TOTAL	LPAR MGMT	EFFECTIVE	TOTAL	
NP1	A	20	100	10	NO	62.2	1.2	CP	00.04.27.302	00.04.27.519	24.86	24.92	0.01	4.24	4.25	
NP2	A	1	0	1	YES	0.0	4	CP	00.00.21.680	00.00.22.083	0.60	0.61	0.01	0.34	0.35	
NP3	A	10	5	8	NO	3.3	1.0	CP	00.03.35.761	00.03.35.859	23.97	23.98	0.02	3.41	3.43	
NP4	A	300	95	155	NO	0.0	6.0	CP	01.12.05.405	01.12.06.405	80.10	80.15	0.01	68.68	68.69	
NP5	A	200	50	52	NO	0.0	4	CP	00.24.11.147	00.24.11.311	40.30	40.31	0.01	23.02	23.03	
PHYSICAL									00.00.03.103				0.05		0.05	
TOTAL									01.44.41.295 01.44.46.280				0.11		99.69 99.80	
CFC1	A	DED					1	ICF	00.14.59.611	00.14.59.625	99.98	99.99	0.01	99.95	99.96	
CFC2	A	DED					1	ICF	00.00.00.000	00.00.00.000	0.00	0.00	0.00	0.00	0.00	
PHYSICAL									00.00.00.065				0.03		0.03	
TOTAL													0.04		99.95 99.99	
CB88	D															
CB89	D															

Figure 42. Monitor I CPU Activity Report - Partition Data Report Section

This report provides data about configured partitions in LPAR mode. Only partitions active at the end of the duration interval are reported.

USING Indicator

Field: PHYSICAL PROCESSORS TOTAL

Description: The total amount of time the physical processor resource was assigned to a partition AND to the management of the LPAR itself. The partition identified by the name *PHYSICAL* is not a configured partition: this is a way to report uncaptured time which was used by LPAR but could not be attributed to a specific logical partition.

Guideline: If the total physical processor complex is 100% busy for one partition, this LPAR could be constrained by other LPARs.

Problem Area: Points out processor consumption.

Potential Solution: Check the partition weights, capping, and number of CPs defined, to be sure you are allocating the CP resources the way you intended.

- See Appendix A. PR/SM LPAR considerations

Is your processor idle due to other bottlenecks?

If your system shows no processor indicators under stress, that does not mean you don't have a potential processor problem: You may have some other bottleneck in the system preventing work from using the processor. If your daily processor-busy peaks start to plateau, before reaching 100%, this could be an indication of such a bottleneck. The major ones are:

- Storage
- I/O
- ENQ
- Operator
- Network

If you think this might be the case, see the other chapters for guidance on finding bottlenecks.

Remember that all CPUs wait at the same speed...

DELAY for processor: Guidelines for improvement

We assume in this section that you have found that your application's primary delay is due to the processor. This means that other workloads are using the processor when your application wanted it. In this section, we will discuss what can be done to decrease processor delay for a workload.

The section "Processor USING: Trimming activities" on page 76 is worth checking when you have a delay problem as well: the delay can be caused by an AS using too much of the processor. Trimming this AS solves the delay problem.

Determine CPU 'holders'

Bringing up the Monitor III Processor Delays report (shown in 39), use the cursor sensitivity. By positioning the cursor on the line describing the job you want to look at, and pressing ENTER, a new panel is displayed, on which the AS delaying the workload are named, as shown in Figure 43.

JOB report

```

RMF V1R10 Job Delays                               Line 1 of 1
Command ==>                                         Scroll ==> HALF
Samples: 100    System: PRD1 Date: 04/07/08 Time: 10.32.00 Range: 100 Sec
Job: BHOLPRO2    Primary delay: Job is waiting to use the processor.
Probable causes: 1) Job BHOLPRO1 may be looping.
                  2) Higher priority work is using the system.
                  3) Improperly tuned dispatching priorities.

----- Jobs Holding the Processor -----
Job:    BHOLPRO1    Job:    BHOL    Job:    *MASTER*
Holding: 4%    Holding: 3%    Holding: 1%
PROC Using: 98%    PROC Using: 3%    PROC Using: 1%
DEV Using: 0%    DEV Using: 8%    DEV Using: 4%
----- 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 0033 BATCH    1  94 94  0  6  0  0  6  0  0  0  0  0  0  BHOLPR01F

```

Figure 43. Monitor III Job Delays Report

Some probable causes of trouble are mentioned as well. This is the best starting point. See which AS are delaying your workload.

Processor analysis

Review dispatching priorities

The *dispatching priority (DP)* of each address space is specified dynamically according to the goals of the respective service class. You should reconsider your service policy definition if the workload delaying your application has a higher dispatching priority.

The Monitor III ASD report shows address space dispatching priorities (DP PR).

ASD report

```
RMF - ASD Address Space Data Report                               Line 1 of 64
Command ==>                                                       Scroll ==> HALF

                                CPU= 67/ 55 UIC=2540 PR=  0           System= SYSF Total

13:43:53          S C R  DP  CS  ESF  CS  TAR  X  PIN  ES  TX  SWAP  WSM
JOBNAME  SRVCLASS P L LS  PR  F      TAR  WSS  M  RT  RT  SC  RV  RV

ANTAS000 STCDEF  1 NS  F1  871
MASTER*  SYSTEM  1 NS  FF  141           0  0.0  0.0  0  0
PCAUTH   STCDEF  1 NS  74  26           0  X  0.0  15  0  175
RASP     STCDEF  1 NS  FF  48           0  X  0.0  15  0  0
TRACE    STCDEF  1 NS  76  78           0  X  0.0  15  0  0
GRS      SYSTEM  1 NS  FF 1381           0  X  0.0  15  0  175
DUMPSRV  SYSTEM  1 NS  FF  31           0  0.0  15  0  175
CONSOLE  SYSTEM  1 NS  FF  62           0  X  0.0  15  1  0
ALLOCAS  SYSTEM  1 NS  71  317           0  X  0.0  15  1  0
LLA      SYSSTC  1 NS  FC  331           0  X  0.0  24  0  0
VTAM     SYSSTC  1 NS  FD  518           0  0.0  44  0  175
NETVS    SYSSTC  1 NS  FC  77           0  X  0.0  32  1  175
NETVIEW  SYSSTC  1 NS  FC  497           0  0.0  30  0  175
SMF      SYSTEM  1 NS  FF  65           0  X  0.0  15  0  175
RMF33    SYSSTC  1 NS  FF  128           0  0.0  38  0  175
DFHSM    SYSTEM  1 NS  7D  674           0  0.0  671  1  175
AMSAQTS  SYSTEM  1 NS  74  37           0  0.0  50  1  175
VLF      SYSTEM  1 NS  78  2297           0  X  0.0  22  0  175
JES2     SYSTEM  1 NS  FE  827           0  0.0  25  0  175
SMS      SYSTEM  1 NS  79  59           0  X  0.0  24  0  175
JOEUSER  BTCHPRD  1 WT TI FF  764           0  X  0.0  15  0  18 -1000
```

Figure 44. Monitor III Address Space State Data Report

Check for CPU delay due to other LPARs

If you are running as an LPAR sharing CPs with other LPARs, and your MVS system is not 100% busy, you might ask: Are the logical CPs not working more because there is no demand (wait time) or because the LPAR is not getting enough CPU?

One way to check if other LPARs are constraining your system is to do the following: Calculate the percentage of time when there was CPU delay by adding up the percentages from lines 'N + 1' to 'N + 80' from the Monitor I CPU Activity report Figure 24 on page 44. If this total is more than your CPU BUSY PERCENTAGE, then you have another LPAR constraining you (you have CPU delay % > CPU busy %). If this is the case, check the partition weights, capping, and number of CPs defined, to be sure you are allocating the CP resources the way you intended.

Check LPAR balancing using the CPC capacity report

You can use the Monitor III CPC Capacity report to check whether the capacity being defined for each partition is sufficient to provide the service that all applications are expecting.

```

RMF V1R10 CPC Capacity
Command ==>
Samples: 59      System: Z201  Date: 11/28/08  Time: 13.16.00  Range: 60  sec
Line 1 of 12
Scroll ==> HALF

Partition: Z2      2094 Model 716
CPC Capacity: 937  Weight % of Max: 74.2  4h Avg: 80  Group: CGRP0010
Image Capacity: 77  WLM Capping %: 5.4  4h Max: 185  Limit: 100*

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

*CP
TZ1      130  122  NO  4.3   11.2  12.5  0.4   3.4   3.8
Z1       150   89  NO  5.2    9.3   9.6   0.1   3.5   3.6
Z2       50   58  NO  2.1   11.5  12.8   0.2   1.7   1.9
Z3       N/A   0   YES 2.4    8.8  10.6   0.3   1.5   1.8
PHYSICAL                                0.1   0.1

*ICF
CF1              1   99.9  99.9  0.0   7.1   7.1
CF2              1    0.0   0.0  0.0   0.0   0.0
PHYSICAL                                0.1   0.1

```

Figure 45. Monitor III CPC Capacity Report

The report provides information about processor resource consumption of each partition, and it displays the 4-hours average consumption. If this average is higher than the defined image capacity, WLM will start capping this partition. In the above example, the 4h average is higher than the image capacity. Therefore, partition Z2 will get less processor resources in the following intervals until the average is below the defined capacity. This might result in processor constraints and performance problems. If this is not just a temporary but a permanent problem, one should consider about increasing the image capacity for this partition.

Processor USING: Trimming activities

This section describes the activities that help decrease the processor consumption. In case of delay, this list of activities can be used to trim the AS that uses too much processor resource.

Check for loops

Check for loops in the following reports:

- "PROC report"
- "Job report" on page 77

PROC report

```

RMF V1R10 Processor Delays                               Line 1 of 71
Command ==>                                           Scroll ==> HALF
Samples: 100      System: PRD1 Date: 04/07/08 Time: 10.32.00 Range: 100 Sec

```

Jobname	Service CX Class	CPU Type	DLY %	USG %	EAppl %	----- Holding Job(s) -----		
						% Name	% Name	% Name
BOE1HCP\$	B BATCH	CP	10	11	9.8	9 BJPI	8 BTFL	4 CATALOG
BTFL	T TSO 2	CP	9	91	83.7	8 BJPI	3 BOE1MGR\$	3 BOE1HCL\$
BOE1TBL\$	B BATCH	CP	9	28	26.7	8 BTFL	7 BJPI	5 JES3
BOE1FIR\$	B BATCH	CP	9	21	20.5	8 BTFL	7 BJPI	7 JES3
BOE1HCL\$	B BATCH	CP	7	19	15.3	7 BJPI	7 BTFL	4 BOE1MGR\$
BOE1DEV\$	B BATCH	CP	7	11	9.7	6 BJPI	5 BTFL	4 TESTRTX
BOE1MGR\$	B BATCH	CP	6	15	12.7	6 BJPI	5 BTFL	3 BOE1HCP\$
BJPI	T TSO	CP	4	96	87.9	4 BTFL	3 JES3	2 XXVTAM
BOE1HIT\$	B BATCH	CP	3	13	10.4	3 BTFL	3 BJPI	2 CATALOG
BOE1HOM\$	B BATCH	CP	3	5	3.0	3 VIBTS	3 BJPI	2 BOE1DEV\$
JES3	S SYSSTC	CP	2	21	22.0	2 BJPI	1 BTFL	1 *MASTER*
BRTV	T TSO	CP	2	6	4.9	2 BTFL	2 BJPI	1 BOE1HCL\$
BRMC	T TSO	CP	2	5	5.0	2 TESTRTX	1 CATALOG	1 BJPI
XXVTAM	S SYSSTC	CP	1	22	15.8	1 JES3	1 BOE1TBL\$	1 AKS

Figure 46. Monitor III Processor Delays Report

Detect loops by inspecting DLY % and USG % in this report for individual address spaces. If the sum of both is 100%, you could be in an enable loop.

You can check this by moving the cursor to jobname BTFL, and press ENTER. This links to the Job Delays report.

Job report

```

RMF V1R10 Job Delays                               Line 1 of 1
Command ==>                                         Scroll ==> HALF

Samples: 100    System: PRD1 Date: 04/07/08 Time: 10.32.00 Range: 100 Sec
Job: BTFL      Primary delay: Job is waiting to use the processor.

Probable causes: 1) Job BJPI may be looping.
                  2) Higher priority work is using the system.
                  3) Improperly tuned dispatching priorities.

----- Jobs Holding the Processor -----
Job:          BJPI          Job:          BOE1MGR$          Job:          BOE1HCL$
Holding:      8%           Holding:      3%           Holding:      3%
PROC Using:   96%         PROC Using:   15%         PROC Using:   19%
DEV Using:    0%          DEV Using:    5%          DEV Using:    8%

----- 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
T 0148 TSO      2   91  91  0   9  0  0   9   0  0  0  0  0   0 BJPI

```

Figure 47. Monitor III Job Delays Report

A good indication of a loop is the fact that no other delays than processor delay are shown.

This is also detected by Monitor III which points to another job that might be looping.

Check with the owner of this address space (in this example, it is a TSO user), before deciding how to proceed - don't cancel it too quickly!

Check for slip traps

```

RMF V1R10 Workflow/Exceptions                               Line 1 of 12
Command ==>>>                                           Scroll ==>> HALF

Samples: 100      System: PRD1 Date: 04/07/08 Time: 10.32.00 Range: 100 Sec

----- Speed (Workflow) -----
                Speed of 100 = Maximum, 0 = Stopped      Average CPU Util: 81 %
Name           Users Active   Speed           Name           Users Active   Speed
*SYSTEM        505         13             54
ALL TSO        433         10             55
ALL BATCH       2           0             42
ALL STC        70          2             55
ALL ASCH                               Not avail

----- Exceptions -----
Name           Reason           Critical val. Possible cause or action
* SLIP *      SLIP PER TRAP                               SLIP ID=SR01 is active.
BEVK          Rate < 2.0       1.220 /sec    Tx rate is 1.220 /s.
BSHR          STOR-COMM        23.1 % delay
CSAHOG        JCSA% > 15       18.3 %        Job CSA usage 18 %, system 57 %.
POK063        DAR > 20         23.22 /sec    I/O rate is 23.22 /s on volume POK063.
DATAPK        Not avail                               Volume DATAPK is not mounted.
    
```

Figure 48. Monitor III Workflow/Exceptions Report

Use the workflow/exceptions report to watch for signs of a SLIP PER active trap. Program Event Recording (PER) is a function of S/390[®] architecture that generates program interrupts and uses the CPU unproductively (when not used to analyze a problem).

PER is started by a SLIP trap and has significant overhead. The WFEX report shows any active SLIP traps using PER. Check to see if the trap is still needed.

The * SLIP * exception line, in yellow, is always shown as the first exception. This line cannot be suppressed when the SLIP is active.

Reduce I/O

Rules of Thumb

- 1000 I/Os per second consume about 7.6% of a 9672 Z87 system.
- Every 160 DASD demand pages/sec uses about 1.0% of a 9672 Z87 system.

You can avoid I/O by:

- Using Data-In-Memory (DIM) techniques
- Increasing block sizes
- Exploring SAM/E chain scheduling possibilities
- Moving I/O bound activities to other shifts
- Decreasing the paging

Avoiding I/O may help in saving CPU. Consequently you improve your performance (response and throughput); however chances are that your CPU is going to be busier, productively, because you decrease I/O delay.

Decrease monitor overhead

- Check for obsolete operations CLISTS.
- Close down duplicate/redundant monitoring.
- Reduce scope and/or frequency of other monitors.
- Control and decrease SMF records generation.
- Increase RMF cycle, control the use of RMFMON and RMF commands.

Tune system address spaces

The system AS that usually consume the most processor are:

- Console AS, due to commands and message traffic in MVS console. Consider message filtering to reduce traffic.
- VTAM AS. Consult a TP specialist.
- Master Scheduler AS, mainly due to paging activities in SRB mode. See Chapter 5, “Analyzing processor storage activity” for ways to reduce paging.
- JES AS. Decrease batch and some TSO activity (drastic action for small payback).

Tracing

GTF traces can be used to see the use of supervisor code in your TCB time (SVCs and PCs).

Sometimes you may see a lot of GETMAIN, EXCP, or CLOSE activities in your GTF trace. This can be a clue for application redesign.

Calculate your capture ratio

If your prime shift capture ratio is below about 80%, check the following:

- How much CPU am I spending with swapping (non-captured)?
- How much CPU am I spending with paging (some non-captured)?

Redesign application

Sometimes the only answer may be to redesign an application completely or partly, or change the user’s expectations.

Install more CPU power

When you have done all you can, and the problem does not go away, you can still resort to buying more CPU power.

Summary

To analyze processor problems you should:

- Determine if this is a *DELAY* due to another user or a *high USING* processor consumption.
- Look at the Monitor III Job Delay report first in case of *DELAY*.
- Reduce all the processor overhead that you can: it will usually be a very small amount.

Chapter 4. Analyzing I/O activity

Let's Analyze Your I/O Subsystem

This chapter discusses how to analyze an I/O problem. We review:

- How to perform an overall health check of your I/O subsystem
- How to improve I/O performance with the Enterprise Storage Server
- How to analyze and solve I/O problems related to
 - Cache
 - DASD
 - Tape

We start by reviewing some of the key concepts, terms, and functions that will be used throughout the chapter.

Do you have an I/O problem?

You can conclude that you have an I/O problem if:

- Service level objectives are being missed
- Users are complaining about response times
- AND, the I/O indicators discussed in this chapter show signs of stress. Or better still, you see high DEV DLY or USG for an important workload directly in Monitor III reports (for example, SYSINFO, DELAY, DEV, or DEVR).

Here are the primary RMF indicators you can use to assess your I/O subsystem. After listing these indicators, we will go on to discuss each one in more detail, with guidelines on values that may indicate contention and suggestions for further action.

RMF I/O INDICATORS

DASD LCU level	DASD Device level	Tape
I/O rate	I/O rate	I/O rate
Response Time	Response Time	Response Time
<ul style="list-style-type: none"> • Disconnect • Connect • Queue • Pend 	<ul style="list-style-type: none"> • Disconnect • Connect • Queue • Pend 	<ul style="list-style-type: none"> • Disconnect • Connect • Queue • Pend
Cache Statistics	Utilization Cache Statistics	
<ul style="list-style-type: none"> • Read hits • DASD fast write hits • Cache Fast write hits • Staging • DFW retries • Bypass • Inhibit • Off • ASYNC 	<ul style="list-style-type: none"> • Read hits • DASD fast write hits • Cache Fast write hits • Staging • DFW retries • Bypass • Inhibit • Off • ASYNC 	

Never start any tuning effort simply because one indicator seems to be a problem. Always check other related indicators to build a good understanding of your I/O subsystem (and overall system). Check to see which workloads are delayed, and by how much. Do all this to make sure you have the big picture well in hand, and only then begin to plan your tuning efforts.

As you analyze I/O problems, remember the following:

- **Manage to service level agreements.** No I/O response time, or other indicator, is inherently good or bad. The basic premise of performance management is that you want good I/O response times for important applications. What is "good" response time? The answer, of course, is that it depends. One way you can determine good response is to measure the response times of the volumes and logical control units (LCUs) when your applications are running well, meeting the formal or informal service level agreements.

- **Keep the balanced systems concept in mind.** Before embarking on I/O subsystem analysis, review the CPU and processor storage resources. Keeping the big picture in mind will help you make better decisions on matters that affect overall system throughput, and on trade-offs such as using additional processor storage to reduce I/O rates.
- **Proceed in a top-down manner.** This implies three things:
 - If you think you see a problem, start with the basic facts. Which workload has the problem? What are the I/O rates and response times? Are service levels really being impacted? A "bad" response time may not be worth spending time on, if the I/O rate is low, or the workload has low priority.
 - Analyze your overall subsystem before trying to tune for specific DASD volumes or tape subsystems.
 - Analyze any response-time problem to see which component is greatest (disconnect, queue, etc.), and address that component first.
- **Optimize DASD data placement.** The best response time is achieved if the I/O never happens. This may be accomplished by a smarter program that avoids the I/O, by buffering, or by data-in-memory techniques where the I/O is resolved somewhere in processor storage. The next best response time is achieved by resolving the I/O directly in a cache control unit, a cache hit. The worst response time occurs when a DASD I/O must be resolved directly from a DASD device.

I/O subsystem health check

This section describes how to perform an overall health check of your I/O subsystem. Use this section to help you establish a baseline measurement of I/O subsystem performance, set up your performance monitoring process, and provide a base of knowledge from which to begin performance tuning.

Understand your business priorities

All devices and workloads do not have the same priority. Before starting any tuning effort, be sure you understand the important business applications and their formal or informal service objectives. Know which volumes serve the important workloads.

Review CPU and processor storage

Alleviating an I/O bottleneck will mean that more work will be delivered to the CPU for processing. Keep the 'balanced system' concept in mind, making sure that CPU and processor storage resources are available. Review processor storage to see how much might be available for increased I/O elimination using data-in-memory techniques.

Establish a baseline measurement

Establishing a set of measurements for when your system is running well, provides a good base for later comparison. This will tell you what numbers are good for YOUR environment, which is better than simply following any rule of thumb guideline.

Measure the I/O rate, response times, and the cache activity at the LCU and device level. Disconnect time is inversely proportional to cache effectiveness and is the key measurement.

Concentrate on the DASD LCUs and individual devices with the highest load or most critical data, and account for all I/Os being resolved either in CACHE or directly to DASD (this gives you the cache hit ratio). Then explode these two categories CACHE and DASD into their components, to gain further insights into how the cache controller is performing and what options may be available to manage the performance of the data.

Using the Spreadsheet Report, you can display these measurements by shift and by hour for further evaluation.

Using the guidelines and recommendations given in this chapter, analyze these measurements to determine overall DASD subsystem health and to identify areas of further opportunity. See *Cache Performance Management* for further details and examples.

Improving I/O performance with the Enterprise Storage Server

The Enterprise Storage Server (ESS) is IBM's most powerful disk storage server, developed using IBM Seascape[®] architecture. The ESS provides functions for the @server family of e-business servers, and also for the non-IBM (that is, Intel-based and UNIX-based) families of servers. Across all of these environments, the ESS features unique capabilities that allow it to meet the most demanding requirements of performance, capacity, and data availability that the computing business may require.

The Seascape architecture design has allowed the IBM TotalStorage Enterprise Storage Server to evolve from the initial E models to the succeeding F models, and to the 800 models, each featuring new, more powerful hardware and functional enhancements.

RMF has incorporated a number of changes to accommodate and report on the advanced ESS functions related to performance, for example, the number of parallel access volumes or information about RAID ranks in the cache subsystem.

z/OS Parallel Sysplex I/O management

In the z/OS Parallel Sysplex, the Workload Manager (WLM) controls where work is run and optimizes the throughput and performance of the total system. Until now, WLM management of the I/O has been limited. With ESS, there are some exciting new functions that allow WLM to control I/O across the sysplex. These functions include parallel access to both single system and shared volumes and the ability to prioritize the I/O based upon the WLM goals. The combination of these features can significantly improve performance in a wide variety of workload environments.

ESS performance features

This chapter covers the performance features of the ESS including PAV, multiple allegiance, and I/O priority queuing.

Concurrent access features

z/OS retains device queuing features that were developed to ensure effective serial access to physical devices. These features were developed at a time when physical devices rather than emulated devices were the norm and long before RAID solutions were developed. They ensured that only one channel program could be active to a disk at any time. This ensured that there was no possibility of interference between channel programs.

The traditional queuing method is illustrated in Figure 49 which shows that Application B is processing a request for volume 2000. If applications A and C issue I/O requests to the same volume, they will get a device busy or UCB busy response.

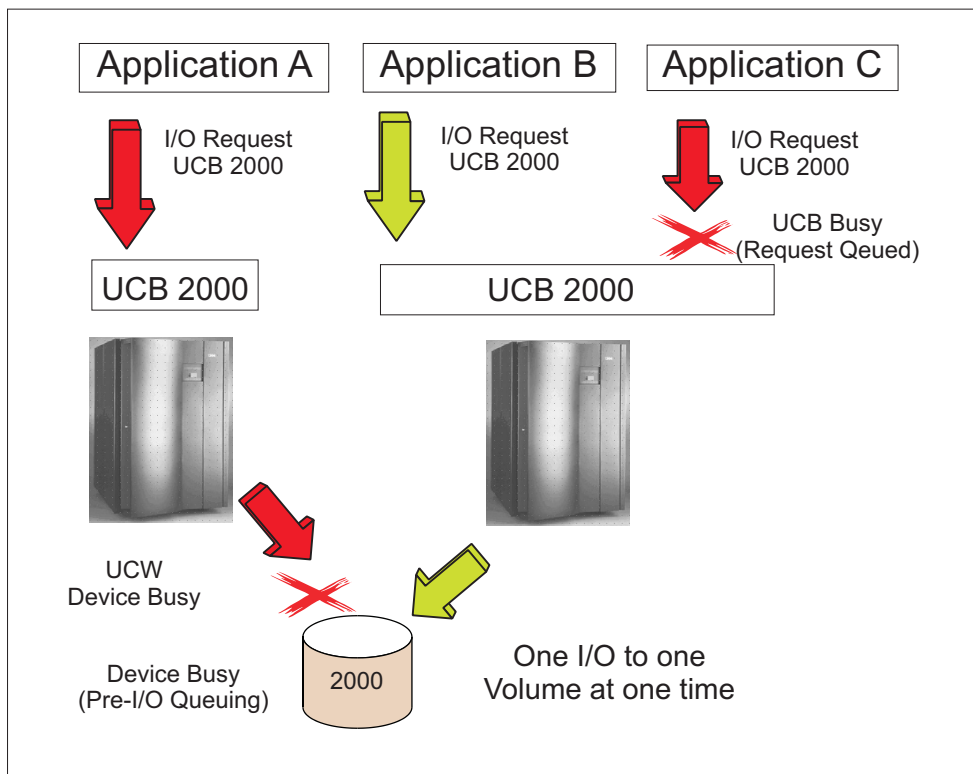


Figure 49. z/OS Traditional Device Serialization

z/OS interacting with the ESS relaxes some of these restrictions. Two features are introduced to allow this.

- Multiple allegiance
- Parallel Access Volumes (PAV)

Although there may be a large number of concurrent operations active to a particular logical volume, the ESS ensures that no I/O operations that have the potential to conflict over access to the same data will be scheduled together. Effectively no data can be accessed by one channel program that has the potential to be altered by another active program. Channel programs that are deemed to be incompatible with an active program are queued within the ESS.

Multiple allegiance

S/390 device architecture has defined that a state of implicit allegiance exists between a device and the channel path group that is accessing it. This allegiance is created in the control unit between the device and a channel path group when an I/O operation is accepted by the device. The allegiance causes the control unit to guarantee access, no busy status presented, to the device for the remainder of the channel program over the set of paths associated with the allegiance. This concept has been expanded to support the ESS with a concept of multiple allegiance.

ESS's concurrent operations capability supports concurrent accesses to or from the same volume from multiple channel path groups, system images. The ESS's multiple allegiance support allows different hosts to have concurrent implicit allegiances provided that there is no possibility that any of the channel programs can alter any data that another channel program might read or write.

Multiple allegiance requires no additional software or host support, other than to support the ESS. It is not externalized to the operating system or operator. Multiple allegiance will reduce contention reported as PEND time.

Resources that will benefit most from multiple allegiance are:

- Volumes that:
 - Have many concurrent read operations
 - Have a high read to write ratio
- Data sets that:
 - Have a high read to write ratio
 - Have multiple extents on one volume
 - Are concurrently shared by many users

Parallel access volumes (PAV)

z/OS systems queue I/O activity on a unit control block (UCB) that represents the physical device. High parallel I/O activity can adversely effect performance, traditionally because high accesses usually correlated to high levels of mechanical motion and in subsystems with large caches and RAID arrays because the volumes were treated as a single resource, serially reused. This contention is worst for large volumes with numerous small data sets. The symptom displayed is extended IOSQ time. The operating system can not attempt to start more than one I/O operation at a time to the device.

The ESS's concurrent operations capabilities also support concurrent data transfer operations to or from the same volume from the same system. A volume accessed in this way is called a Parallel Access Volume (PAV).

Figure 50 illustrates multiple allegiance and PAV allowing concurrent I/O processing.

PAV exploitation requires both software enablement and an optional feature on your ESS. PAV support must be installed on each ESS. It enables the issuing of multiple channel programs to a volume from a single system, and allows simultaneous access to the logical volume by multiple users or jobs. Reads, as well as writes to different extents, can be satisfied simultaneously. The domain of an I/O consists of the specified extents to which the I/O operation applies. Writes to the same domain still have to be serialized to maintain data integrity.

Support is implemented by defining multiple UCBs for volumes. The UCBs are of two types.

Base address

This is the actual unit address of the volume. There is only one base address for any volume.

Alias address

Alias addresses are mapped back to a base device address. I/O scheduled for an alias is physically performed against the base by the ESS. No physical disk space is associated with an alias address, however, they do occupy storage within z/OS. Alias UCBs are stored above the 16MB line.

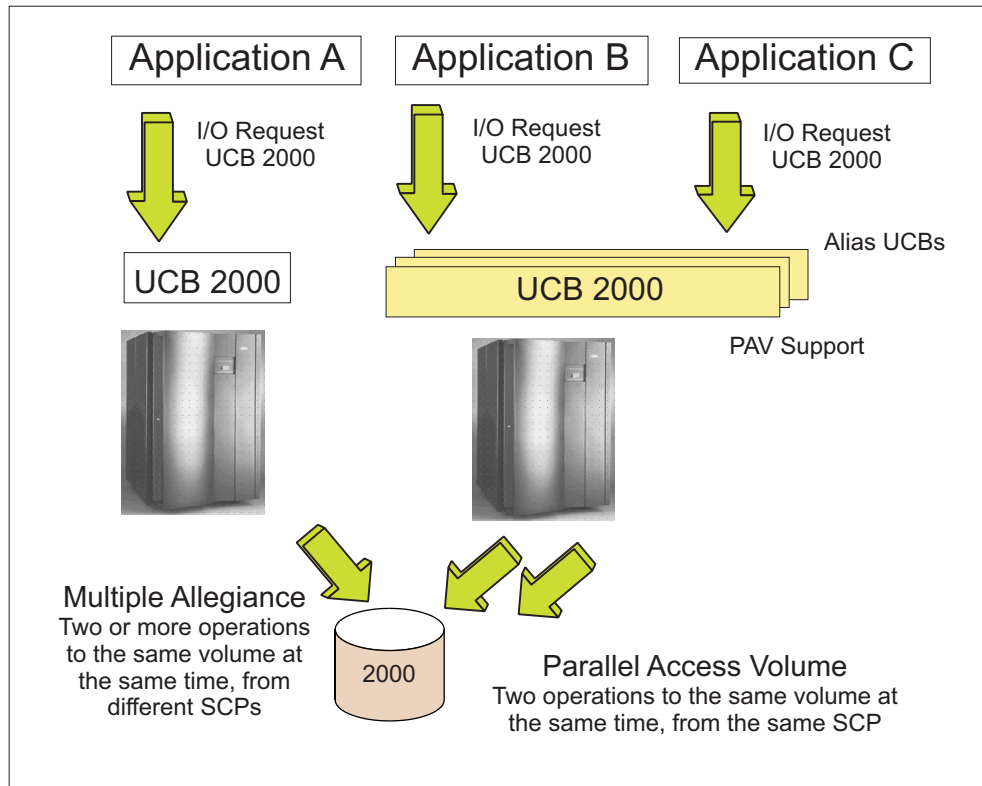


Figure 50. Device Queuing in a Parallel Access Volume Environment

The workloads that are most likely to benefit from the PAV function being available include:

- Volumes that have many concurrently open data sets, for example, volumes in a work pool
- Volumes that have high read to write ratio per extent
- Volumes reporting high IOSQ times

Candidate data types are:

- High read to write ratio
- Many extents on one volume
- Concurrently shared by many readers
- Accessed using media manager or allocated as VSAM extended format

PAVs can be assigned to base UCBs either manually or automatically by WLM. PAVs assigned manually are called static, while those movable by WLM are called dynamic. If WLM is used to manage I/O priorities, then you can use WLM dynamic PAV management as well. Otherwise, WLM can only use dynamic PAVs to balance device utilizations, not to directly help work achieve its goal.

I/O priority queuing

If I/Os cannot run in parallel, for example, due to extent conflicts, the ESS will internally queue I/Os. This reduces operating system overheads incurred by having to post "device busy" and redriving channel programs. The ESS will queue I/Os in the order in which they are received. This helps to reduce problems that occur when one processor can respond to interrupts faster than a sharing one, and so monopolize a device.

You also have the option to enable priority queuing of I/Os to the ESS. WLM sets a priority bit in the CCW. Priority queuing is within a sysplex and queuing is at a volume level. The ESS queues I/O requests in the order specified by WLM. I/O may be queued in the following situations:

- An extent conflict exists for a write operation.
- To allow servicing of a cache miss, device will reconnect when data has been staged to cache.
- A reserve request is issued and other accesses are current with a different path group ID.

Cache performance

Below are the CACHE and DASD components which you should measure. They are obtained from the Monitor I Cache Subsystem Activity report. In addition, you should collect the Monitor I Device Activity report data for the same period. You also need to know which devices are used by your important workloads: you can get this information from Monitor III DEVR or DEV reports.




Figure 51. Cache Performance Management: The Key

CACHE

RHIT	Normal and sequential read hits
DFWHIT	DASD fast write (DFW) hits
CFWHIT	Cache fast write (CFW) write and read hits

DASD

STAGE	Read misses and DFW misses. The record was not found in the cache. It is read or written, and that record with the rest of the track is staged into the cache. In the ESS, stages no longer occur
--------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Cache performance

when data for a DFW request is absent from cache. Instead, this is handled as a write "promote" (and counts as a hit).

DFWRETRY	DASD fast write retry. The record is in cache and is about to result in a cache hit, but non-volatile storage (NVS) is full. The operation is not really retried but written directly to DASD.
THRU	Write 'THRU' to DASD. These are writes to devices behind storage controls that are not enabled for DFW.
BYPASS	The bypass mode was specified in the define extent. The I/O is sent directly to DASD, the track is not promoted and the record is invalidated in cache. The ESS does not actually bypass the cache, but ensures that data which specifies bypass mode is destaged quickly.
INHIBIT	The inhibit mode was specified in the define extent and the record was not in the cache. Examples of inhibit include DFDSS for reads and DFSMS dynamic cache management (DCM or DCME) "maybe cache" and "never cache" storage classes. The I/O will retrieve the record in cache if it is there; if it is not in the cache, the I/O is sent directly to DASD, but no staging of the rest of the track occurs. If inhibit mode was specified and the result was a cache hit, the I/O would fall into one of the cache hit categories.
OFF	Device is turned off for caching.
ASYNC	Asynchronous destage of records to DASD. This includes anticipatory destage to write new or updated data from cache to DASD if NVS or cache is full. The unit is in tracks per second. All the other values given above are in I/Os per second. Async is not actually part of the I/O rate. It is the consequence of a write hit which must eventually be written to DASD. Counting it in the I/O rate would cause the I/O to be accounted for twice, once on the write hit and again when the track was destaged. Async can offer insight into the DFW.

It should be noted that the causes of bypass and inhibit listed above may not be all inclusive. In particular, you should consult your software vendor to understand which caching mode, if any, they use. The default mode is normal or sequential. However, bypass and inhibit mode can explicitly be set in software on an I/O basis. For example, it is possible for some I/Os to a data set to use inhibit and some normal caching mode, even within the execution of a particular application program. The same is true for DB2. Within DB2, one plan may access a table space in normal caching mode and another plan may optimize to access in prefetch mode, utilizing the bypass mode to accomplish this.

Cache management with ESS

The ESS provides performance improvements over those provided by raw disk by using caching. Caching algorithms are executed by the storage directors and determine what data occupies cache. Attached hosts can offer information about their data access intentions that the storage directors will use to help select the best algorithm to use for cache management. The following caching features are provided.

Read caching

Read hits occur when all of the data requested for a data access is located in cache. The ESS improves the performance of read caching by storing data in cache tracks that have the greatest probability of being accessed by a read operation.

An I/O operation that results in a read hit will not disconnect from the channel and data will be immediately available to the application. If the requested data is not located in cache, a write miss occurs and data is read from disk, returned to the program, and loaded to cache. This is called staging. While records are being read from disk, the channel program is disconnected allowing other applications to access the channel.

There are three types of staging:

Record staging

Only those records accessed by the channel program are staged into cache.

Partial track staging

The required records and the rest of the track are staged to cache. This is the default mode of cache operation.

Full track staging

Based on the prediction of sequential I/O processing. This can be predicted either by the previous behavior of the application or by the application signalling in the I/O request that sequential access will be used.

Data transferred using inhibit cache load or bypass cache attributes will be loaded into cache but eligible for accelerated destaging.

The ESS offers new capabilities to support the optimization of sequential performance; improved and more sensitive pre-fetching algorithms, and new channel commands that improve overheads and provide increased bandwidth.

Write caching

Two forms of write caching are supported: DASD Fast Write (DFW) and Cache Fast Write (CFW). Cache fast write is a z/OS function that is intended for data that does not need to be written to disk. CFW is used only when explicitly requested by the application and should only be selected for transient data, data that is not required beyond the end of the current job step and that can more easily be recreated from scratch than recovered from a checkpoint. An example of this type of data is intermediate work files, such as used by a sort program.

DFW is the more usual form of write caching. With DFW the application is told that an I/O operation is complete once data has successfully been written to cache and Non-Volatile Storage (NVS). Data integrity and availability is maintained by retaining two copies of the data until it is hardened to disk, one copy in cache on one cluster and the second in NVS of the other cluster. NVS is protected by battery backup. Normal access to the data is from the copy retained in cache.

Destaging of data that is backed up in NVS from cache to disk is based on a Least Recently Used (LRU) algorithm. Data may be retained in cache after being written to disk based on the cache activity. Destaging from NVS and cache is anticipatory and threshold based. The intention is to always have NVS and cache resources available to accept new data.

Tracks at the top of the cache LRU list are checked for updates in NVS that have not been destaged to disk. The ESS schedules tracks at the top of the NVS LRU for destaging, so that they can be allocated without a delay during destaging.

Write performance

Caching benefits write performance as almost all writes are at cache speeds. DFW minimizes any potential penalty of RAID 5 generation of parity. This performance benefit is clearly demonstrated by the published results of ESS performance tests.

Cache performance

In addition, write performance is enhanced by striping. The ESS automatically stripes logical volumes across all the drives in the RAID array. This provides automatic load balancing across the disk in the array, and an elimination of hot spots. This design should reduce the amount of effort that storage administrators spend, hand placing data, at the same time offering performance improvements.

The ESS RAID 5 implementation gives a minimal RAID 5 write penalty for sequential writes. When writing a sequential 'stripe' across the disks in an array, the ESS generates the parity only once for all the disks. This is sometimes called a RAID 3-like implementation, and it provides high performance in sequential operations.

Cache Subsystem Activity report

The report provides cache statistics on a subsystem basis as well as on a detailed device-level basis. You may start the analysis of your cache subsystem with the Summary report.

Subsystem summary

The report allows 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 V1R10				SYSTEM ID OS04						DATE 06/05/2008				INTERVAL 14.59.996						
				RPT VERSION V1R10 RMF						TIME 09.30.00										
SSID	CU-ID	TYPE	CACHE	NVS	I/O RATE	OFF RATE	--CACHE READ	HIT DFW	RATE-CFW	-----DASD STAGE	I/O RATE	-----DASD ICL	BY OTHER	ASYNC RATE	TOTAL H/R	READ H/R	WRITE H/R	% READ		
0600	0600	3990-006	256	64	3.1	0.0	1.0	2.2	0.0	0.0	0.0	0.0	0.0	0.0	1.000	1.000	1.000	30.8		
06C0	06E0	3990-006	512	32	38.6	0.0	16.5	22.0	0.0	0.0	0.0	0.0	0.0	1.7	0.999	0.997	1.000	42.9		
0A80	0ABE	3990-006	512	32	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		
6801	4000	2105-F20	8192	192	55.3	0.0	0.0	55.2	0.0	0.0	165.5	0.0	0.0	486.2	1.000	1.000	1.000	0.1		
6802	4100	2105-F20	8192	192	49.5	0.0	0.2	49.3	0.0	0.0	145.1	0.0	0.0	500.8	1.000	1.000	1.000	0.4		
6803	4200	2105-F20	8192	192	56.9	0.0	0.2	56.7	0.0	0.0	169.9	0.0	0.0	489.1	1.000	1.000	1.000	0.4		
6804	4300	2105-F20	8192	192	50.7	0.0	0.2	50.5	0.0	0.0	149.0	0.0	0.0	465.6	1.000	1.000	1.000	0.4		
6805	4400	2105-F20	8192	192	53.0	0.0	0.2	52.8	0.0	0.0	158.3	0.0	0.0	467.6	1.000	1.000	1.000	0.4		
6806	4500	2105-F20	8192	192	42.0	0.0	0.3	41.8	0.0	0.0	123.6	0.0	0.0	434.6	1.000	1.000	1.000	0.6		
6807	4600	2105-F20	8192	192	50.3	0.0	0.2	50.0	0.0	0.0	150.2	0.0	0.0	454.1	1.000	1.000	1.000	0.4		

Figure 52. Cache Subsystem Activity report - Summary

Top-20 device lists

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 V1R10			SYSTEM ID OS04			DATE 06/05/2008			INTERVAL 14.59.996								
			RPT VERSION V1R10 RMF			TIME 09.30.00											
*** DEVICE LIST BY DASD I/O RATE ***																	
VOLUME SERIAL	DEV NUM	SSID	% I/O	I/O RATE	---CACHE READ	HIT DFW	RATE-- CFW	-----DASD STAGE	I/O RATE-- DFWP	----- ICL	----- BYP	----- OTHER	ASYNC RATE	TOTAL H/R	READ H/R	WRITE H/R	% READ
PRD440	077E	00B1	25.7	53.1	21.3	16.7	0.0	13.9	0.0	1.3	0.0	0.0	1.6	0.714	0.604	1.000	67.9
PPDS14	0220	00CC	6.9	29.1	1.0	15.7	0.0	0.1	0.0	12.4	0.0	0.0	2.0	0.573	0.943	1.000	6.2
PRD437	0214	00CC	4.5	19.0	12.5	2.0	0.0	4.4	0.0	0.0	0.0	0.0	2.0	0.765	0.738	0.998	89.2
PPD026	0B76	00F1	27.4	71.4	65.9	1.1	0.0	4.3	0.0	0.0	0.0	0.0	0.5	0.939	0.939	0.975	98.4
PRD554	06B9	00FE	59.1	32.4	28.9	0.0	0.0	3.5	0.0	0.0	0.0	0.0	0.0	0.893	0.893	1.000	99.9
PRD339	0231	00CC	2.4	10.3	7.2	0.0	0.0	3.1	0.0	0.0	0.0	0.0	0.0	0.698	0.698	N/A	0.0
...																	
*** DEVICE LIST BY TOTAL I/O RATE ***																	
VOLUME SERIAL	DEV NUM	SSID	% I/O	I/O RATE	---CACHE READ	HIT DFW	RATE-- CFW	-----DASD STAGE	I/O RATE-- DFWP	----- ICL	----- BYP	----- OTHER	ASYNC RATE	TOTAL H/R	READ H/R	WRITE H/R	% READ
PPD026	0B76	00F1	27.4	71.4	65.9	1.1	0.0	4.3	0.0	0.0	0.0	0.0	0.5	0.939	0.939	0.975	98.4
PRD440	077E	00B1	25.7	53.1	21.3	16.7	0.0	13.9	0.0	1.3	0.0	0.0	1.6	0.714	0.604	1.000	67.9
PRD327	0200	00CC	11.8	49.8	3.0	46.8	0.0	0.1	0.0	0.0	0.0	0.0	6.3	0.998	0.973	1.000	6.2
PRD343	0515	00E4	17.5	48.9	24.8	24.0	0.0	0.1	0.0	0.0	0.0	0.0	5.5	0.998	0.996	0.999	50.8
PBV321	022C	00CC	11.2	47.3	47.1	0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.1	1.000	1.000	1.000	99.7
PRD307	0507	00E4	15.3	42.8	4.1	38.2	0.0	0.5	0.0	0.0	0.0	0.0	13.0	0.988	0.891	1.000	10.8
...																	

Figure 53. Cache Subsystem Activity Report - Top-20 Device Lists

Based on these reports, you can decide whether further investigation is necessary. The next step could be the creation of reports which show some more details. This can be done on subsystem level as well as on device level.

Subsystem-level reporting

This generates a report with three sections.

- Cache Subsystem Status
- Cache Subsystem Overview
- Cache Subsystem Device Overview

The subsystem-level report gives an overall view of the storage control, that is the amount of cache storage and non-volatile storage installed, as well as the current status of the cache. In addition, the performance analyst finds the number of I/O requests sent to the control unit and their resolution in the cache (*hits*). The report is completed by a list of all volumes attached to the subsystem, showing their specific utilization of the cache.

Device-level reporting

This generates, in addition to the above described report, 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 report is intended to analyze the cache usage in detail on the basis of the information about the applications that access these volumes.

Based on the data in the Summary report (Figure 52), one might decide that the subsystem SSID=00B1 needs further investigation.

Cache subsystem overview

C A C H E S U B S Y S T E M A C T I V I T Y													PAGE	1
z/OS V1R10		SYSTEM ID OS04			DATE 06/05/2008		INTERVAL 14.59.996							
		RPT VERSION V1R10 RMF			TIME 09.30.00									
SUBSYSTEM	3990-06	CU-ID	074C	SSID	00B1	CDATE	06/05/2008	CTIME	09.30.03	CINT	15.00			
TYPE-MODEL	9396-001													

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	1024.0M	CONFIGURED	32.0M	CACHING		- ACTIVE								
AVAILABLE	1019.9M	PINNED	0.0	NON-VOLATILE STORAGE		- ACTIVE								
PINNED	0.0			CACHE FAST WRITE		- ACTIVE								
OFFLINE	0.0			IML DEVICE AVAILABLE		- YES								

C A C H E S U B S Y S T E M O V E R V I E W														

TOTAL I/O	186363	CACHE I/O	181922	CACHE OFFLINE		0								
TOTAL H/R	0.865	CACHE H/R	0.886											
CACHE I/O REQUESTS	-----READ I/O REQUESTS-----		-----WRITE I/O REQUESTS-----										%	
	COUNT	RATE	HITS	RATE	H/R	COUNT	RATE	FAST	RATE	HITS	RATE	H/R	READ	
NORMAL	116003	128.9	95313	105.9	0.822	15150	16.8	15150	16.8	15134	16.8	0.999	88.4	
SEQUENTIAL	16655	18.5	16617	18.5	0.998	34114	37.9	34114	37.9	34114	37.9	1.000	32.8	
CFW DATA	0	0.0	0	0.0	N/A	0	0.0	0	0.0	0	0.0	N/A	N/A	
TOTAL	132658	147.4	111930	124.4	0.844	49264	54.7	49264	54.7	49248	54.7	1.000	72.9	
REQUESTS	-----CACHE MISSES-----		-----MISC-----		-----NON-CACHE I/O-----									
	READ	RATE	WRITE	RATE	TRACKS	RATE	COUNT		RATE	COUNT		RATE		
NORMAL	20690	23.0	16	0.0	12993	14.4	DFW BYPASS		289	0.3	ICL 4435 4.9			
SEQUENTIAL	38	0.0	0	0.0	870	1.0	CFW BYPASS		0	0.0	BYPASS 6 0.0			
CFW DATA	0	0.0	0	0.0			DFW INHIBIT		0	0.0	TOTAL 4441 4.9			
							ASYNC (TRKS)		7239	8.0				
TOTAL	20744	RATE	23.0											
---CKD STATISTICS---														
---RECORD CACHING---														
WRITE	268	READ MISSES	14502											
WRITE HITS	268	WRITE PROM	12938											

Figure 54. Cache Subsystem Activity Report - Status and Overview

This section provides, at a glance, all relevant data for one subsystem. It distinguishes between I/O requests to be handled by the cache (*cachable I/Os*) and non-cache I/O requests.

Cachable I/O Requests

These requests are shown in three categories:

NORMAL

Cache will be managed by *least-recently-used (LRU)* algorithm for making cache space available.

SEQUENTIAL

Tracks following the track assigned in the current CCW chain are promoted, they will be transferred from DASD to cache in anticipation of a near-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.

Non-Cache I/O Requests

These requests that switched off cache processing, are shown in two categories:

ICL Inhibit cache load — number of I/O requests that inhibited load of data into cache, and the data was not found in the cache.

Note: If the data was in the cache, it has been counted as cache hit. Therefore, this is actually the number of ICL misses.

BYPASS

Number of I/O requests that explicitly bypassed the cache, no matter whether the data is in the cache or not.

I/O Requests

You find the following numbers in the sample report:

Cachable I/O Requests

TOTAL READS	132658			
READ HITS	111930	→	READ HIT RATE (H/R)	0.844
TOTAL WRITES	49264			
WRITE HITS	49248	→	WRITE HIT RATE (H/R)	1.000
TOTAL CACHE I/O	181922			
CACHE HITS	161178	→	CACHE HIT RATE (H/R)	0.886

Non-Cache I/O Requests

ICL	4435
BYPASS	6
TOTAL	4441

Total I/O Requests

The sum of these two groups is shown as total number of I/O requests for cached devices of this subsystem:

CACHE I/O	181922			
NON-CACHE I/O	4441			
TOTAL I/O	186363			
CACHE HITS	161178	→	TOTAL HIT RATE (H/R)	0.865

While hit percentage by itself (CACHE H/R) does not drive the cache performance management process, hit percentage should be derived by dividing the hits by all I/Os (TOTAL H/R).

Hit percents derived by dividing the hits by cachable I/Os are misleading. Cachable I/Os refers to those I/Os that are eligible for caching. At first glance, the distinction between cachable I/Os and all I/Os seems rather insignificant since most I/Os are eligible for caching. However, some software (e.g. DFSMS Dynamic Cache Management and DB2) will selectively disable I/O operations for caching, using inhibit and bypass modes of operations. This could result in high cachable hit percentages but lower hit percentages. In this case the disconnect time would be higher than what you would expect for a high cachable hit percent, but the

Cache performance

disconnect time would be reasonable for a lower hit percent. Where all I/Os are eligible for caching, the hit percent and cachable hit percent will be equal.

You can also get from the report the relationship between READ and WRITE requests as %READ. Sometimes, performance experts work with the READ/WRITE ratio — this value is not shown in the report because of arithmetic problems in case of very less (or zero) WRITE requests, but can be calculated easily, otherwise.

READ/WRITE Ratio

TOTAL CACHE I/O	181922			
TOTAL CACHE READS	132658	→	% READ	72.9
TOTAL CACHE WRITES	49264	→	R/W ratio	2.69

The R/W ratio can also be calculated as

$$\text{R/W ratio} = 72.9 / (100 - 72.9) = 2.69$$

Cache subsystem device overview

C A C H E S U B S Y S T E M A C T I V I T Y														PAGE 2			
z/OS V1R10		SYSTEM ID OS04				DATE 06/05/2008		INTERVAL 14.59.996									
		RPT VERSION V1R10 RMF				TIME 09.30.00											
SUBSYSTEM	2105-01	CU-ID	4000	SSID	6801	CDATE	06/05/2008	CTIME	09.30.03	CINT	15.00						
TYPE-MODEL	2105-F20																

C A C H E S U B S Y S T E M D E V I C E O V E R V I E W																	

VOLUME SERIAL	DEV NUM	RRID	% I/O	I/O RATE	---CACHE READ	HIT DFV	RATE-- CFV	-----DASD STAGE	I/O DFWP	RATE-- ICL	----- BYP	OTHER	ASYNC RATE	TOTAL H/R	READ H/R	WRITE H/R	% READ
*ALL			100.0	55.3	0.0	55.2	0.0	0.0	165.5	0.0	0.0	0.0	486.2	1.000	1.000	1.000	0.1
*CACHE-OFF			0.0	0.0													
*CACHE			100.0	55.3	0.0	55.2	0.0	0.0	165.5	0.0	0.0	0.0	486.2	1.000	1.000	1.000	0.1
SR4B00	4000	0000	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	N/A	N/A	N/A	N/A
SR4B01	4001	0000	17.5	9.7	0.0	9.7	0.0	0.0	29.0	0.0	0.0	0.0	39.1	1.000	N/A	1.000	0.0
SR4B02	4002	0000	17.9	9.9	0.0	9.9	0.0	0.0	29.7	0.0	0.0	0.0	36.3	1.000	N/A	1.000	0.0
SR4B03	4003	0000	10.3	5.7	0.0	5.6	0.0	0.0	16.9	0.0	0.0	0.0	16.8	1.000	1.000	1.000	0.9
SR4B04	4004	0000	17.7	9.8	0.0	9.8	0.0	0.0	29.4	0.0	0.0	0.0	39.2	1.000	N/A	1.000	0.0
SR4B05	4005	0000	18.1	10.0	0.0	10.0	0.0	0.0	30.0	0.0	0.0	0.0	33.8	1.000	N/A	1.000	0.0
SR4B06	4006	0000	18.4	10.2	0.0	10.2	0.0	0.0	30.5	0.0	0.0	0.0	34.6	1.000	N/A	1.000	0.0

Figure 55. Cache Subsystem Activity Report - Device Overview

This section of the report provides an overview for one control unit with all attached devices. You can use this report for investigating the cache effectiveness of each volume.

During this analysis, you should use the values I/O RATE and % I/O as indicators to decide whether a volume has an I/O activity that makes it worth for further performance investigation.

If you come to the conclusion that further analysis of a specific volume might be necessary, you should request the report on device level. It has the same structure as the Subsystem Overview report, but provides the details for each volume.

RAID Rank Activity

C A C H E S U B S Y S T E M A C T I V I T Y														PAGE 3			
z/OS V1R10		SYSTEM ID OS04				DATE 06/05/2008		INTERVAL 14.59.996									
		RPT VERSION V1R10 RMF				TIME 09.30.00											
SUBSYSTEM	2105-01	CU-ID	4000	SSID	6801	CDATE	06/05/2008	CTIME	09.30.03	CINT	15.00						
TYPE-MODEL	2105-F20																

R A I D R A N K A C T I V I T Y																	

ID	RAID TYPE	DA	HDD	----- READ REQ -----				----- WRITE REQ -----				----- HIGHEST UTILIZED VOLUMES -----					
				RATE	AVG MB	MB/S	RTIME	RATE	AVG MB	MB/S	RTIME						
*ALL			14	265	0.055	14.6	150	243	0.111	27.0	249						
0000	RAID-5	01	7	257	0.055	14.2	150	125	0.089	11.1	266	SR4B06	SR4B05	SR4B02	SR4B04	SR4B01	SR4B03
0001	RAID-5	01	7	7	0.055	0.4	148	118	0.134	15.9	232						

Figure 56. Cache Subsystem Activity Report - RAID Rank Activity

This section of the report provides an overview about the activities of all RAID ranks belonging to the storage server. For each rank, you get the list of those six logical volumes which have the highest utilization, this might help you in deciding whether the mapping between RAID ranks and logical volumes is well balanced, or whether some changes should be performed.

Cache performance

Cache Device Activity

C A C H E D E V I C E A C T I V I T Y														PAGE	2
z/OS V1R10		SYSTEM ID		SYS1	DATE 11/28/2008		INTERVAL 15.00.055								
		RPT VERSION		V1R10	TIME 00.30.00										
SUBSYSTEM	2105-01	CU-ID	1E91	SSID	3007	CDATE	11/28/2008	CTIME	00.30.00	CINT	15.00				
TYPE-MODEL	2105-E20	MANUF	IBM	PLANT	75	SERIAL	000000016374								
VOLSER	NP1MHI	NUM	1E85	RRID	0700										

C A C H E D E V I C E S T A T U S															

CACHE STATUS				DUPLEX PAIR STATUS											
CACHING	- ACTIVE			DUPLEX PAIR	- NOT ESTABLISHED										
DASD FAST WRITE	- ACTIVE			STATUS	- N/A										
PINNED DATA	- NONE			DUAL COPY VOLUME	- N/A										

C A C H E D E V I C E A C T I V I T Y															

TOTAL I/O	32141	CACHE I/O	32141	CACHE OFFLINE	N/A										
TOTAL H/R	0.785	CACHE H/R	0.785												
CACHE I/O	-----READ I/O REQUESTS-----			-----WRITE I/O REQUESTS-----										%	
REQUESTS	COUNT	RATE	HITS	RATE	H/R	COUNT	RATE	FAST	RATE	HITS	RATE	H/R	READ		
NORMAL	26728	29.6	19998	22.1	0.748	0	0.0	0	0.0	0	0.0	N/A	100.0		
SEQUENTIAL	5413	6.0	5241	5.8	0.968	0	0.0	0	0.0	0	0.0	N/A	100.0		
CFW DATA	0	0.0	0	0.0	N/A	0	0.0	0	0.0	0	0.0	N/A	N/A		
TOTAL	32141	35.6	25239	27.9	0.785	0	0.0	0	0.0	0	0.0	N/A	100.0		

-----CACHE MISSES-----						-----MISC-----				-----NON-CACHE I/O-----					
REQUESTS	READ	RATE	WRITE	RATE	TRACKS	RATE	COUNT	RATE	DFW BYPASS	0	0.0	ICL	0	0.0	
NORMAL	6730	7.4	0	0.0	6899	7.6	CFW BYPASS	0	0.0	BYPASS	0	0.0	0	0.0	
SEQUENTIAL	172	0.2	0	0.0	6026	6.7	DFW INHIBIT	0	0.0	TOTAL	0	0.0	0	0.0	
CFW DATA	0	0.0	0	0.0			ASYNC (TRKS)	0	0.0						
TOTAL	6902	RATE	7.6												

---CKD STATISTICS---			---RECORD CACHING---			---HOST ADAPTER ACTIVITY---				-----DISK ACTIVITY-----					
WRITE	0	READ MISSES	0			BYTES	BYTES			RESP	BYTES	BYTES			
WRITE HITS	0	WRITE PROM	28			/REQ	/SEC			TIME	/REQ	/SEC			
						READ	736	146			READ	14.132	63.1K	743.9K	
						WRITE	15.9K	471.6K			WRITE	13.812	50.9K	648.1K	

Figure 57. Cache Subsystem Activity Report - Cache Device Activity (device-level reporting)

This report shows the statistics for volume NP1MHI. For example, you might look at this data if users are complaining about the response time of a specific application, and Monitor III reports flag a certain volume as a major cause of delay.

If you prefer to see data for more than one interval, you can create an Overview report with the selection of those data you are interested in. Let us assume that you want to get the following data for volume PRD440 (with SSID=00B1 and address 077E):

Exception	Meaning
CADRT	Total READ I/O rate
CADWT	Total WRITE I/O rate
CADRHRT	Total READ hit ratio
CADWHRT	Total WRITE hit ratio
CADMT	Total staging rate
CADDFWB	Total DFW bypass rate
CADNCICL	Non-cache ICL rate

You find a listing of all conditions in the *z/OS RMF User's Guide*.

You get the Overview report by specifying the following Postprocessor options:

```
OVERVIEW(REPORT)
OVW(READRATE(CADRT(SSID(00B1),DEVN(077E))))
OVW(WRTRATE(CADWT(SSID(00B1),DEVN(077E))))
OVW(READHR(CADRHRT(SSID(00B1),DEVN(077E))))
OVW(WRTHR(CADWHRT(SSID(00B1),DEVN(077E))))
OVW(STGRT(CADMT(SSID(00B1),DEVN(077E))))
OVW(DFWBPRT(CADDFWB(SSID(00B1),DEVN(077E))))
OVW(ICLRT(CADNCICL(SSID(00B1),DEVN(077E))))
```

This creates the following report:

R M F O V E R V I E W R E P O R T									
z/OS V1R10			SYSTEM ID OS04		DATE 06/05/2008		INTERVAL 14.59.996		
			RPT VERSION V1R10 RMF		TIME 09.30.00		CYCLE 1.000 SECONDS		
NUMBER OF INTERVALS 5			TOTAL LENGTH OF INTERVALS 01.14.57						
DATE	TIME	INT	READRATE	WRTRATE	READHR	WRTHR	STGRT	DFWBPRT	ICLRT
MM/DD	HH.MM.SS	HH.MM.SS							
06/05	10.44.00	01.15.00	35.213	16.671	0.604	1.000	13.957	0.030	1.263
06/05	10.59.00	00.14.59	2.799	12.444	0.514	1.000	1.359	0.089	0.021
06/05	11.14.00	00.15.00	2.769	0.159	0.318	1.000	1.889	0.000	0.063
06/05	11.29.00	00.14.59	3.692	0.225	0.377	1.000	2.301	0.001	0.034
06/05	11.44.00	01.14.59	0.535	0.261	0.696	0.996	0.164	0.003	0.022

Figure 58. Overview Report for Cached DASD Device

In this report, you see that the volume was very active just in one interval, this means that further investigation is required to highlight the critical volumes of this Cache subsystem.

Using the Spreadsheet Reporter, you can create reports with key performance data you are interested in.

If you would like to see all SSIDs at glance, you can create the Cache Subsystem Report which displays all important cache characteristics. To get all details about cache rates for a longer period of time, you would use the Cache Trend Report.

Cache performance

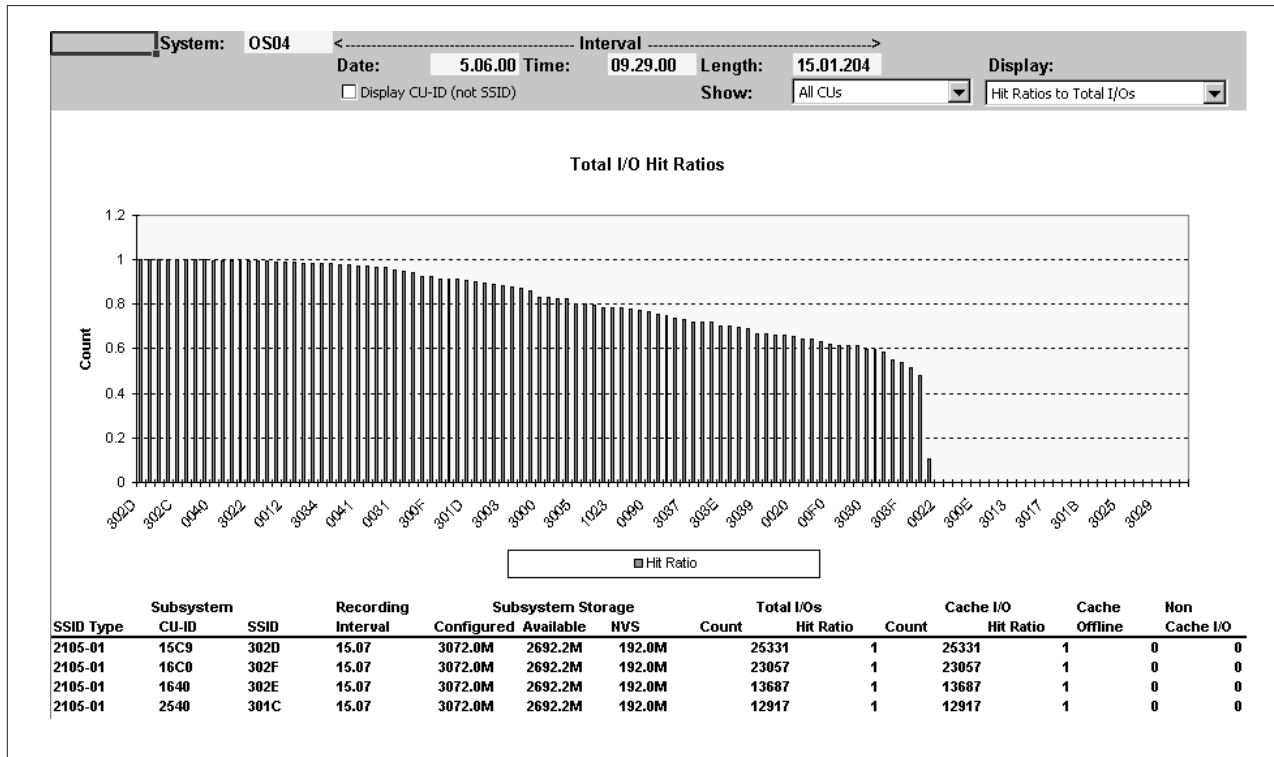


Figure 59. Cache Hits Overview Report. Spreadsheet Reporter macro RMFR9CAC.XLS (Cache Subsystem Report - SSID0vw)

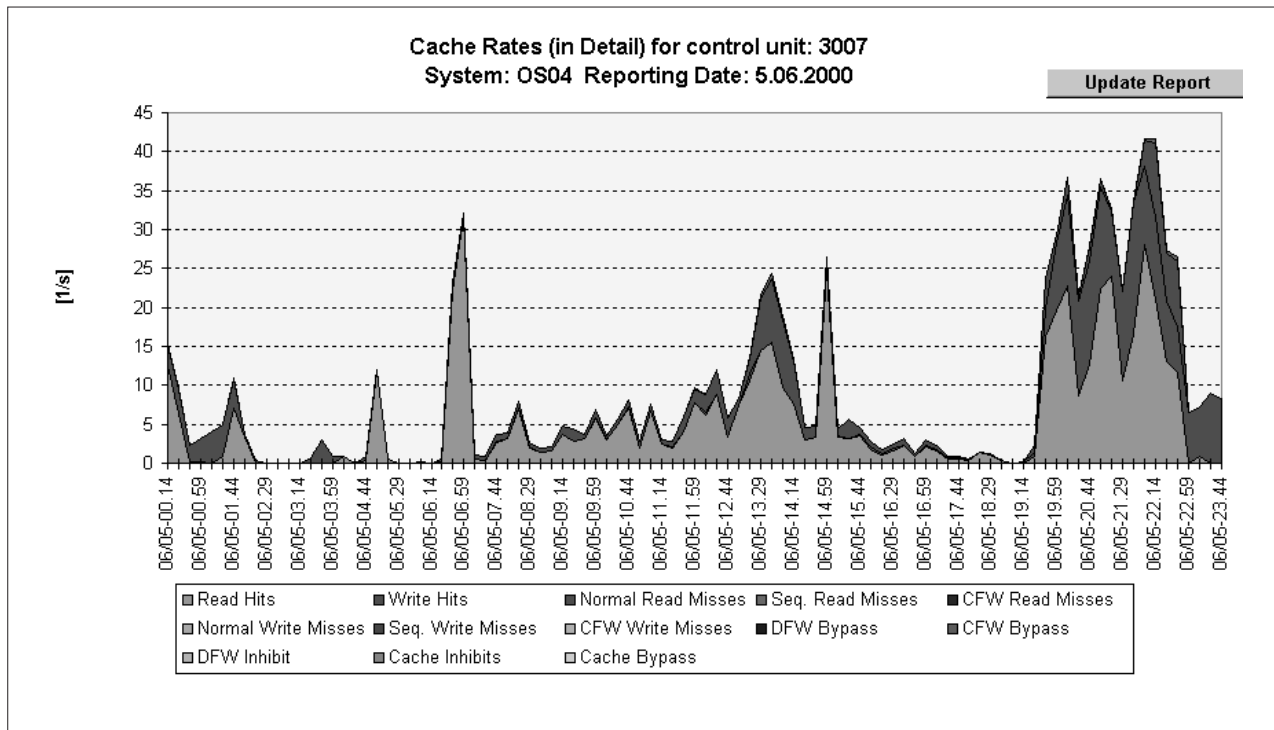


Figure 60. Cache Trend Report. Spreadsheet Reporter macro RMFX9CAC.XLS (Cache Trend Report - Summary)

Notes:

1. Cache controller information is collected by individual device addresses without segregation by issuing system. Therefore, the Monitor I Data Gatherer automatically collects the total for each device behind the storage controller. There is no way to differentiate the I/Os by each sharing system. Thus, data gathering is only required to be activated on one of the sharing systems.
2. The report uses cache storage subsystem ID (SSID) numbers to identify control units. The Monitor I DASD Activity report uses logical control unit (LCU) numbers to report on these same control units. LCU and SSID often will not be the same value, even when referring to the same control unit. You may need to compare the address range of devices on a given controller to match LCU to SSID.
3. Also, I/O rates shown in the Cache Subsystem Activity report may not exactly match I/O rates given in Shared Device Activity report (the sum of device activity from all sharing systems). This is because the storage controller counts each locate in a multiple locate CCW chain, whereas RMF only counts one I/O to start the CCW. An example of a CCW with multiple locates is DB2 list prefetch. Conversely, control unit commands, such as standalone RESERVEs, would get counted for the DASD Activity report but not for the Cache Subsystem Activity report.

Indicator

Field: Use the following fields:

- TOTAL H/R - CACHE H/R in the Subsystem Overview report and Device Activity report.
- TOTAL H/R - READ H/R - WRITE H/R in the Device Overview report

Guideline: Check whether all volumes (pay special attention to those that do a lot of I/O) are experiencing a good hit ratio (80% or higher). If this is the case, further analysis is not necessary. Note that the occurrence of a low hit ratio for one or more devices cannot be interpreted as a sure sign of the need for more cache (other analytical tools, such as the IBM Cache Analysis Aid, will be required for more thorough cache sizing). The low hit ratio may be caused by data sets which do not cache well, bypass cache, etc.

Problem Area: Low cache hit ratios will mean higher DISC times for that device. Also, devices with low cache hit ratios may use cache storage that other devices could have used more profitably.

Potential Solution: If you have DFSMS/MVS[®] and its *Dynamic Cache Management Extended (DCME)* capability, DCME will dynamically manage cache for those data sets you specify (via storage class). Thus no manual cache tuning should be required on your part.

If your data is not SMS-managed, you may wish to identify volumes which are not achieving a good cache hit ratio, and are doing a significant amount of staging, and turn them off for caching. This will improve the use of cache for volumes that benefit more from cache. A trial and error approach will be needed here; you will need to make a judgment call on whether to cache devices that are benefitting marginally from cache (maybe 30-80 % hit ratios).

Using Monitor III cache reporting

Monitor III offers two reports which can assist you in monitoring online the performance of your Cache subsystem:

CACHSUM Cache Summary report

CACHDET Cache Detail report

Both reports provide similar information as you can find in the Postprocessor Cache reports. Therefore, the performance indicators are the same, and you might refer to the description of the Postprocessor reports for further evaluation.

Cache Summary report

The report 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.

```

RMF V1R10 Cache Summary - SYSPLEX Line 1 of 21
Command ==> _ Scroll ==> HALF
Samples: 120 Systems: 5 Date: 04/15/08 Time: 08.25.00 Range: 120 Sec
CDate: 04/15/08 CTime: 08.24.55 CRange: 120 Sec

```

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

Figure 61. CACHSUM Report

Cache Detail report

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

RMF V1R10 Cache Detail - SYSPLEX													
Line 1 of 20													
Command ==> _													
Scroll ==> HALF													
Samples: 120		Systems: 5		Date: 11/28/08		Time: 16.11.30		Range: 120		Sec			
				CDate: 11/28/08		CTime: 16.11.25		CRange: 119		Sec			
Volume	/Num	SSID	I/O	I/O	Hit	Cache	Hit	Rate	-	DASD	I/O	Seq	Async
			%	Rate	%	Read	DFW	CFW	Total	Stage	Rate	Rate	Rate
*ALL			100	42.9	99.5	30.8	11.9	0.0	0.2	0.2	0.6	1.8	
*NOCAC			0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
*CACHE			100	42.9	99.5	30.8	11.9	0.0	0.2	0.2	0.6	1.8	
SYSPXC	05CB	0043	27.9	12.0	100	10.4	1.6	0.0	0.0	0.0	0.0	0.0	
SYSSM3	05C3	0043	11.6	5.0	100	3.9	1.1	0.0	0.0	0.0	0.3	0.4	
SYSAXC	05CA	0043	11.2	4.8	100	3.2	1.6	0.0	0.0	0.0	0.0	0.0	
SYSSM6	05C6	0043	10.8	4.6	98.9	3.7	0.9	0.0	0.1	0.1	0.0	0.3	
SYSSMS	05C0	0043	9.7	4.2	97.6	0.4	3.7	0.0	0.1	0.1	0.0	0.0	
SYSSM8	05C8	0043	6.8	2.9	100	1.6	1.4	0.0	0.0	0.0	0.1	0.1	
SYSSM5	05C5	0043	5.4	2.3	100	2.3	0.1	0.0	0.0	0.0	0.0	0.2	
SYSSM2	05C2	0043	3.7	1.6	100	0.6	1.0	0.0	0.0	0.0	0.0	0.5	
SYSSMB	05DA	0043	3.6	1.6	100	1.4	0.1	0.0	0.0	0.0	0.1	0.0	
SYSSMA	05CD	0043	3.6	1.5	99.5	1.5	0.0	0.0	0.0	0.0	0.1	0.0	
SYSSM7	05C7	0043	3.3	1.4	98.2	1.2	0.2	0.0	0.0	0.0	0.0	0.1	
SYSOPE	05CC	0043	1.6	0.7	97.5	0.6	0.1	0.0	0.0	0.0	0.0	0.0	
SYSSM4	05C4	0043	0.6	0.3	96.9	0.0	0.2	0.0	0.0	0.0	0.0	0.0	
SYSSM9	05C9	0043	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
SYSSM1	05C1	0043	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
SYSSMC	05DB	0043	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	
MVSWR1	05FF	0043	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	

Figure 62. CACHDET Report

DASD indicators and guidelines

The following sections discuss performance topics and tuning guidelines that can help you to resolve DASD I/O throughput issues.

Response time components

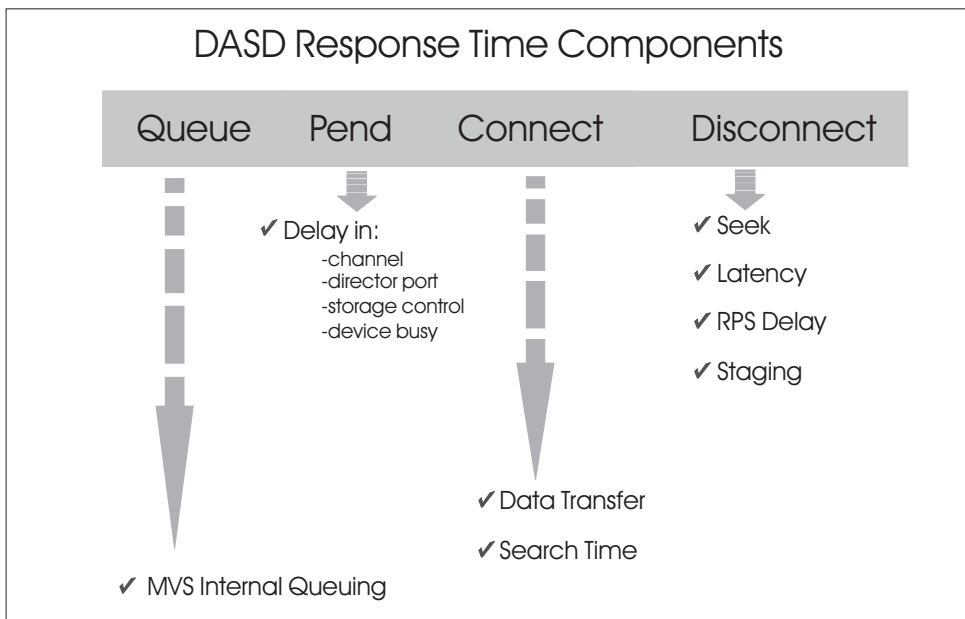


Figure 63. DASD Response Time Components

Response time is a key measure of how well your I/O is performing, and a key indicator of what can be done to speed up the I/O. It is important to understand the components of this response time, and what they mean. DASD response time is the sum of queue, pend, connect and disconnect time. The following lists the response time components and what causes them:

Connect The part of the I/O during which data is actually transferred, protocol, search and data transfer time.

Disconnect Time that an I/O request spends freed from the channel. This is the time that the I/O positions for the data that has been requested. It includes

SEEK and SET SECTOR

Moving the device to the requested cylinder and track

Latency

Waiting for the record to rotate under the head

Rotational Position Sensing (RPS)

Rotational delay, as the device waits to reconnect to the channel. The term "RPS delay" traditionally means waiting for an extra disk rotation because a transfer could not occur when the data was under the head. Delays of this type are now obsolete with the ESS.

Pend Time that the I/O is delayed in the path to the device. Pend time

may be attributable to the channel, control unit, or wait for director (director port delay time), although it is often caused by shared DASD.

IOS Queue Represents the average time that an I/O waits because the device is already in use by another task on this system, signified by the device's UCBBUSY bit being on.

For most customers, reporting is based on one or both of the following measurements:

- Service Time** Connect time plus disconnect time
- Response Time** Connect time plus disconnect time plus pend time plus IOSQ time

Spreadsheet reports

In Chapter 2, "Diagnosing a problem: The first steps," on page 23, you could see the DASD Summary report (Figure 30 on page 51) to be used for first steps in understanding the DASD performance of your system. If you feel that further investigation is required, you can create additional reports.

You might start with a second Summary report that provides data on the most heavily used LCU and DASD devices:

RMF DASD System Summary													
System Id: UIG1		Operating System: z/OS V1R9 -		Report Range(seconds): 3600								Sort Lists	
Reporting Date: 02.10.2007		Reporting Time: 09.05.11		Report Range(hh.mm.ss): 001.00.00									
System Summary													
#of LCU's	# of DASD I/O	Intens.	ST Intens.	Path Int.	Act. Rt.	Resp. Tm	Serv. Tm	IOSQ Tm	Pend. Tm	Disc. Tm	Conn. Tm	RT/ST	
29	1669	5875,37	5536,31	3597,10	1270,41	4,62	4,36	0,03	0,24	1,53	2,83	1,06	
LCU Summary													
	LCU	I/O Intens.	ST Intens.	Path Int.	Act. Rt.	Resp. Tm	Serv. Tm	IOSQ Tm	Pend. Tm	Disc. Tm	Conn. Tm	RT/ST	
Top 5	00CA	1075,06	1023,86	435,14	255,97	4,30	4,00	0,00	0,30	2,30	1,70	1,08	
	00CD	798,51	759,24	602,16	130,90	6,10	5,80	0,10	0,20	1,20	4,60	1,05	
	00C2	533,76	467,04	346,94	133,44	3,80	3,50	0,10	0,20	0,90	2,60	1,09	
	00C9	406,71	395,09	197,54	116,20	3,60	3,40	0,00	0,20	1,70	1,70	1,06	
	00C7	388,47	377,53	153,20	54,71	7,20	6,90	0,00	0,30	4,10	2,80	1,04	
sorted by I/O Intensity													
Device Summary Top 10													
LCU	VolSer	I/O Intens.	ST Intens.	Path Int.	Act. Rt.	Resp. Tm	Serv. Tm	IOSQ Tm	Pend. Tm	Disc. Tm	Conn. Tm	RT/ST	
00CA	UDB921	353,72	329,61	144,71	80,39	4,40	4,10	0,00	0,30	2,30	1,80	1,07	
00CA	UDB925	274,59	261,81	102,17	63,86	4,40	4,10	0,00	0,30	2,50	1,60	1,07	
00C8	SYSUC1	230,66	199,21	94,36	104,85	2,30	1,90	0,10	0,30	1,00	0,90	1,21	
00CD	UDB334	159,58	155,92	154,08	18,34	8,70	8,50	0,00	0,20	0,10	8,40	1,02	
00C2	UDB905	150,12	128,68	107,23	71,49	2,00	1,80	0,00	0,20	0,30	1,50	1,11	
00CD	UDB337	146,12	135,12	133,55	15,71	9,40	8,60	0,60	0,20	0,10	8,50	1,09	
00CA	UDB950	129,18	122,72	51,67	32,30	4,00	3,80	0,00	0,20	2,20	1,60	1,05	
00C4	UIDBSR	125,91	116,22	82,33	48,43	2,60	2,40	0,00	0,20	0,70	1,70	1,08	
00CD	UDB333	95,26	94,18	93,10	10,83	8,90	8,70	0,00	0,20	0,10	8,60	1,02	
00C2	UIRG14	94,00	92,26	88,78	8,70	10,80	10,60	0,00	0,20	0,40	10,20	1,02	
sorted by I/O Intensity													
Averages													
Devices	Top5	I/O Intens.	ST Intens.	Path Int.	Act. Rt.	Resp. Tm	Serv. Tm	IOSQ Tm	Pend. Tm	Disc. Tm	Conn. Tm	RT/ST	
	Top5	233,74	215,04	120,51	67,79	3,45	3,17	0,02	0,26	1,39	1,78	1,09	
	Top10	175,92	163,57	105,20	45,49	3,87	3,60	0,06	0,21	1,28	2,31	1,08	

Figure 64. DASD Summary Report. Spreadsheet Reporter macro RMFR9DAS.XLS (DASD Activity Report - Summary)

DASD performance

If you have seen some average data for most active volumes, you might be interested in getting some more details:

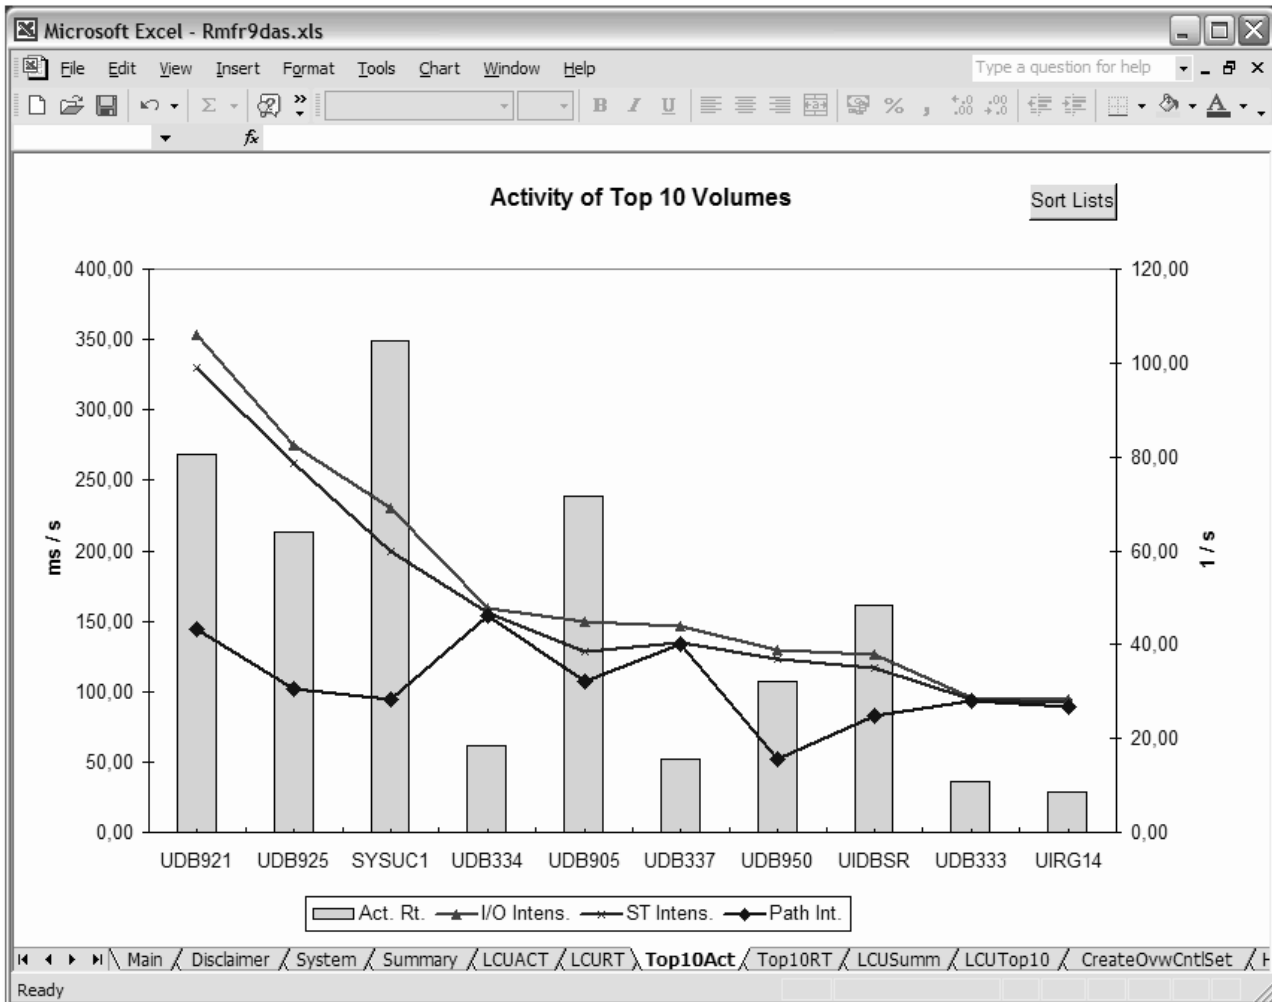


Figure 65. Activity of Top-10 Volumes. Spreadsheet Reporter macro RMFR9DAS.XLS (DASD Activity Report - Top10Act)

The report shows four values: in addition to the activity rate, you get the values for the I/O Intensity, the Service Time Intensity, and the Path Intensity.

I/O Intensity is the product of response time multiplied with the activity rate. The unit of measurement is ms/s.

It allows you to examine how many milliseconds per second applications waited for the device. The value can exceed 1000, which is an indicator that the device experiences a very heavy load and requires a long time to satisfy the requests. The value also includes the IOS Queue Time. If the activity rate is very low (e.g. less than 1 SIO/s) and the IOS Queue Time is high, the value is not very meaningful. You should use this measurement for devices with a considerable load (more than 2 IOs/s).

I/O Intensity is not a common name in the literature. Other references may use the same measurement with a different name, for example

Response Time Volume
DASD MPL (typically divided by 1000)

If you are interested in getting some details about the response times of the top-10 volumes, you can get this report:

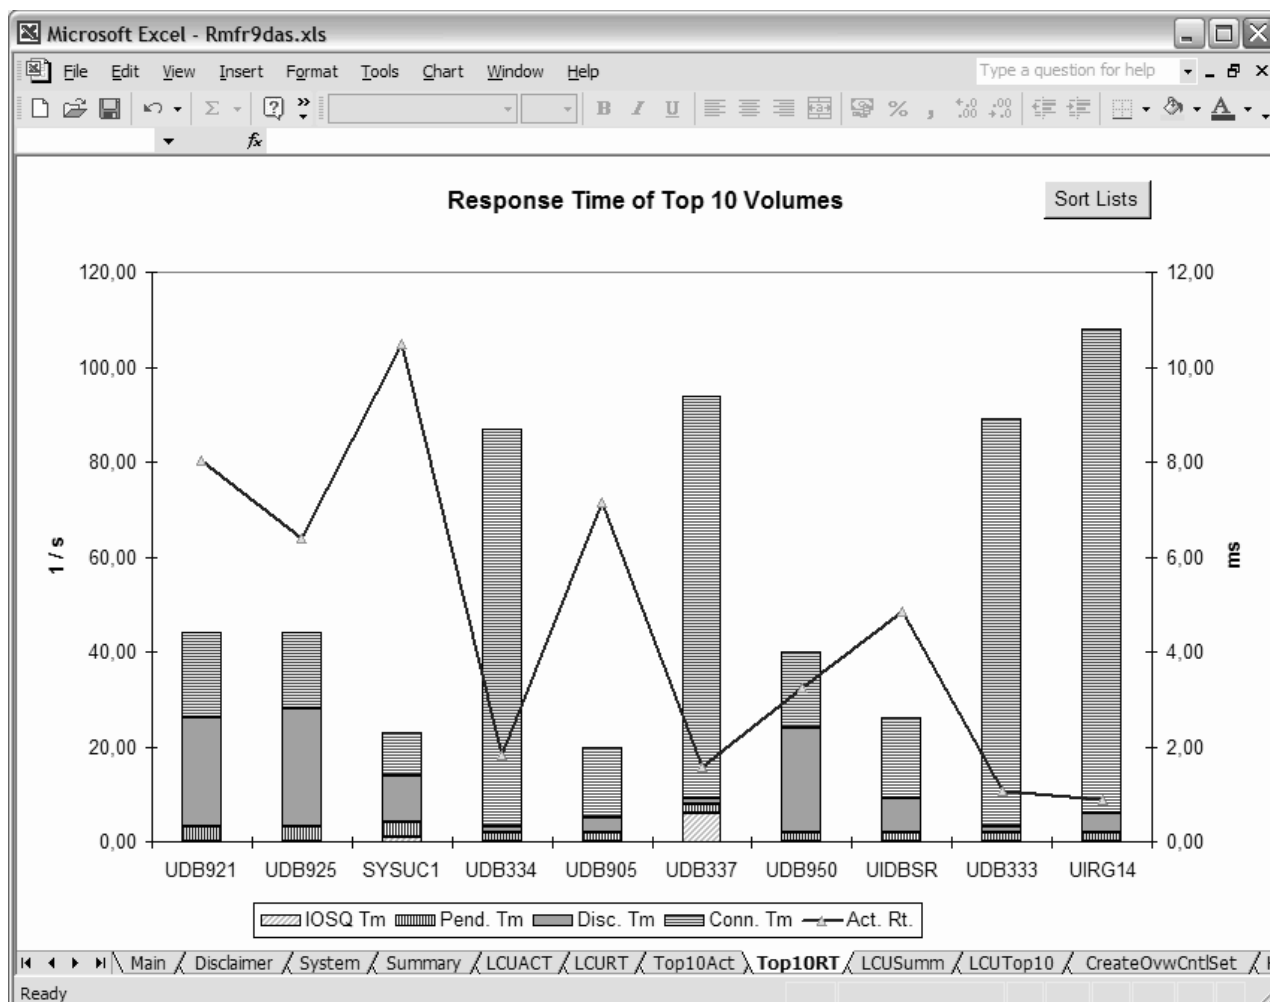


Figure 66. Response Time of Top-10 Volumes. Spreadsheet Reporter macro RMFR9DAS.XLS (DASD Activity Report - Top10Rt)

General DASD guidelines

Following are some guidelines for expected values for typical DASD volumes. There are two practical ways to monitor DASD activities with RMF:

- The Monitor I DASD Activity report
- The Monitor III DEV and DEVR reports

The most common procedure is to analyze the Monitor I DASD Activity report.

Note: With any RMF reports, the I/O information reported is from one MVS system only. If you are sharing I/O across multiple systems, you will need to review RMF reports from each of the sharing systems in order to see the complete I/O activity to the LCUs and devices (this is not true for the Cache RMF Reporter, however; since CRR data is obtained from the control unit, it does include I/O activity from all sharing systems). Contention from a sharing system will generally be seen as increased pend and disconnect times.

DASD performance

DASD Activity 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	2		
z/OS V1R9			SYSTEM ID UIG1			DATE 02/10/2007			INTERVAL 14.59.998											
			RPT VERSION V1R9 RMF			TIME 11.17.31			CYCLE 1.000 SECONDS											
TOTAL SAMPLES =		900	IODF =	00	CR-DATE:	10/16/2008	CR-TIME:	10.31.31	ACT: ACTIVATE											
STORAGE GROUP	DEV NUM	DEVICE TYPE	NUMBER OF CYL	VOLUME SERIAL	PAV	LCU	DEV 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	AVG NUMBER ALLOC	% ANY ALLOC	% MT PEND
	732A	33903	200	SYSXCA		002C	0.037	4.87	.000	.178	.000	.287	.000	4.58	0.02	0.02	0.0	0.0	100.0	0.0
	733A	33903	200	SYSAXA		002C	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0	0.0
	733D	33903	200	SYSDXA		002C	1.354	1.86	.000	.133	.000	.274	.052	1.54	0.21	0.22	0.0	14.0	100.0	0.0
				LCU		002C	1.391	1.94	.000	.134	.000	.275	.051	1.62	0.07	0.08	0.0	14.0	100.0	0.0
	969F	33909	10017	SYSIMS	1	0040	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0	0.0
	969F	33909	10017	SYSIMS	1	0040	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0	0.0
	9691	33909	10017	SYSUSR	1	0040	1.707	.536	.000	.133	.037	.239	.000	.297	0.05	0.05	0.0	4.0	100.0	0.0
	9692	33909	10017	SYSOPE	1	0040	0.041	2.22	.000	.135	.000	.214	.000	2.01	0.01	0.01	0.0	3.0	100.0	0.0
	9693	33909	10017	SYSCDS	1	0040	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0	0.0
	9694	33909	10017	SYSHM1	1	0040	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0	0.0
SMS	9695	33909	10017	SYSSMS	1	0040	0.012	.593	.000	.163	.000	.256	.000	.337	0.00	0.00	0.0	1.0	100.0	0.0
SMS	9696	33909	10017	SYSST1	1	0040	0.001	1.79	.000	.512	.000	.640	.000	1.15	0.00	0.00	0.0	0.0	100.0	0.0
SMS	9697	33909	10017	SYSST2	1	0040	0.001	2.94	.000	1.41	.000	1.54	.000	1.41	0.00	0.00	0.0	0.0	100.0	0.0
	9699	33909	10017	SYSHM2	1	0040	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0	0.0
DB2	969A	33909	10017	SYSSD1	1	0040	0.001	1.66	.000	.384	.000	.512	.000	1.15	0.00	0.00	0.0	0.0	100.0	0.0
DB2	969B	33909	10017	SYSSD2	1	0040	0.001	1.02	.000	.256	.000	.384	.000	.640	0.00	0.00	0.0	0.0	100.0	0.0
DB2	969C	33909	10017	SYSSD3	1	0040	0.001	2.05	.000	.640	.000	.768	.000	1.28	0.00	0.00	0.0	0.0	100.0	0.0
	969E	33909	10017	SYSDB2	1	0040	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0	0.0
	969F	33909	10017	SYSIMS	1	0040	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0	0.0
				LCU		0040	3.047	1.05	.365	.143	.036	.259	.002	.428	0.01	0.01	0.0	46.0	100.0	0.0

Figure 67. Direct Access Device Activity Report

In this report, the fields of most interest are highlighted.

Indicator**Field:** AVG RESP TIME**Description:** The average response time, in milliseconds, for an I/O to this device or LCU.**Related Fields:** DEVICE ACTIVITY RATE, AVG PEND TIME, AVG DISC TIME, AVG CONN TIME, AVG IOSQ TIME**Guideline:** For most situations, this measurement is the best way to quickly determine how well a device is performing.

AVG RESPONSE TIME can be considered as the overall measure of the health of a device's operation.

It is the sum of the AVG IOSQ TIME, AVG PEND TIME, AVG DISCONNECT TIME, and AVG CONNECT TIME. If response time is too high, then you can look further at these response time components and take remedial action.

What is too high? If your service level objectives are being met, then your response time is OK. If not, analyze what job is being delayed, and for how long (see prior section). Check the I/O rate to the device: if the I/O rate is insignificant, why bother tuning for response time?

Problem Area: I/O response time is a component of user response time; often it is the dominant component. High I/O response can delay user jobs and increase response times to unacceptable levels.**Potential Solution:** Determine the dominant response time component, and see "Improving your DASD performance" on page 124 for suggestions.

Indicator

Field: AVG IOSQ TIME

Description: This is the time measured when a request is being delayed in the MVS system.

Related Fields: DEVICE ACTIVITY RATE, AVG RESP TIME

Guideline: If IOSQ is greater than 1/2 times *service time* (where service time is DISC + CONN) then it warrants a closer look. In practice, there should be little queue time (less than 5 ms).

Problem Area: High IOSQ times contribute to high AVG RESP TIME. Combined with high I/O rates for important workloads, this could become a user response time problem.

Potential Solution: If your problem is high IOSQ time then you have the following options. Traditionally IOSQ problems are usually resolved by data movement either by separating active data sets or by moving active data to faster storage, for example, a coupling facility structure.

On an ESS subsystem, you have additional options.

- If PAV is not enabled for the device, enable it.
- If you are using static PAVs, assign more aliases to the device.
- If you are using dynamic PAV, then increase the number of PAVs associated in the pool for the subsystem.
- Check to ensure that all PAVs that should be bound to the device are online and operational. You can use the DEVSERV QP and DS QP,xxxx,UNBOX commands to do this.

Indicator**Field:** AVG PEND TIME**Description:** This value represents the average time (ms) an I/O request waits in the hardware (channel subsystem).**Related Fields:** AVG RESP TIME, DEVICE ACTIVITY RATE**Guideline:** A high PEND time suggests that the channel subsystem is having trouble initiating the I/O operation. There is a blockage somewhere in the path to the device. That might be due to:

- AVG DPB DELAY. Delay due to ESCON director ports being busy.
- AVG CUB DELAY. Delay due to the DASD control unit being busy, due to I/O from another sharing MVS system.
- AVG DB DELAY. Delay due to the device being busy, due to I/O from another sharing MVS system.
- Channel path wait. Whatever PEND time is not accounted for by the above three measures is due to delay for channel paths. This is the measure of channel delay that matters - not channel busy. If you think your channels are too busy, track this component of response time, for volumes that serve important work, to see if it is really a problem.

If you think shared DASD causes the problem, look at the DASD Activity report from the other system, taken at the same time. PEND time should generally be less than 4 ms.

Problem Area: High PEND times contribute to high AVG RESP TIME. Combined with high I/O rates for important workload, this could become a user response time problem.**Potential Solution:** High PEND times are usually caused by shared DASD contention or high channel path utilization.

On an ESS subsystem, multiple allegiance should reduce PEND times. If problems with PEND times exist you have the following options.

- Change the mix of data on the volume to reduce contention. If you can identify one data set that is contributing to most of the problem this may be eligible to be moved to a custom volume or moved into storage.
- Check channel utilization, especially if not using eight-path path groups. Changes have been made to CCWs to reduce channel overheads for the ESS. This will tend to lower channel utilization and increase throughput, however, using multiple allegiance and PAV successfully will cause increases in channel utilization.

Indicator

Field: AVG DISC TIME

Description: For cache controllers, DISC includes the time waiting while staging completes for the previous I/O to this particular device or until one of the four lower interfaces frees up from either transferring, staging or destaging data for other devices.

DISC can also be caused by a device which looks for the correct location on the disk to read or write the data.

Related Fields: AVG RESP TIME, DEVICE ACTIVITY RATE

Guideline: Generally, DISC time should be less than 17 ms. If the data is cached with reasonable hit ratios there should be little DISC time (3-5 ms).

Once your DISC time gets above 12 ms (about the amount of time you would expect for latency and typical seek times), you are probably experiencing RPS misses due to path contention.

Problem Area: High DISC times contribute to high AVG RESP TIME. Combined with significant I/O rates for important workloads, this could become a user response time problem.

Potential Solution: If the major cause of delay is disconnect time then you will need to do some further research to find the cause. You might use the *ESS Expert* for creating more detailed reports. The most probable cause of high disconnect time is having to read or write data from disks rather than cache. You need to check for the following conditions.

- High DISK to CACHE transfer rate, check the ESS Expert Disk to Cache Transport report and use the drill-down functions of the ESS Expert to identify the logical volumes that are experiencing response time problems.
- High disk utilization, check the ESS Expert Disk Utilization report, if the problem is limited to one physical disk, or array the best solution is to move data to balance performance across the subsystem. The other option is to move data to another subsystem, or to another storage medium, for example data in memory, or to change the way that the application uses this data.
- Low CACHE hit ratio — seen in the Cache Subsystem Device Overview report. If you are suffering from poor cache hit ratios there is little that you can do, you should look at the ESS as a whole, using ESS Expert and check that activity is balanced across both clusters of the ESS.
- NVS full condition, if the NVS is overcommitted you will see an increase in the values reported in the DFW BYPASS field of the Cache Subsystem Overview report. If this is a persistent performance indicator you may want to spread activity across more ESS subsystems.

Indicator**Field:** AVG CONN TIME

Description: This is time as measured by the channel subsystem during which the device is actually connected to the CPU through the path (channel, control unit, DASD) and transferring data. This time is considered good in most cases, because it is transferring data.

Related Fields: AVG RESP TIME, DEVICE ACTIVITY RATE

Guideline: Connect time is a function of block size and channel speed. Typical connect times often fall in the 2-6 ms range.

Problem Area: There are times when doing large searches in catalogs, VTOC, and PDS directories causes the connect time to become excessive. This results in connect time being considered "expensive" time in that no real data is transferred during this activity. Also during connect time, all parts related to an I/O operation (the entire path) must be available for use by that I/O. So no other user can do I/O to or from a device while that device is in connect mode.

One word of caution here - chained scheduling and seldom ending channel programs may generate high connect time which is not indicative of a performance problem.

Potential Solution: If you see devices with poor performance and high connect times then the cause is probably application related, high connect times are associated with large data transfers, either due to the use of large block sizes or to activities like DB2 prefetch that schedule large I/O transfers. If the application is not reporting poor response and other users of the volume are not impacted, then no further work is required. High connect time is an indication of large amounts of data being transferred. An application, such as DB2, transfers large I/Os for maximum efficiency.

Indicator

Field: % DEV UTIL

Description: This is the percentage of time the UCB is "busy" (see IOSQ time). This includes both the time when the device was involved in an I/O operation (CONN and DISC) as well as any time it was reserved but not involved in an I/O operation.

Related Field: AVG IOSQ TIME

Guideline: A device overload condition is not always easy to detect. The device utilization as reported by RMF is not a measure of goodness or badness. If there is only one active data set on the device, and presuming that this data set cannot be made resident in processor storage, there is nothing wrong with a very high device utilization.

However, if multiple online users are all trying to get to data sets on the same device, they can be delayed due to the device being busy (from the other users). This will show up as increased IOSQ time.

If the utilization rises above 35%, response times may start to suffer, due to increased queuing for the device. If you have online response-sensitive data on a non-cached volume, you may want to investigate if utilization rises above the 25-30% range.

Problem Area: High device utilization can lead to increased IOSQ time, if users are delayed while the device is in use by others.

Potential Solution: Potential ways to decrease device utilization include:

- Caching the device (if not already cached). With a good hit ratio service time may be reduced significantly, which will reduce device utilization dramatically.
- Distribute heavily used data sets onto different volumes.
- Remove poor cache users. Turn off caching for data sets or volumes that do not achieve good hit ratios. As discussed before, this frees up cache for users that benefit more.

Indicator

Field: Multiply DEVICE ACTIVITY RATE by AVG CONN TIME

Description: Utilization (% busy) of the paths to DASD devices.

Related Fields: AVG PEND TIME, AVG DISC TIME

Guideline: The maximum allowable utilization is not a fixed number, but it can be higher as the overall hit ratio increases. In an average 4-path system (DLSE), cached or non-cached, a general guideline is not to exceed 50%, if the goal is to maintain good response for interactive users.

If the workload using these paths is not response-time sensitive, you may not care about high path utilization. For example, if you have only four jobs accessing four paths, then there is no delay and high path utilization is desired.

Since the storage path utilization is not directly reported, an estimate can be made using the RMF reported LCU I/O rate and the average LCU connect time. This amounts to a simple calculation, which you should execute for each system that accesses devices attached to the storage control unit:

1. Multiply I/O rate by average connect time
2. Divide the result by four (a 3990 has four storage paths)
3. Divide by 10 to convert to a percentage (because I/O rate is given in I/O per second, while connect time is reported in millisecond time units)

Problem Area: Higher utilizations will be an issue if they are causing delay to important work. This delay will show up as increased pend and disconnect times for volumes sharing these paths.

Potential Solution: If your storage path utilization is above 50% and causing delay to important work, a potential solution is to balance activity across control units by moving data sets.

Indicator

Field: AVG DISC TIME, ASYNC and STAGE rates

Description: These fields give you an indication of "lower interface" contention. The lower interfaces are the connections between the DASD devices and the DASD control unit.

Related Fields: LCU cache hit ratio

Guideline: Direct measure of lower interface utilization is not available from any reporting tool, but normally you should not have to worry about the lower interfaces. If the utilization of the storage paths and the hit ratios are within acceptable limits, the lower interface utilization should be fine.

The lower interfaces can be overloaded on cache controllers, however, if the hit ratio for the storage control unit falls below about 70% and a high subsystem throughput is imposed.

Problem Area: Problems here would be reflected in increased DISC times. If you suspect a problem, check your STAGE and ASYNC rates from the Cache RMF Reporter. Correlate both STAGE and ASYNC rates to your DISC times (for example, plot DISC time vs. STAGE + ASYNC rates).

Potential Solution: If you see a significant correlation between STAGE + ASYNC and DISC, investigate cache tuning or workload movement to reduce cache overload.

DASD performance in the sysplex

The Shared Device Activity report is a sysplex report that is available for DASD and tape devices. In this context we will discuss the DASD version of this report.

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.

Shared DASD Activity report

SHARED DIRECT ACCESS DEVICE ACTIVITY																	PAGE 1		
z/OS VIR10			SYSPLEX RMFPLEX1				DATE 04/10/2008				INTERVAL 014.59.946								
			RPT VERSION VIR10 RMF				TIME 16.30.00				CYCLE 1.000 SECONDS								
TOTAL SAMPLES(AVG) = 900.0 (MAX) = 900.0 (MIN) = 900.0																			
DEV NUM	DEVICE TYPE	VOLUME SERIAL	PAV	SMF SYS ID	IODF SUFF	LCU	DEVICE ACTIVITY RATE	AVG RESP TIME	AVG IOSQ TIME	AVG DPB DLY	AVG CUB DLY	AVG DB DLY	AVG PEND TIME	AVG DISC TIME	AVG CONN TIME	% DEV CONN	% DEV UTIL	% DEV RESV	AVG NUMBER ALLOC
06AA	3380E	DPVOL1		*ALL			14.589	3.5	0.0	0.0	0.0	0.2	0.4	0.6	2.5	3.63	4.45	0.1	651
				MVS1	01	0043	4.550	4.8	0.0	0.0	0.0	0.0	0.2	0.5	4.1	1.86	2.10	0.0	205
				MVS2	02	0143	5.489	3.1	0.0	0.0	0.0	0.3	0.8	0.6	1.7	0.91	1.24	0.1	210
				MVS3	03	0043	4.550	3.8	1.0	0.0	0.0	0.2	0.3	0.6	1.9	0.86	1.11	0.0	236
077A	3380E	ERBDAT		*ALL			15.186	9.3	0.5	0.0	0.0	0.2	1.2	5.2	2.4	3.62	11.48	0.4	1078
077A				MVS1	01	0043	4.069	5.2	2.0	0.0	0.0	0.8	1.1	0.4	1.7	0.68	0.84	0.0	343
0872				MVS2	02	0143	6.438	14.8	0.5	0.0	0.0	0.0	1.0	10.2	3.1	2.00	8.58	0.0	357
077A				MVS3	03	0052	4.679	8.0	2.0	0.0	0.0	0.1	1.6	2.4	2.0	0.94	2.06	0.4	378

Figure 68. Shared DASD Activity Report

This example reports about a sysplex consisting of three systems (MVS1, MVS2, and MVS3). Only two devices are shown.

The second device does not have the same device number on all three systems.

Indicator
Field: *ALL

Description: The system line shows the device activity contributed by all systems in the sysplex.

Guideline: The report gives you an overall performance picture of DASD devices that are shared between MVS systems in a sysplex.

For each shared DASD device the report contains one line for each system that has access to it.

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.

Analyzing specific DASD problems

This section will focus on how to analyze specific DASD problems. It is assumed that you have a good overall knowledge of your DASD subsystem (as discussed in “I/O subsystem health check” on page 84) before proceeding.

One approach to investigating I/O problems can be summarized:

1. Is there a problem?
2. Who’s got it?
3. Do I care?
4. What can be done about it?

Is there a problem? The Monitor III SYSINFO report shows whether there is any device delay. It also measures the size of the problem by the number of users affected.

SYSINFO report

```

RMF V1R10 System Information                               Line 1 of 27
Command ==>>>                                           Scroll ==>> HALF

Samples: 100      System: PRD1 Date: 11/28/08 Time: 10.32.00 Range: 100 Sec

Partition:  MVS1   2094 Model 716      App1%:    63 Policy: STANDARD
CPs Online:   4    Avg CPU Util%: 73    EApp1%:   65 Date: 11/28/08
AAPs Online:  0    Avg MVS Util%: 84    App1% AAP:  0 Time: 14.05.07
IIPs Online:  0                                App1% IIP:  0

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

*SYSTEM 34 664 26      13.95 5.1 5.0    1.9 4.1 7.0 2.6 2.0 2.0
*TSO    50 534 8       13.95 2.6 2.1    0.4 1.5 2.0 0.8 0.0 0.0
*BATCH  26 11 10      0.00 1.5 1.4    1.4 1.7 0.5 1.8 1.0 2.0
*STC    27 115 8      0.00 1.1 1.5    0.1 1.0 4.5 0.1 1.0 0.0
*ASCH   3 0 0        0.00 0.0 0.0    0.0 0.0 0.0 0.0 0.0 0.0
*OMVS   2 0 0        0.00 0.0 0.0    0.0 0.0 0.0 0.0 0.0 0.0
*ENCLAVE 5 4 N/A      N/A 0.2 N/A    3.7 N/A 0.0 N/A N/A N/A
PRIMEBAT W 26 11 10 46.0 0.06 1.5 1.4 1.4 1.7 0.5 1.8 1.0 2.0
NRPRIME S 26 11 10 46.0 0.06 1.5 1.4 1.4 1.7 0.5 1.8 1.0 2.0
        1 23 9 9 27.9 0.06 0.9 1.4 0.8 1.6 0.5 1.8 1.0 2.0
        2 29 0 0 54.2 0.02 0.1 0.0 0.2 0.0 0.0 0.0 0.0 0.0
        3 59 1 1 .000 0.00 0.6 0.0 0.4 0.0 0.0 0.0 0.0 0.0
PRIMETSO W 50 527 8 .759 13.98 2.6 2.1 0.4 1.5 2.0 0.8 0.0 0.0
TSOPRIME S 50 527 8 .759 13.98 2.6 2.1 0.4 1.5 2.0 0.8 0.0 0.0
        1 48 526 8 .403 13.98 2.1 1.9 0.3 1.3 2.0 0.8 0.0 0.0
        2 75 1 1 30.6 0.08 0.3 0.2 0.1 0.1 0.0 0.0 0.0 0.0
        3 75 0 0 126 0.02 0.1 0.0 0.0 0.0 0.0 0.0 0.0 0.0
    
```

Figure 69. Monitor III System Information Report

Who’s got it? Look at the Delay report shown in Figure 70 on page 119. Who (which address space) has a device delay problem? This is more important than the number of users affected. Remember, I/O service is not democratic. One important job or transaction being delayed can be much more significant than several other transactions or jobs. The Delay report shows where delays exist by specific resource for specific users. This aids in the assessment of performance impact: “Who will be helped? How much?”

DELAY report

RMF V1R10 Delay Report													Line 21 of 49				
Command ==>													Scroll ==> HALF				
Samples: 100													System: PRD1	Date: 04/07/08	Time: 10.32.00	Range: 100	Sec
Name	Service CX Class	WFL Cr	USG %	DLY %	IDL %	UKN %	----- PRC	% Delayed for	-----	Primary Reason							
							DEV	STR	SUB	OPR	ENQ						
MISTYDFS	B NRPRIME	0	0	100	0	0	0	0	0	100	0	Message					
BHOLEQB	B NRPRIME	0	0	49	0	0	0	0	0	0	49	SYSDSN					
BHOLW01	B NRPRIME	0	0	33	0	1	0	0	0	33	0	Message					
BRHI	T TSOPRIME	0	0	4	0	0	0	4	0	0	0	COMM					
CONSOLE	S SYSSTC	0	0	3	0	97	0	3	0	0	0	LOCL					
VTMLCL	S SYSSTC	0	0	2	0	98	0	2	0	0	0	LOCL					
STARTING	T TSOPRIME	0	0	2	0	7	0	2	0	0	0	COMM					
CATALOG	S SYSSTC	0	0	1	0	99	0	1	0	0	0	LOCL					
INIT	S SYSSTC	0	0	1	49	0	0	1	0	0	0	LOCL					
BHOLDEV3	B NRPRIME	20	20	80	0	1	79	0	0	0	0	BBVOL1					
BTEUDASD	B NRPRIME	22	22	77	0	1	76	0	0	0	0	BBVOL1					
BAJU	T TSOPRIME	24	24	75	0	1	75	0	0	0	0	BBVOL1					
KLSPRINT	B NRPRIME	25	25	75	0	0	75	0	0	0	0	BBVOL1					
BHIM	T TSOPRIME	31	11	25	64	0	18	5	0	0	0	BBVOL1					
BHOLPRO1	B NRPRIME	52	13	12	0	0	12	0	0	0	0	BHOLPRO2					

Figure 70. Monitor III Delay Report

Do I care? Who is being delayed, and how much, often determines whether any action is necessary. Very important work requires immediate attention. Less important work? Small delays? Well, maybe you can look at it tomorrow.

What can be done about it? To identify the solution to an I/O problem, you must identify where the I/O pain is.

Remember, Monitor III can direct you to any device that is causing delays. This point of view is important; just because a device might have a high activity rate or a high response time, it MAY NOT be causing significant delays, or it may be delaying work you do not care about.

The reports that are useful here include:

DASD performance

DEV report

```
RMF V1R10 Device Delays                               Line 1 of 9
Command ==>                                           Scroll ==> HALF
Samples: 100      System: PRD1 Date: 04/07/08 Time: 10.32.00 Range: 100 Sec

Jobname  C Class      DLY  USG  CON  ----- Main Delay Volume(s) -----
          %    %    %    % VOLSER  % VOLSER  % VOLSER  % VOLSER
BHOLDEV3 B NRPRIME  79  15   4   72 BBVOL1
BTEUDASD B NRPRIME  76  20   4   75 BBVOL1
KLSPRINT B NRPRIME  75  23   4   75 BBVOL1
BAJU     T TSOPRIME 75  19   1   71 BBVOL1
BHIM     T TSOPRIME 18   6   1   16 BBVOL1  1 430D13
BHOL     T TSOPRIME  4  11   6   9  BBVOL1
JES2     S SYSSTC   4   3   0   6  SYSPAG
BHOLNAM4 B NRPRIME   1   0   0   1  BBVOL1
```

Figure 71. Monitor III Device Delays Report

The Device Delays report shows jobs in your system that are delayed due to device contention.

Jobnames (address space names) are listed in descending order in delay percentages. For each jobname, the names of up to 4 device volumes are listed indicating which volumes are causing the delay. If a job is being delayed, you can immediately see which volume is causing the most delay.

DEVR report

```

RMF V1R10 Device Resource Delays                               Line 1 of 16
Command ==>                                                    Scroll ==> HALF

Samples: 100      System: PRD1  Date: 04/07/08  Time: 10.32.00  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  %  %
BBVOL1   41.8 .099  88  21  66  PND  1  3380K      KLSPRINT B  NRPRIME  6  92
  0222                3990-2      BHOLDEV3 B  NRPRIME  2  92
                BTEUDASD B  NRPRIME  5  91
                BAJU      T  TSOPRIME  3  91
                BHIM      T  NRPRIME  1  22
                BHOL      T  TSOPRIME  3  10
                RMFGAT   S  SYSSTC  0  1
                BHOLNAM4 B  NRPRIME  0  1
                BHOLPRO1 B  NRPRIME  0  1
SYSPAG    10.9 .028  19  6  13  PND  0  3380K      BVAU      T  TSOPRIME  6  14
  0225                3990-2      JES2      S  SYSSTC  0  7
430D13    1.1 .032  4  2  2  PND  0  3380K      BENK      T  TSOPRIME  0  1
  0220                3990-2      BHIM      T  TSOPRIME  0  1
                BHOL      T  TSOPRIME  2  0
SYSCAT    0.3 .017  0  0  0  PND  0  3380K      INIT      S  SYSSTC  0  1
  0221                3990-2

```

Figure 72. Monitor III Device Resource Delays Report

The DEVR report is used to evaluate the performance of each volume. It differs from the DEV report in that it lists a volume and all of the jobs that are being delayed for that volume. Thus, if you find a volume that RMF suggests is performing poorly, you can use Monitor III to determine which jobs it is affecting and how badly it is affecting any particular job.

You can get similar information from the Monitor I DASD Activity report.

Report Analysis

- In this case, from the SYSINFO report (shown in 69) we know that 1.7 jobs in service class NRPRIME are delayed because of a device.
- From the Delay report (shown in 70) we know the jobs in NRPRIME are waiting for volume BBVOL1.
- From the DEVR report (shown in 72) we identify that BBVOL1 is a 3380K that is not cached.

Replacing this disk by a faster device would be appropriate.

Performance analysis on data set level

The following sample reports do not show specific problems but can illustrate how performance analysis can be done on data set level.

If you need further analysis for a specific volume that you see in the DEVR report, you can navigate through cursor-sensitive control to some reports which provide information on data set level.

```

RMF V1R10 Device Resource Delays
Command ==> _
Line 1 of 19
Scroll ==> HALF

Samples: 76      System: D3      Date: 05/03/08 Time: 12.39.00 Range: 120 Sec

Volume S/  Act  Resp  ACT  CON  DSC  PND %,  DEV/CU
/Num PAV  Rate Time  %  %  %  Reasons  Type      Jobname  C Class  %  %

COUPLB S   2.4 .115  21  9  0  PND 12 33903  XCFAS  S SYSSTC  16  7
9191      DB  8 3990-3
          DPB  3
COUPLP S   0.8 .224  18  1  0  PND 17 33903  XCFAS  S SYSSTC  22  5
91CD      DB 15 3990-3
SGT103 S   0.7 .056  3  0  1  PND 2 33903  U015074 T TSOSLOW  1  1
9182      DPB 1 3990-3  IXGLOGR S SYSTEM  1  1
          CATALOG S SYSTEM  0  1
          U015041 T TSOSLOW  0  1
CMNAF8 S   0.5 .105  5  0  0  PND 5 33903  CATALOG S SYSTEM  0  4
BAF8      DB  5 3990-6  RMFGAT  S SYSSTC  1  1
HDCAT S   2.3 .009  0  0  0  PND 0 33903  CATALOG S SYSTEM  1  1
910E      3990-3  *MASTER* S SYSTEM  1  0
    
```

Figure 73. Monitor III - Device Resource Delays Report

If you select volume SGT103 and click on Enter, you get this report:

DSNV Report

```

RMF V1R10 Data Set Delays - Volume
Command ==> _
Line 1 of 5
Scroll ==> HALF

Samples: 76      System: D3      Date: 05/03/08 Time: 12.39.00 Range: 120 Sec

----- Volume SGT103 Device Data -----
Number: 9182      Active: 3%      Pending: 2%      Average Users
Device: 33903    Connect: 0%     Delay DB: 0%     Delayed
Shared: Yes      Disconnect: 1%  Delay CU: 0%     0.1
                  Delay DP: 1%

----- Data Set Name ----- Jobname ASID DUSG% DDLY%
IXGLOGR.ATR.UTCPLXHD.RM.DATA.A0000010.DATA IXGLOGR 0021 1 1
U015074.CALENDAR.DATA U015074 0110 0 1
U015041.COURSEIN.DATA U015041 0125 0 1
SYS1.VVDS.VSGT103 SYS1.CATALOG 0051 0 1
SYS1.VTOCIX.VSGT103 U015074 0110 1 0
    
```

Figure 74. Monitor III - Data Set Delays by Volume

In addition to general information about this volume, you see a list with all data sets which are currently used. Selecting a specific job, you navigate to the next report.

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.

It might happen that -- N/A -- is provided instead of a data set name. This happens if only those I/O instructions have been detected for which no data set information is provided by the SMS subsystem, like

- I/Os to system data sets (like paging or spooling)
- I/Os to any data set which was opened prior to SMS subsystem initialization
- I/Os like SENSE or RELEASE
- System I/Os not done by an access method

DSNJ report

RMF V1R10 Data Set Delays - Job				Line 1 of 2			
Command ==> _				Scroll ==> HALF			
Samples: 76	System: D3	Date: 05/03/08	Time: 12.39.00	Range: 120	Sec		
Jobname: U015074	EXCP Rate	0.1	Connect:	0%			
ASID	-----	Data Set Name	-----	Volume	Num	DUSG%	DDLY%
0110		U015074.CALENDAR.DATA		SGT103	9182	0	1
		SYS1.VTOCIX.VSGT103		SGT103	9182	1	0

Figure 75. Monitor III - Data Set Delays by Job

This report shows all active data sets for a specific job. If you now are interested to know whether some other job is working with this data set, again, you can navigate to the corresponding report.

DSND report

```
RMF V1R10 Data Set Delays                               Line 1 of 1
Command ==> _                                           Scroll ==> HALF
Samples: 76      System: D3      Date: 05/03/08  Time: 12.39.00  Range: 120  Sec
Input Data Set Name: SYS1.VTOCIX.VSGT103
----- Data Set Name ----- Volume  Jobname  ASID  DUSG%  DDLY%
SYS1.VTOCIX.VSGT103          SGT103  U015074  0110    1    0
```

Figure 76. Monitor III - Data Set Delays by Data Set

In this sample, no other job than U015074 is using the data set.

Improving your DASD performance

This section will suggest actions you may want to consider, to improve the performance of your DASD.

- Obviously, the best response time is achieved if the DASD I/O never happens. Eliminate I/O using buffering or data-in-memory techniques wherever possible.
- Next best is resolving the I/O in cache.
- Last is an I/O to DASD itself.

Also remember ways to increase the "parallelism" of your I/O processing.

- This could include the Data Striping capability of DFSMS, where sequential data is distributed across multiple I/O devices to allow concurrent I/O.
- As a last resort, parallelism could also mean manually splitting data across volumes, or duplicating data on multiple volumes.

In addition, here are actions to consider to address the following symptoms:

High DISC time

- Cache the volume, if it is a suitable cache candidate.
- Un-cache the volume, if it is a poor cache candidate.
- Tune cache.
 - Best to use DCME to dynamically manage cache for you.
 - If necessary, manually tune cache by turning off poor cache users.
- Review block size. Increase where applicable.
- Add paths (if not already 4-path).
- Use faster device.
- Move or copy data sets to other volumes to reduce contention.
- Reduce arm movement via data set placement (seek analysis required).

High IOSQ time

- Cache the volume, if it is a suitable cache candidate.
- Tune cache.
 - Best to use DCME to dynamically manage cache for you.
 - If necessary, manually tune cache by turning off poor cache users.
- Decrease DISC, CONN, PEND times. Decreasing the other response time components will also decrease IOSQ time proportionally.
- Move or copy data sets to other volumes to reduce contention.
- Run requests sequentially to reduce contention.

High CONN time

- Cache the volume, if it is a suitable cache candidate.
- Use indexed VTOCs.
- In a cached environment, use faster channels.
- Use PDSEs if you suspect directory search time is a problem.

High PEND time

- Address the biggest subcomponent of PEND time (see page 111).
 - Use faster channels, or add more channels (if not already 4-path).
 - Move data sets to avoid contention from sharing systems.

Tape indicators and guidelines

This section discusses guidelines for analyzing and resolving tape performance issues for the 3480 and 3490 families of tape devices.

From a performance perspective, you can think of tape as being like DASD with long connect times. All of the things we used to do for single path DASD are true for tape: short strings, many control units and channels. If the subsystem is configured with one channel per control unit function, the only reason to add channels or cross communication of control units is availability, not performance.

Identifying tape-bound jobs

An indication of a tape-bound job is when a large proportion of the total time for the job is spent in tape activity and a small amount in CPU (or other) activity. The job termination messages, NUMBER OF MOUNTS and AVG MOUNT TIME from the Magnetic Tape Device Activity report can help determine if a job is tape-bound.

Look at the job termination messages for the ratio of CPU time to the total elapsed time as indication for tape-bound jobs. You can also use the Monitor III Delay report, to verify that no other (non-I/O) significant delays exist for the job in question and that tape delay is dominant.

Once a job is identified as tape-bound then further evaluation and measurements may be justified.

RMF measurements for tape monitoring

The options required to measure and monitor the tape subsystems may not be the default options used in most user environments. The reason is that the jobs and applications using tape are normally batch tasks, more concerned with throughput than response time.

Some RMF options that must be enabled to have meaningful RMF data for tape subsystem measurements are discussed. These options may be modified in the ERBRMF00 parmlib member or may be modified online for current sessions:

Monitor I Session Options

INTERVAL	This value needs to be small (preferably no more than 15 minutes) if the details of the tape subsystem performance are to be evaluated.
DEVICE	TAPE must be one of the device types specified.
IOQ	To request I/O queuing activity for tape LCUs, IOQ(TAPE) must be specified.
CHAN	This option specifies that the channel activity is to be measured. The channel utilization may be of interest in evaluating multiple control units on the same channel, or for evaluating the requirement for additional channel paths.

Magnetic Tape Device Activity report

MAGNETIC TAPE DEVICE ACTIVITY																	PAGE	1	
z/OS V1R10		SYSTEM ID SYSE				DATE 11/28/2008				INTERVAL 15.00.027									
				RPT VERSION V1R10 RMF				TIME 23.45.00				CYCLE 1.000 SECONDS							
TOTAL SAMPLES =		810	IODF = 00		CR-DATE: 05/13/2008				CR-TIME: 12.03.30				ACT: ACTIVATE						
DEV	DEVICE	VOLUME	LCU	DEVICE	AVG	AVG	AVG	AVG	AVG	AVG	AVG	%	%	%	NUMBER	AVG	TIME		
NUM	TYPE	SERIAL		ACTIVITY	RESP	IOSQ	CMR	DB	PEND	DISC	CONN	CONN	UTIL	RESV	OF	MOUNT	DEVICE		
				RATE	TIME	TIME	DLY	DLY	TIME	TIME	TIME	CONN			MOUNTS	TIME	ALLOC		
1A71	3590L		001D	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0	0	0		
1A72	3590L		001D	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0	0	0		
1A73	3590L		001D	0.005	157	.000	.000	.000	.448	156	1.09	0.00	0.08	0.0	0	0	0		
			LCU	001D	0.005	157	.000	.000	.448	156	1.09	0.00	0.03	0.0	0	0	0		
1A78	3590L		001F	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0	0	0		
1A79	3590L		001F	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0	0	0		
1A7A	3590L		001F	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0	0	0		
1A7B	3590L		001F	0.005	414	.000	.000	.000	.448	413	1.09	0.00	0.20	0.0	0	0	0		
			LCU	001F	0.005	414	.000	.000	.448	413	1.09	0.00	0.05	0.0	0	0	0		
1A7C	3590L		0020	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0	0	0		
1A7F	3590L		0020	0.006	126	.000	.000	.000	1.79	122	1.61	0.00	0.08	0.0	0	0	0		
			LCU	0020	0.006	126	.000	.000	1.79	122	1.61	0.00	0.04	0.0	0	0	0		
1A85	3490L		0021	0.007	260	.000	.000	.000	.896	258	1.32	0.00	0.19	0.0	0	0	0		
			LCU	0021	0.007	260	.000	.000	.896	258	1.32	0.00	0.19	0.0	0	0	0		

Figure 77. Magnetic Tape Device Activity Report

Response time components

Response time is a key measure of how well your I/O is performing, and a key indicator of what can be done to speed up the I/O. It is important to understand the components of this response and what they mean. Tape response time is the sum of queue, pend, connect and disconnect time (see page 109). The following lists the response time components and what causes them:

Queue Time

AVG IOSQ TIME is a measure of an I/O request waiting in software for the device to become available to service the request. For tape this usually means a REWIND/UNLOAD is in progress.

Pend Time

AVG PEND TIME reflects the time an I/O spends waiting to get a path to the control unit. This is the time spent waiting for an available channel path.

Disconnect Time

AVG DISC TIME is one of the more significant measures of a heavily loaded tape subsystem. Disconnect time is that time spent waiting for the buffer in the control unit to become available to transfer data; it includes the tape rewind time. One example of a disconnected I/O operation is a buffer-full condition.

Connect Time

AVG CONN TIME is predominantly time spent transferring data between the control unit and the processor. Connect time is directly related to data block size, and to the amount of data transferred for each I/O.

General tape indicators

Indicator

Field: AVG DISC TIME (LCU), AVG PEND TIME (LCU), AVG CONN TIME (LCU)

Description: Ratio of DISC + PEND time to CONN time. Lower ratios indicate less contention for channels and control units.

Related Fields: DEVICE ACTIVITY RATE. If the I/O rate to the controller is negligible, don't bother with any further analysis.

Guideline: Percent channel busy is not a good indicator of tape performance. Instead look at the ratio of LCU DISC + PEND time to CONN time. If the ratio exceeds 2 to 1 for extended periods of time (with an I/O rate high enough to be of interest), this is an indication of channel and/or control unit contention.

Problem Area: This is an indicator that response times are increasing (increased DISC and PEND) due to an overloaded control unit or channels.

Potential Solution: If you have significant DISC time and little PEND time, add more or faster control units. If you have significant DISC **and** PEND times, add more/faster control units and more/faster channels. As a guideline, you probably want to configure no more than four active drives per control unit. Beyond four **concurrently active** drives behind a single control unit function, data throughput for the additional drives drops off dramatically.

Indicator

Field: AVG CONN TIME (LCU)

Description: Predominantly, data transfer time between the control unit and the CPU.

Related Fields: AVG RESP TIME, DEVICE ACTIVITY RATE

Guideline: CONN times in the 20-40 ms range are typical. If your CONN time is significantly lower (for example, 4-5 ms), you may have smaller than optimal block sizes, resulting in more I/Os required to transfer the data.

Problem Area: This could mean delay to the application, if a high number of I/Os are being generated to transfer the data. Decreasing the I/O rate by increasing the amount of data transferred per I/O will speed up the application run-time.

Potential Solution: Increase block size and/or the number of host buffers where applicable.

Improving your tape performance

This section will suggest actions you may want to consider to improve tape performance.

The key to good performance is to maximize data transfer, minimize control unit communication and command processing, and minimize PEND and DISC times. This can be accomplished by using large block sizes, allocating no more than four concurrently active drives per control unit and not cross-coupling control units.

As a checklist, review the following:

- Review block size. Increase where applicable.
- Allocate data sets to DASD. TMM (tape mount management) is a method to accomplish this and significantly reduce tape mount requirements.
- Increase host buffers for priority jobs. Increase buffers from default of 5, to as high as 20 (but no more than 20).
- Use VIO for temporary data sets.
- Use faster channels.
- Add channels (if not already one channel per control unit function).
- Reduce mount delay.
 - use ACLs/ICLs (cartridge loaders) for nonspecific mounts
 - use an automated tape library for specific mounts
 - review physical logistics for manual mount efficiencies (e.g. are tapes stored near the operators, and near the drives?)
- Add control unit functions (probably want no more than 4 concurrently active drives per control unit).
- Use enhanced-capacity cartridges and tape drives.
- Restructure for more parallelism (for example, to allow application to use multiple drives).
- Use the *Improved Data Recording Capability (IDRC)* to reduce the amount of data transfer required.

Summary

Typically, I/O processing is the largest single component of transaction response time. So, while I/O performance management can be complicated and time consuming, the return in terms of response time and throughput can be significant.

Look for ways to avoid I/O altogether: Through optimized use of buffers, newer data-in-memory techniques, and eliminating unnecessary I/O.

Cache the I/O that will benefit from cache, and keep cache "tuned" to the extent that poor cache users are not allowed to get in the way of better cache users.

As with all performance management you need to:

- Understand the overall system picture (CPU, processor storage) and the overall I/O subsystem picture before tuning.
- Measure the delay to important workloads.
- Review "related fields" before tuning. For example, if response time for a device seems high, also check the I/O rate, type of workload using the device, and % DLY to users, before deciding what (if any) action to take.
- Determine where your leverage for improvement exists. Start by seeing which component of response time (DISC, CONN, etc.) is dominant, and take further action accordingly.

Chapter 5. Analyzing processor storage activity

Let's Analyze Your Storage

This topic discusses how to analyze a processor storage problem. Topics include:

- What are the indicators in RMF to review, to verify there is a problem?
- Some examples of common problems that RMF can help you resolve.
- Potential tuning actions.

Simply put, resolving processor storage problems generally means prioritizing which workloads get access to processor storage, finding ways to move fewer pages, or speeding up the movement of pages.

Do you have a processor storage problem?

Here are the primary RMF indicators you may wish to use in assessing your processor storage configuration. After listing these indicators, we will then discuss each one in more detail, with guidelines for each value that may indicate contention and suggestions for further action.

We will refer to the components of storage as *central storage (CS)* and *auxiliary storage (AUX)*.

Never start any tuning effort simply because one indicator seems to be a problem. Always check other related indicators to build a good understanding of your overall processor storage situation and overall system situation, and only then begin to plan your tuning efforts. For example, if your UIC is low, also look at demand paging from auxiliary storage and storage delays in Monitor III.

Table 3 shows for each monitor, which fields in which report can most often help detect and solve a processor storage problem.

Table 3. Processor Storage Indicators in Monitor III and Monitor I

MONITOR	REPORT	FIELDS
Monitor III	SYSINFO	Average Number Delayed for STOR
	DELAY	%Delayed for STR
	STOR	DLY%
Monitor I	Paging Activity	PAGE MOVEMENT WITHIN CENTRAL STORAGE PAGE-IN EVENTS AVG HIGH UIC TOTAL FRAMES AVAILABLE
	Page/Swap Activity	SLOTS ALLOC SLOTS USED % IN USE
	DASD Activity	AVG RESP TIME

Monitor III indicators

This section reviews the most frequently used processor storage indicators in the Monitor III reports. We will start by looking at SYSINFO and DELAY reports, which can lead you to the STOR report.

SYSINFO report

```

RMF V1R10 System Information                               Line 1 of 23
Command ==>                                             Scroll ==> HALF

Samples: 100      System: PRD1 Date: 04/07/08 Time: 10.32.00 Range: 100 Sec

Partition:  MVS1    2094 Model 716      Appl%:      63 Policy: STANDARD
CPs Online:   4     Avg CPU Util%: 73    EApp%:     65 Date:  11/28/08
AAPs Online:  0     Avg MVS Util%: 84    App% AAP:  0 Time:  14.05.07
IIPs Online:  0     App% IIP:   0
    
```

Group	T	WFL %	--Users-- TOT ACT	RESP Time	TRANS /SEC	-AVG PROC	USG- DEV	-Average PROC	Number DEV	Delayed STOR	For SUBS OPER	- ENQ
*SYSTEM		34	664 26		13.95	5.1	5.0	1.9	4.1	7.0	2.6	2.0
*TSO		50	534 8		13.95	2.6	2.1	0.4	1.5	2.0	0.8	0.0
*BATCH		26	11 10		0.00	1.5	1.4	1.4	1.7	0.5	1.8	1.0
*STC		27	115 8		0.00	1.1	1.5	0.1	1.0	4.5	0.1	1.0
*ASCH			3 0		0.00	0.0	0.0	0.0	0.0	0.0	0.0	0.0
*OMVS			2 0		0.00	0.0	0.0	0.0	0.0	0.0	0.0	0.0
*ENCLAVE		5	4 N/A		N/A	0.2	N/A	3.7	N/A	0.0	N/A	N/A
PRIMEAPP	W		3 0	.000	0.00	0.0	0.0	0.0	0.0	0.0	0.0	0.0
APPRIME	S		3 0	.000	0.00	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	1		2 0	.000	0.00	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	2		1 0	.000	0.00	0.0	0.0	0.0	0.0	0.0	0.0	0.0
PRIMEBAT	W	26	11 10	46.0	0.06	1.5	1.4	1.4	1.7	0.5	1.8	1.0
NRPRIME	S	26	11 10	46.0	0.06	1.5	1.4	1.4	1.7	0.5	1.8	1.0
	1	23	9 9	27.9	0.06	0.9	1.4	0.8	1.6	0.5	1.8	1.0
	2	29	0 0	54.2	0.02	0.1	0.0	0.2	0.0	0.0	0.0	0.0
	3	59	1 1	.000	0.00	0.6	0.0	0.4	0.0	0.0	0.0	0.0

Figure 78. Monitor III System Information Report

Indicator

Field: Average Number Delayed For STOR

Description: This shows you the average number of AS delayed for reasons related to processor storage.

Guideline: Look at the value for the user community you are trying to help (for example, a specific service class, all TSO, or overall SYSTEM). If the STOR delay is larger than the other delays listed, you probably have some leverage here. If the STOR delay is zero or near zero, then your leverage is elsewhere (for example, I/O or CPU).

Problem Area: This points out delays for paging, swapping, VIO, and out-ready.

Potential Solution: See the Storage Delays report to find out which AS are delayed, and the type of storage delay they are having. You can get there by positioning the cursor on the storage value you are interested in, and pressing ENTER.

DELAY report

RMF V1R10 Delay Report													Line 1 of 43				
Command ==>													Scroll ==> HALF				
Samples: 100													System: PRD1	Date: 04/07/08	Time: 10.32.00	Range: 100	Sec
Name	Service CX Class	Cr	WFL %	USG %	DLY %	IDL %	UKN %	----- PRC	% Delayed for				----- ENQ	Primary Reason			
BHOLEQB	B BATCH		0	0	51	0	49	0	0	0	0	0	51	SYSDSN			
JES2	S SYSSTC		0	0	2	0	98	0	2	0	0	0	0	SYSPAG			
SMF	S SYSSTC		0	0	1	0	99	0	1	0	0	0	0	SYSPAG			
BAJU	T TSO		29	9	22	67	2	1	0	21	0	0	0	LOCL			
RMFGAT	S SYSSTC		50	1	1	0	98	1	0	0	0	0	0	BAJU			
BHOL	T TSO		78	18	5	73	4	1	3	1	0	0	0	RMFUSR			
BHOLPRO2	B BATCH		93	93	7	0	0	7	0	0	0	0	0	BHOLPRO1			
BHOLPRO1	B BATCH		95	95	5	0	0	5	0	0	0	0	0	BHOLPRO2			
MASTER	S SYSTEM		96	24	1	0	76	0	1	0	0	0	0	SYSPAG			
CATALOG	S SYSSTC		100	2	0	0	98	0	0	0	0	0	0				
BHOLSTOA	B BATCH		100	1	0	99	0	0	0	0	0	0	0				
PCAUTH	S SYSSTC			0	0	0	100	0	0	0	0	0	0				

Figure 79. Monitor III Delay Report

Indicator**Field:** % Delayed for STR**Description:** This will give you the percentage of time the job was delayed for processor storage related reasons.**Guideline:** See if the STR delay is larger than the others, and is significant (over 10%). If so, go to the Storage Delays panel (enter STOR at the command line) and see the guidelines described there.**Problem Area:** Delays due to paging, swapping, VIO, and out-ready.**Potential Solution:** See the "STOR report" on page 136 topic of this information unit.

Storage analysis

STOR report

```
RMF V1R10 Storage Delays                                     Line 1 of 44
Command ==>                                                Scroll ==> HALF

Samples: 100      System: PRD1  Date: 04/07/08  Time: 10.32.00  Range: 100  Sec

Jobname  C  Service  DLY  % Delayed for  Working Set
          C  Class   %    COMM LOCL  VIO  SWAP  OTR  Central Expanded
BAJU     T  TSO      21    0    18    0    3    0    273
BHOL     T  TSO       1    1     0    0    0    0    1111
*MASTER* S  SYSTEM   0     0     0    0    0    0    73
PCAUTH   S  SYSSTC  0     0     0    0    0    0    30
RASP     S  SYSSTC  0     0     0    0    0    0    60
TRACE    S  SYSSTC  0     0     0    0    0    0    35
XCFAS    S  SYSSTC  0     0     0    0    0    0    173
GRS      S  SYSSTC  0     0     0    0    0    0    68
SMXC     S  SYSSTC  0     0     0    0    0    0    20
SYSBMAS  S  SYSSTC  0     0     0    0    0    0    21
DUMPSRV  S  SYSSTC  0     0     0    0    0    0    33
CONSOLE  S  SYSSTC  0     0     0    0    0    0    74
ALLOCAS  S  SYSSTC  0     0     0    0    0    0    26
SMF      S  SYSSTC  0     0     0    0    0    0    47
LLA      S  SYSSTC  0     0     0    0    0    0    213
BHOLST09 B  BATCH   0     0     0    0    0    0     0
```

Figure 80. Monitor III Storage Delays Report

Indicator

Field: Check DLY % first, then find biggest component (e.g. COMM, LOCL...).

Description: This report shows AS in descending order of delay.

Guideline: Check if the value in the DLY % field is higher than 10% before proceeding. The DLY % fields in the various Monitor III reports help you to estimate how much you can reduce response time by eliminating delay. For example, by entirely eliminating 10% delay, you can improve response time by about 10%.

Problem Area: Paging, swapping, VIO, etc.

Potential Solution:

- COMM
 - See “Auxiliary storage tuning” on page 144
- LOCL
 - See “Prioritize access to processor storage” on page 144 - storage isolation and criteria age
 - See “Auxiliary storage tuning” on page 144
 - If swappable, consider swap set size tuning
- VIO
 - See “Prioritize access to processor storage” on page 144 - storage isolation and criteria age
 - See “Auxiliary storage tuning” on page 144
- SWAP
 - See “Prioritize access to processor storage” on page 144 - storage isolation and criteria age
 - See “Auxiliary storage tuning” on page 144
 - Swap set size tuning
- OUTR
 - Check RTO setting

Monitor I indicators

This section will review the most frequently used processor storage indicators in the Monitor I reports.

Paging Activity - Central Storage Paging Rates

PAGING ACTIVITY									
z/OS V1R10		SYSTEM ID OS04			DATE 06/05/2008		PAGE 1		
		RPT VERSION V1R10 RMF			TIME 09.29.00		INTERVAL 14.59.996		
OPT = IEADPT02 MODE=ESAME		CENTRAL STORAGE PAGING RATES - IN PAGES PER SECOND							
CATEGORY	PAGE IN				PAGE OUT				
	SWAP	NON SWAP, BLOCK	NON SWAP, NON BLOCK	% OF TOTL RATE	SWAP	NON SWAP	TOTAL RATE	% OF TOTL SUM	
PAGEABLE SYSTEM AREAS (NON VIO)									
LPA		0.00	0.01	0.01	0				
CSA		0.00	0.04	0.04	0	0.00	0.00	0	
SUM		0.00	0.05	0.05	0	0.00	0.00	0	
ADDRESS SPACES									
HIPERSPACE		0.06		0.06	1	0.00	0.00	0	
VIO		0.00		0.00	0	0.00	0.00	0	
NON VIO	0.28	7.01	2.01	9.30	99	0.00	0.01	0.01	100
SUM	0.28	7.07	2.01	9.37	100	0.00	0.01	0.01	100
TOTAL SYSTEM									
HIPERSPACE		0.06		0.06	1	0.00	0.00	0	
VIO		0.00		0.00	0	0.00	0.00	0	
NON VIO	0.28	7.01	2.06	9.35	99	0.00	0.01	0.01	100
SUM	0.28	7.07	2.06	9.41	100	0.00	0.01	0.01	100
SHARED			0.00	0.00		0.00	0.00		
PAGE MOVEMENT WITHIN CENTRAL STORAGE				20.84	PAGE MOVEMENT TIME % 0.0				
AVERAGE NUMBER OF PAGES PER BLOCK				5.6					
BLOCKS PER SECOND				1.26					
PAGE-IN EVENTS (PAGE FAULT RATE)				3.31					

Figure 81. Paging Activity report - Central Storage Paging Rates

Indicator

Field: PAGE MOVEMENT WITHIN CENTRAL STORAGE

Description: Rate of page movement above and below 16MB in CS. For example, to keep long term fixed pages out of reconfigurable storage.

Guideline: This is rarely a problem. If you see large numbers here (maybe over 100/sec), it may be worth looking at. Or, if PAGE MOVEMENT TIME % reports a high value, this may indicate a problem.

Problem Area: Some CPU time is required to move these pages. Generally not enough to worry about.

Potential Solution: Review the setting of your RSU parameter for reconfigurable storage. Reduce if possible (but not at the expense of your disaster recovery planning).

Indicator**Field:** PAGE-IN EVENTS (PAGE FAULT RATE)**Description:** This is the rate of demand paging in from auxiliary storage. The AS waits while these pages are obtained.**Guideline:** No overall system guideline for this number would make much sense. More important than the single-number here, is the amount of paging delay to important work (see Monitor III Storage Delays report). You may wish to track this one over time, or turn to this for specific MPL tuning.**Problem Area:** Problems here would show up as paging delay to important work.**Potential Solution:** Paging problems are probably best addressed by analyzing which workloads are paging, then addressing the problem for those workloads specifically (see Monitor III Storage Delays report). Use Monitor III to get to the AS level, or the Workload Activity report to get to the service class. There, the following solutions could apply:

- Isolate storage
- Adjust swap set size
- Adjust the criteria age value
- As a last resort, see “Auxiliary storage tuning” on page 144 for ways to speed up the paging.

Storage analysis

Paging Activity - Central Storage Movement Rates / Frame and Slot Counts

PAGING ACTIVITY							PAGE	2
z/OS V1R10		SYSTEM ID TRX1		DATE 11/28/2008		INTERVAL 00.59.999		
		RPT VERSION V1R10 RMF		TIME 10.13.00		CYCLE 1.000 SECONDS		
OPT = IEAOPT00		MODE = ESAME		CENTRAL STORAGE MOVEMENT RATES - IN PAGES PER SECOND				

HIGH UIC (AVG) = 65535 (MAX) = 65535 (MIN) = 65535								
		WRITTEN TO	READ FROM	*---	CENTRAL STORAGE	FRAME COUNTS	----*	
		CENTRAL STOR	CENTRAL STOR	MIN	MAX	AVG		
HIPERSPACE	RT	0.00	0.00	2	2	2		
PAGES								
VIO	RT	0.00	0.00	0	0	0		
PAGES								

FRAME AND SLOT COUNTS								

CENTRAL STORAGE				LOCAL PAGE DATA SET SLOT COUNTS				
		MIN	MAX	AVG		MIN	MAX	AVG

(7 SAMPLES)								
AVAILABLE		744,235	746,444	745,345	AVAILABLE SLOTS	593,999	593,999	593,999
SQA		10,771	10,771	10,771	VIO SLOTS	0	0	0
LPA		4,830	4,830	4,830	NON-VIO SLOTS	0	0	0
CSA		5,804	5,814	5,810	BAD SLOTS	0	0	0
LSQA		27,562	27,571	27,566	TOTAL SLOTS	593,999	593,999	593,999
REGIONS+SWA		220,490	223,193	220,962				
TOTAL FRAMES		1048836	1048836	1048836				
FIXED FRAMES				SHARED FRAMES AND SLOTS				
		MIN	MAX	AVG		MIN	MAX	AVG

NUCLEUS		2,581	2,581	2,581	CENTRAL STORAGE	2,627	2,627	2,627
SQA		9,884	9,884	9,884	FIXED TOTAL	21	21	21
LPA		77	77	77	FIXED BELOW 16 M	0	0	0
CSA		7,370	9,930	9,564	AUXILIARY SLOTS	0	0	0
LSQA		8,472	8,474	8,472	TOTAL	2,627	2,627	2,627
REGIONS+SWA		19,145	21,754	21,009				
BELOW 16 MEG		36	36	36				
BETWEEN 16M-2G		20,333	20,396	20,369				
TOTAL FRAMES		47,592	52,698	51,588				
MEMORY OBJECTS AND FRAMES								
		MIN	MAX	AVG		MIN	MAX	AVG

OBJECTS COMMON		3	5	5				
SHARED		1	1	1				
LARGE		9	11	11				
FRAMES COMMON		512	512	512				
COMMON FIXED		0	0	0				
SHARED		39,442	43,854	41,637				
1 MB		65	85	81				

Figure 82. PAGING Report - Central Storage Movement Rates / Frame and Slot Counts

Indicator**Field:** HIGH UIC (AVG)

Description: Amount of time a page remains in central storage without being referenced. This measures stress on central storage and ranges from a minimum of 0 to a maximum of 65535.

Guideline: Small values might indicate storage constraints, however, some systems run perfectly at low UICs.

Problem Area: Problems here would show up as delays caused by demand paging or swapping, or increased CPU overhead due to page movement. Also throughput of the system could drop, due to MPL adjustment.

Potential Solution: Decrease storage demand.

Page/Swap Data Set Activity report

PAGE / SWAP DATASET ACTIVITY													
PAGE 1													
z/OS VIRI0			SYSTEM ID OS04		DATE 06/05/2008		INTERVAL 14.59.996						
			RPT VERSION VIRI0 RMF		TIME 09.29.00		CYCLE 1.000 SECONDS						
NUMBER OF SAMPLES = 898			PAGE DATA SET USAGE										
PAGE SPACE TYPE	VOLUME SERIAL	DEV NUM	DEVICE TYPE	SLOTS	SLOTS USED		BAD SLOTS	% IN USE		PAGE TRANS TIME	NUMBER IO REQ	PAGES XFER'D	V I DATA SET NAME
				ALLOC	MIN	MAX		AVG					
PLPA	PAGP11	0378	33903	45000	16466	16466	16466	0	0.11	0.200	6	5	PAGE.OS04.APLPA
COMMON	PAGP10	0377	33903	21240	5345	5378	5355	0	0.11	0.027	38	37	PAGE.OS04.ACOMMON
LOCAL	PAGP12	0379	33903	566820	67670	70965	69140	0	2.67	0.004	711	5,697	N PAGE.OS04.LOCAL01
LOCAL	PAGP13	037A	33903	566820	65887	69007	67190	0	3.01	0.005	771	5,711	N PAGE.OS04.LOCAL02
LOCAL	PAGP14	037B	33903	566820	66611	69723	67928	0	3.01	0.005	821	5,625	Y PAGE.OS04.LOCAL03
LOCAL	PAGP15	038D	33903	566820	66296	69563	67618	0	3.23	0.005	726	5,523	Y PAGE.OS04.LOCAL04
LOCAL	PAGP16	039F	33903	566820	66390	69311	67560	0	3.12	0.005	766	5,548	Y PAGE.OS04.LOCAL05

Figure 83. Monitor I Page/Swap Data Set Activity Report

Note: Starting with OS/390 2.10, swap data sets are no longer supported. Therefore, the report contains information only about page data sets.

Indicator**Field:** SLOTS ALLOC and AVG SLOTS USED

Description: This tells you how much of your auxiliary storage space is used.

Guideline: Keep the percent of space used to around 30% (always below 50%) of the space allocated.

Storage analysis

Indicator

Field: % IN USE

Description: This is the busy percentage for the data set.

Guideline: For page data sets, you may see response time increase if this number rises above 30%. If you have swap data sets defined, you may see response time increases if they are above 15%.

Problem Area: Increased time to satisfy a page request from DASD.

Potential Solution: See "Auxiliary storage tuning" on page 144

DASD Activity 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	2		
z/OS V1R10			SYSTEM ID SYSF		DATE 05/27/2008			INTERVAL 14.59.998												
			RPT VERSION V1R10 RMF		TIME 06.00.00			CYCLE 1.000 SECONDS												
TOTAL SAMPLES =		900	IODF = 00		CR-DATE: 10/16/2008		CR-TIME: 10.31.31		ACT: ACTIVATE											
STORAGE GROUP	DEV NUM	DEVICE TYPE	NUMBER OF CYL	VOLUME SERIAL	PAV	LCU	DEVICE ACTIVITY RATE	AVG RESP TIME	AVG IOSQ	AVG CMR DLY	AVG DB DLY	AVG PEND TIME	AVG DISC TIME	AVG CONN TIME	% DEV CONN	% DEV UTIL	% DEV RESV	AVG NUMBER ALLOC	% ANY ALLOC	% MT PEND
	732A	33903	200	SYSXCA		002C	0.037	4.87	.000	.178	.000	.287	.000	4.58	0.02	0.02	0.0	0.0	100.0	0.0
	733A	33903	200	SYSAXA		002C	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0	0.0
	733D	33903	200	SYSDXA		002C	1.354	1.86	.000	.133	.000	.274	.052	1.54	0.21	0.22	0.0	14.0	100.0	0.0
				LCU		002C	1.391	1.94	.000	.134	.000	.275	.051	1.62	0.07	0.08	0.0	14.0	100.0	0.0
	9690	33909	10017	SYSLIB	1	0040	1.281	1.71	.867	.154	.038	.284	.005	.549	0.07	0.07	0.1	38.0	100.0	0.0
	9691	33909	10017	SYSUSR	1	0040	1.707	.536	.000	.133	.037	.239	.000	.297	0.05	0.05	0.0	4.0	100.0	0.0
	9692	33909	10017	SYSOPE	1	0040	0.041	2.22	.000	.135	.000	.214	.000	2.01	0.01	0.01	0.0	3.0	100.0	0.0
	9693	33909	10017	SYSCDS	1	0040	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0	0.0
	9694	33909	10017	SYSHM1	1	0040	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0	0.0
SMS	9695	33909	10017	SYSSMS	1	0040	0.012	.593	.000	.163	.000	.256	.000	.337	0.00	0.00	0.0	1.0	100.0	0.0
SMS	9696	33909	10017	SYSST1	1	0040	0.001	1.79	.000	.512	.000	.640	.000	1.15	0.00	0.00	0.0	0.0	100.0	0.0
SMS	9697	33909	10017	SYSST2	1	0040	0.001	2.94	.000	1.41	.000	1.54	.000	1.41	0.00	0.00	0.0	0.0	100.0	0.0
	9699	33909	10017	SYSHM2	1	0040	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0	0.0
DB2	969A	33909	10017	SYSSD1	1	0040	0.001	1.66	.000	.384	.000	.512	.000	1.15	0.00	0.00	0.0	0.0	100.0	0.0
DB2	969B	33909	10017	SYSSD2	1	0040	0.001	1.02	.000	.256	.000	.384	.000	.640	0.00	0.00	0.0	0.0	100.0	0.0
DB2	969C	33909	10017	SYSSD3	1	0040	0.001	2.05	.000	.640	.000	.768	.000	1.28	0.00	0.00	0.0	0.0	100.0	0.0
	969E	33909	10017	SYSDB2	1	0040	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0	0.0
	969F	33909	10017	SYSIMS	1	0040	0.000	.000	.000	.000	.000	.000	.000	.000	0.00	0.00	0.0	0.0	100.0	0.0
				LCU		0040	3.047	1.05	.365	.143	.036	.259	.002	.428	0.01	0.01	0.0	46.0	100.0	0.0

Figure 84. Direct Access Device Activity Report

Indicator

Only DASD volumes used for paging and swapping are of interest in this section. See Chapter 4, "Analyzing I/O activity" for analysis of other I/O.

Field: AVG RESP TIME

Description: This will show overall response time (ms) for these volumes.

Related Field: DEVICE ACTIVITY RATE, Monitor III page/swap delay % values.

Guideline: Typical response times for page/swap volumes are in the 30-40 ms range. If your response times are significantly higher you may want to investigate.

Problem Area: Problems here would add page or swap delay to your important workloads.

Potential Solution: See "Auxiliary storage tuning" on page 144

General storage recommendations

This section gives actions you may want to consider to improve your processor storage performance. Before you start, bear in mind the following;

- The first priority of tuning is to make the best use of your resources.
- After this, prioritizing workload is even more critical as you will be stealing from one workload to give to another.

Increase storage

- Take central storage from another partition, or install more.

Auxiliary storage tuning

- Reduce paging (see above).
- Dedicate volumes to page data sets.
- Distribute page data sets among channel paths and control units.
- Limit use of VIO=YES.
- Add more LOCALs. Make the sum of all the page space 2 to 4 times the number of slots used.
- Use faster devices.

Prioritize access to processor storage

- Make better use of the PLPA by removing unused modules and verify that concatenated libraries are needed (LPALSTxx).
- Decrease the page fixing level. RMF tells you the cause of the problem at PRIV-FF in the ARD Monitor II report.
- Use the storage protection specification in the service policy.

Tune DASD

- Tuning 'normal' DASD I/O reduces storage requirements. Transactions will complete faster, needing less virtual and processor storage.

Summary

To analyze storage problems you should:

- View the system as a whole and not rely on individual guidelines.

As we have said before the main thing is: Are your SLA targets being met? If not, and storage is the problem then:

- Ensure the storage is being used fully and start looking at prioritizing access.
- Use the recommendation checklist with the *z/OS MVS Initialization and Tuning Reference*.
- Only change one value at a time and monitor its effect before making any other change.

Chapter 6. Analyzing sysplex activity

Let's Analyze Your Sysplex

This information unit discusses some special performance aspects in a sysplex:

- How to understand CICS and IMS workload reports
- Analyzing coupling facility activities

Understanding Workload Activity data for IMS or CICS

This information unit provides an explanation of some Workload Activity report data for work managers. RMF provides data for the subsystem work managers that support workload management. In z/OS, these are IMS and CICS. You find these values either in the Monitor III Work Manager Delays report or in the Postprocessor Workload Activity report. The following examples will describe the Workload Activity report.

This section provides some sample reports for CICS and IMS, followed by some potential problem data and explanations. Based on the explanations, you may decide to alter your service class definitions. In some cases, there may be some actions that you can take. In those cases, you can follow the suggestion. In other cases, the explanations are provided only to help you better understand the data.

Interpreting the Workload Activity report

The following samples of a Workload Activity report with CICS transactions are discussed:

- “CICS service class accessing DBCTL”
- “CICS service class accessing DBCTL with IMS V5” on page 148
- “CICS accessing data via FORs” on page 149

CICS service class accessing DBCTL

Figure 85 shows an example of the work manager state section for a service class representing CICS transactions accessing Database Control (DBCTL).

REPORT BY: POLICY=HPTSPOLI WORKLOAD=PRODWKLD SERVICE CLASS=CICSHR RESOURCE GROUP=*NONE PERIOD=1 IMPORTANCE=1										
TRANSACTIONS	TRANS.-TIME	HHH.MM.SS.TTT								
AVG	0.00	ACTUAL			114					
MPL	0.00	EXECUTION			78					
ENDED	216	QUEUED			36					
END/S	0.24	R/S AFFINITY			0					
#SWAPS	0	INELIGIBLE			0					
EXCTD	216	CONVERSION			0					
AVG ENC	0.00	STD DEV			270					
REM ENC	0.00									
MS ENC	0.00									
----- STATE SAMPLES BREAKDOWN (%) -----										
SUB	P	RESP	--ACTIVE--				-----WAITING FOR-----			-----STATE-----
		TIME	READY	IDLE						SWITCHED SAMPL(%)
		(%)	SUB	APPL	CONV	PROD				LOCAL
CICS	BTE	93.4	10.9	0.0	0.0	0.0	89.2	0.0	89.2 0.0 0.0	
CICS	EXE	67.0	19.7	0.0	10.6	0.0	0.0	69.7	0.0 0.0 0.0	

Figure 85. Hotel Reservations Service Class

The fields in this report describe the CICSHR service class. CICS transactions have two phases:

- The *begin-to-end phase* (CICS BTE) takes place in the first CICS region to begin processing a transaction. Usually this is the terminal owning region (TOR). The TOR is responsible for starting and ending the transaction.
 - The **ENDED** field shows that 216 hotel reservation transactions completed.
 - The **ACTUAL** time shows that the 216 transactions completed in an average transaction time of 0.114 seconds.
- The *execution phase* (CICS EXE) can take place in an application owning region (AOR) and a file owning region (FOR). In this example, the 216 transactions were routed by a TOR to an AOR.

- The **EXCTD** field shows that the AORs completed 216 transactions in the interval.
- The **EXECUTION** time shows that on average it took 0.078 seconds for the AORs to execute the 216 transactions. The **EXECUTION** time applies only to the **EXCTD** transactions.

While executing these transactions, the CICS subsystem records the states the transactions are experiencing. RMF reports the states in the STATE SAMPLES BREAKDOWN (%) section of the report. Because there is a CICS BTE and a CICS EXE field, you can assume that the time spent in the TOR represents the BTE phase, and the time spent in the AOR represents the EXE phase. There is one EXE phase summarizing all the time spent in one or more AORs. Figure 86 shows the actual response time break down of the CICSHR service class.

Response Time = 0.114 seconds = 100%

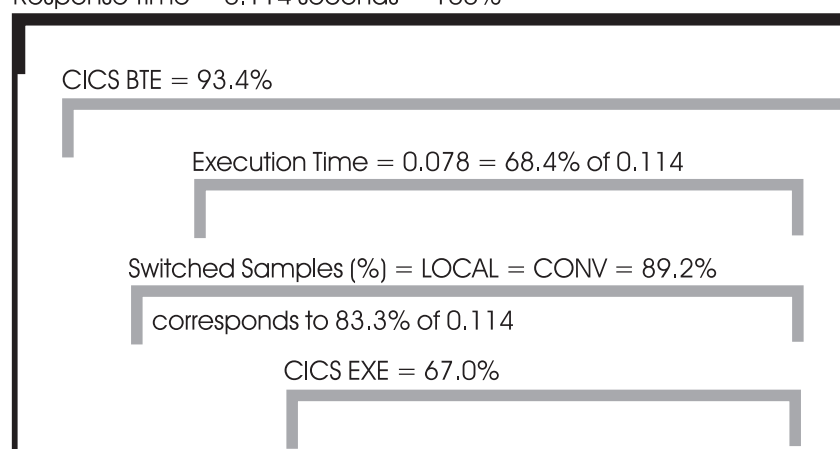


Figure 86. Response Time Breakdown of CICSHR accessing DBCTL

The CICS BTE field shows that the TORs have information covering 93.4% of the response time. RMF does not have information covering 100% of the 0.114 seconds response time, because it takes some time for the system to recognize and assign incoming work to a service class before it can collect information about it.

For most of the 93.4% of the time, the transactions did not run in the TOR, but had been routed locally to an AOR on the same MVS image. You can see this by the SWITCHED SAMPL (%) LOCAL field, which is 89.2% of the total state samples. This value accounts for 83.3% of the response time, because 100% of the total state samples correspond to 93.4% of the response time ($89.2 \times 93.4 / 100 = 83.3\%$). This value of 89.2% is close, if not equal, to the WAITING FOR CONV field, which indicates that there is no delay in the TOR once the AOR has returned the transactions.

The total execution time is some percentage of the total response time. It is the EXECUTION transaction time (0.078), divided by ACTUAL transaction time (0.114), which is 68.4%. The CICS execution phase (CICS EXE field) covers 67% of the response time. Some of that time the work is active in the AOR, sometimes it is waiting behind another task in the region, but 69.7% of the total state samples in the PROD field (which corresponds to $69.7 \times 67 / 100 = 46.7\%$ of the response time) were found outside of the CICS subsystem, waiting for another product to provide some service to these transactions. Based on the configuration of the system, the transactions are accessing DBCTL.

Understanding WLM

The LOCL, SYSP and REMT state percentages appear in the WAITING FOR section if greater than zero and show the percentages of the total state samples the service class was delayed in these states when CICS was waiting to establish a session. The STATE SWITCHED SAMPL(%) fields LOCAL, SYSPL, and REMOT show the percentage of the state samples in which transactions were routed via MRO, MRO/XCP, or VTAM connections.

CICS service class accessing DBCTL with IMS V5

Figure 87 shows an example of a Workload Activity report for a CICS service class representing point-of-sale transactions (CICSPS). This example builds on the example in Figure 85 on page 146. It shows how the RMF report changes when IMS is upgraded to Version 5.

REPORT BY: POLICY=WSTPOL02		WORKLOAD=CICS		SERVICE CLASS=CICSPS		RESOURCE GROUP=*NONE		PERIOD=1 IMPORTANCE=1					
TRANSACTIONS	TRANS.-TIME	HHH.MM.SS.TTT											
AVG	0.00	ACTUAL	503										
MPL	0.00	EXECUTION	399										
ENDED	579	QUEUED	104										
END/S	0.95	R/S AFFINITY	0										
#SWAPS	0	INELIGIBLE	0										
EXCTD	559	CONVERSION	0										
AVG ENC	0.00	STD DEV	847										
REM ENC	0.00												
MS ENC	0.00												
----- STATE SAMPLES BREAKDOWN (%) -----													
SUB	P	RESP	--ACTIVE--				-----WAITING FOR-----			-----STATE-----			
TYPE		TIME (%)	SUB	APPL	READY	IDLE	CONV	PROD	I/O	LOCK	LOCAL	SYSPL	REMT
CICS	BTE	96.6	2.5	0.0	0.1	0.0	97.4	0.0	0.0	0.0	97.4	0.0	0.0
CICS	EXE	29.2	48.6	0.0	14.0	0.0	0.0	23.6	13.7	0.0	0.0	0.0	0.0
IMS	EXE	43.7	92.2	0.0	0.0	0.0	0.0	0.0	0.0	7.8	0.0	0.0	0.0

Figure 87. Point-of-Sale Service Class

The ENDED field shows that 579 point-of-sale transactions completed. The ACTUAL time shows that those transactions completed in an average response time of 0.503 seconds.

The EXCTD field shows that the AORs completed 559 transactions in the interval. This is less than the ENDED value because a TOR may not always route all transactions to an AOR. Those non-routed transaction are not counted as EXCTD. The EXECUTION field shows that on average it took the AORs 0.399 seconds to execute the 559 transactions.

While executing these transactions, the CICS subsystem reports the states the transactions are experiencing. Figure 88 on page 149 shows the response time breakdown of the CICSPS service class.

Response Time = 0.503 seconds = 100%




Figure 88. Response Time Breakdown of CICSPS accessing DBCTL with IMS V5

The CICS BTE field shows that the TORs have information covering 96.6% of the response time. Most of that time, the transactions were in fact not being run in the TOR, but had been routed locally to an AOR on the same MVS image. You can see this by the SWITCHED SAMPL(%) LOCAL field, which is 97.4%. This value accounts for 94.1% of the response time ($97.4 \times 96.6 / 100 = 94.1\%$)

The EXECUTION transaction time (0.399 seconds) divided by the ACTUAL transaction time (0.503 seconds) is 79.32%. There are two execution phases shown: a CICS EXE, and an IMS EXE. The CICS execution phase covers only 29.2% of the response time. That is lower than in the previous example, because now IMS is providing information. The IMS EXE row shows the 43.7% of the response time for which it is responsible.

CICS accessing data via FORs

Figure 89 shows an example of a Workload Activity report for some CICS user transactions which have data accessed via file owning regions.

REPORT BY: POLICY=EQUALIMP		WORKLOAD=CICS		SERVICE CLASS=CICUSRTX		RESOURCE GROUP=*NONE		PERIOD=1 IMPORTANCE=1																																																	
TRANSACTIONS	TRANS.-TIME	HHH.MM.SS.TTT																																																							
AVG	0.00	ACTUAL		93																																																					
MPL	0.00	EXECUTION		69																																																					
ENDED	14119	QUEUED		24																																																					
END/S	47.17	R/S AFFINITY		0																																																					
#SWAPS	0	INELIGIBLE		0																																																					
EXCTD	13067	CONVERSION		0																																																					
AVG ENC	0.00	STD DEV		63																																																					
REM ENC	0.00																																																								
MS ENC	0.00																																																								
<table border="1"> <thead> <tr> <th rowspan="2">SUB TYPE</th> <th rowspan="2">P</th> <th rowspan="2">RESP TIME (%)</th> <th colspan="6">STATE SAMPLES BREAKDOWN (%)</th> <th colspan="3">-----STATE-----</th> </tr> <tr> <th>--ACTIVE--</th> <th>READY</th> <th>IDLE</th> <th colspan="3">I/O CONV LOCK</th> <th>SWITCHED SAMPL(%)</th> <th>LOCAL</th> <th>SYSPL</th> <th>REMOT</th> </tr> </thead> <tbody> <tr> <td>CICS</td> <td>BTE</td> <td>90.5</td> <td>3.0</td> <td>0.0</td> <td>3.9</td> <td>0.2</td> <td>4.0</td> <td>88.8</td> <td>0.0</td> <td>92.4</td> <td>0.0</td> <td>0.0</td> </tr> <tr> <td>CICS</td> <td>EXE</td> <td>67.2</td> <td>7.0</td> <td>0.0</td> <td>3.9</td> <td>0.0</td> <td>89.0</td> <td>0.0</td> <td>0.1</td> <td>60.9</td> <td>0.0</td> <td>0.0</td> </tr> </tbody> </table>										SUB TYPE	P	RESP TIME (%)	STATE SAMPLES BREAKDOWN (%)						-----STATE-----			--ACTIVE--	READY	IDLE	I/O CONV LOCK			SWITCHED SAMPL(%)	LOCAL	SYSPL	REMOT	CICS	BTE	90.5	3.0	0.0	3.9	0.2	4.0	88.8	0.0	92.4	0.0	0.0	CICS	EXE	67.2	7.0	0.0	3.9	0.0	89.0	0.0	0.1	60.9	0.0	0.0
SUB TYPE	P	RESP TIME (%)	STATE SAMPLES BREAKDOWN (%)						-----STATE-----																																																
			--ACTIVE--	READY	IDLE	I/O CONV LOCK			SWITCHED SAMPL(%)	LOCAL	SYSPL	REMOT																																													
CICS	BTE	90.5	3.0	0.0	3.9	0.2	4.0	88.8	0.0	92.4	0.0	0.0																																													
CICS	EXE	67.2	7.0	0.0	3.9	0.0	89.0	0.0	0.1	60.9	0.0	0.0																																													

Figure 89. CICS User Transactions Service Class

In this example, 14,119 transactions completed with a response time of 0.093 seconds. Only 13,067 transactions were executed by AORs in the interval. Others ran completely in the TOR.

Understanding WLM

The CICS execution phase contains information about the transactions while they were running in either the AOR or FOR. This shows that 89.0% of the state samples for these transactions were waiting for I/O completion. This accounts for 59.8% of the response time ($67.2 \times 89.0 / 100 = 59.8\%$). It is not possible to describe whether this time was for I/O initiated by the FOR, or for I/O initiated within the AOR. However, the SWITCHED SAMPL(%) LOCAL value in the CICS EXE line says that for 60.9% of the state samples (which is 40.9% of the response time, that is, $60.9 \times 67.2 / 100$), one of those regions was waiting for the other region to complete some processing for the transaction before the original region could proceed.

Problem: Very large response time percentage

Figure 90 shows an example of a work manager state section for the CICSPROD service class. In column RESP TIME (%), both the CICS EXE and the CICS BTE rows show inflated percentages: 78.8K and 140.

REPORT BY: POLICY=HPTSPOL1 WORKLOAD=PRODWKLD SERVICE CLASS=CICSPROD RESOURCE GROUP=*NONE PERIOD=1 IMPORTANCE=1													
TRANSACTIONS	TRANS.-TIME	HHH.MM.SS.TTT											
AVG	0.00	ACTUAL	111										
MPL	0.00	EXECUTION	123										
ENDED	1648	QUEUED	0										
END/S	1.83	R/S AFFINITY	0										
#SWAPS	0	INELIGIBLE	0										
EXCTD	1009	CONVERSION	0										
AVG ENC	0.00	STD DEV	351										
REM ENC	0.00												
MS ENC	0.00												
----- STATE SAMPLES BREAKDOWN (%) -----													
SUB	P	RESP	--ACTIVE--		READY	IDLE	-----WAITING FOR-----			-----STATE-----			
TYPE		TIME	SUB	APPL			MISC	PROD	CONV	I/O	SWITCHED SAMPL(%)		
		(%)									LOCAL	SYSPL	REMOT
CICS	BTE	78.8K	0.2	0.0	0.3	2.5	96.7	0.0	0.3	0.0	0.3	0.0	0.0
CICS	EXE	140	65.6	0.0	2.2	0.0	0.0	32.4	0.0	0.1	0.0	0.0	0.0

Figure 90. Response Time Percentages Greater than 100

Possible Explanations

Long-running transactions: The report above shows how long-running transactions can inflate the value for RESP TIME (%). While the following example does not explain the exact values in the figure, it explains why this is possible.

Suppose 100 transactions have ended within 1 second, and one transaction has been running for 5 minutes and is still executing when the RMF interval expires. The ACTUAL transaction time shows an average response time of 1 second, and RMF shows the breakdown into the states recorded by CICS or IMS. The subsystem, however, recorded a total of 6 minutes and 40 seconds (5 minutes plus 100 seconds) worth of data. That is an average of 4 seconds worth of data for each completed transaction — which is 4 times the 1 second response time. The state samples breakdown, however, shows information representing 100% of the state samples.

Also, when the one long-running transaction completes, it could easily distort the average response time during that interval. The RMF standard deviation and distribution of response times emphasizes when this occurs.

The long-running transactions could be either routed or non-routed transactions. Routed transactions are transactions that are routed from a TOR to any AOR. Long-running routed transactions could result in many samples of WAITING FOR CONV (waiting for a conversation) in the CICS BTE phase, as well as states recorded from the AOR in the execution phase.

Long-running non-routed transactions execute completely in a TOR, and have no CICS EXE phase data, and could inflate any of the state data for the CICS BTE phase.

Never-ending transactions: Never-ending transactions differ from long-running transactions in that they persist for the life of a region. For CICS, these could include the IBM reserved transactions such as CSNC and CSSY, or customer defined transactions. Never ending transactions are reported similarly to long-running transactions explained in “Long-running transactions” on page 150. However, for never-ending CICS transactions, RMF might report high percentages in the IDLE, WAITING FOR TIME, or the WAITING FOR MISC (miscellaneous) fields.

Conversational transactions: Conversational transactions are considered long-running transactions. CICS marks the state of a conversational transaction as IDLE when the transaction is waiting for terminal input. Terminal input often includes long end-user response time, so you might see percentages close to 100% in the IDLE state for completed transactions.

Service class includes dissimilar work: A service class that mixes customer and IBM transactions, short and long or never-ending transactions, routed and non-routed transactions, or conversational and non-conversational transactions can expect to have RMF reports showing that the total states sampled account for more than the average response time. This could be true for both IMS and CICS, and can be expected if the service class is the subsystem default service class. The default service class is defined in the classification rules. It is the service class to which all work in a subsystem is assigned that is not assigned to any other service class.

Possible actions

Group similar work into service classes: Make sure your service classes represent a group of similar work. This could require creating additional service classes. For the sake of simplicity, you can have only a small number of service classes for CICS or IMS work. If there are transactions for which you want the RMF state samples breakdown data, consider including them in their own service class.

Do nothing: For service classes representing dissimilar work such as the subsystem default service class, understand that the response time percentage could include long-running or never-ending transactions. RMF data for such service classes may not make immediate sense.

Problem: Response time is zero

Figure 91 on page 152 shows an example of a work manager state section for the CICSLONG service class. The RESP TIME (%) fields shows a 0.0 value.

Understanding WLM

```

REPORT BY: POLICY=HPTSPOL1  WORKLOAD=PRODWKLD  SERVICE CLASS=CICSLONG  RESOURCE GROUP=*NONE  PERIOD=1  IMPORTANCE=1

                                CICS Long Running Internal Trxs

TRANSACTIONS  TRANS.-TIME  HHH.MM.SS.TTT.SS.TTT
AVG          0.00  ACTUAL                0
MPL          0.00  EXECUTION             0
ENDED        0     QUEUED                0
END/S        0.00  R/S AFFINITY          0
#SWAPS       0     INELIGIBLE            0
EXCTD        0     CONVERSION            0
AVG ENC      0.00  STD DEV               0
REM ENC      0.00
MS ENC       0.00

SUB   P   RESP  ----- STATE SAMPLES BREAKDOWN (%) -----  -----STATE-----
TYPE  TIME  --ACTIVE--  READY  IDLE  -----WAITING FOR-----  SWITCHED SAMPL(%)
      (%)  SUB  APPL  CONV  I/O  PROD  DIST  REMT  LOCK  LOCAL  SYSPL  REMOT
CICS  BTE  0.0  70.8  0.0  1.4  0.7  11.2  9.2  0.3  5.3  1.2  0.0  11.2  0.0  0.0
CICS  EXE  0.0  43.2  0.0  0.2  0.1  31.8  10.4  8.7  0.0  2.9  2.8  0.0  0.0  0.0

```

Figure 91. Response Time Percentages all Zero

Possible explanations

No transactions completed: While a long-running or never-ending transaction is being processed, RMF stores the service class state samples in SMF 72.3 records. But when no transactions have completed, the average response time is 0. However, the calculations for the state samples breakdown will result in values greater than 0.

RMF did not receive data from all systems in the sysplex: The Postprocessor may have been given SMF records from only a subset of the systems running in the sysplex. The report may represent only a single MVS image. If that MVS image has no TOR, its AORs receive CICS transactions routed from another MVS image or from outside the sysplex. Since the response time for the transactions is reported by the TOR, there is no transaction response time for the work, nor are there any ended transactions, on this MVS image.

Possible actions

Do nothing: You may have created this service class to prevent the state samples of long-running transactions from distorting data for your production work.

Combine all SMF records for the sysplex: When a single MVS image that does not have TORs is combined with another MVS image that does have TORs and therefore does report response times, the states and response time from the first image will be combined by RMF with the states and response time from the second.

Problem: More executed transactions than ended transactions

In an IMS TM environment, the control region counts the number of ended transactions, and provides the actual response time information. The message processing program (MPP) counts the number of executed transactions, and provides the execution time. IMS provides only EXE phase information. There is no BTE phase information for IMS work.

Figure 92 on page 153 shows an example of a work manager data section for the IMSTRX service class. The EXCTD field shows that 5312 transactions have finished executing while the ENDED field shows that only 5212 have been ended by the CTL

region.

TRANSACTIONS		TRANS.-TIME		HHH.MM.SS.TTT	
AVG	0.00	ACTUAL			133
MPL	0.00	EXECUTION			107
ENDED	5212	QUEUED			26
END/S	5.79	R/S AFFINITY			0
#SWAPS	0	INELIGIBLE			0
EXCTD	5312	CONVERSION			0
AVG ENC	0.00	STD DEV			512
REM ENC	0.00				
MS ENC	0.00				

SUB P		RESP	STATE SAMPLES BREAKDOWN (%)					STATE			
TYPE		TIME (%)	---ACTIVE---	READY	IDLE	---	WAITING FOR	---	SWITCHED	SAMPL (%)	
			SUB	APPL	LOCK				LOCAL	SYSPL	REMO
IMS	EXE	82.4	98.1	0.0	0.0	0.0	1.9		0.0	0.0	0.0

Figure 92. Executed Transactions greater than Ended Transactions

Possible explanations

IMS program-to-program switches: When an IMS transaction makes a program-to-program switch, the switched work can be considered either a new transaction, or a continuation of the originating transaction. The system checks the classification rules for the service class to be associated with the switched work.

If the service class is the same as the originating transaction, it is processed in the MPP, and it is counted as an executed transaction (EXCTD). When the response is sent back to the network, RMF shows two executed transactions: the originating and the switched work, but only a single ended transaction (ENDED). Here, the number of executed transactions exceeds the number of ended transactions.

If the service class is different from the original transaction, then RMF counts it as a new transaction with its own EXCTD and ENDED. Here, the executed transaction value and ended value agree.

RMF processed only part of the entire sysplex data: It is possible to invoke the Postprocessor, giving it SMF records from just part of a sysplex — for example a single MVS image. Suppose that MVS image has no TOR, but its AORs just receive CICS transactions routed from another MVS image. Since the transaction ends in the TOR, but is executed in the AOR, there will be no ended transactions for the CICS service class from this MVS image.

Snapshot data is not always complete: The Workload Activity report contains "snapshot data", which is data collected over a defined interval. For IMS, in a given interval, several transactions may have already executed in an MPP, but the control region may not have reported their completion yet. Similarly for CICS, the AOR could have finished executing a transaction, but the TOR may not have completed processing and reporting the completion. Thus, RMF may show more executed transactions than completions.

Possible actions

Classify IMS transactions uniquely: When an IMS transaction makes a program-to-program switch, the system checks the classification rules for the service class to be associated with the transaction. You can classify the switched transactions to a service class different from the service class of the originating transaction. RMF then counts a new transaction, with its own execution and end.

Understanding WLM

This way, the execution and ended values agree, and the RMF data is consistent with the data reported in the IMS Performance Analysis and Reporting System (IMSPARS). IMSPARS includes transaction timings, analysis reports, and detailed reports tracing individual transaction and database change activity. Classifying the program-to-program switched transaction to a unique service class might be useful for asynchronous program-to-program switches.

Combine all SMF records for the sysplex: When a single MVS image that does not have TORs is combined with another MVS image that does, the executed transactions from the first image will be combined by RMF with the ended transactions from the second so that the sysplex-wide report is consistent.

Problem: Execution time is greater than response time

Figure 93 shows an example of a work manager state section for the CICSPROD service class. In the example, there is a response time of .091 seconds, but an execution time of .113 seconds. The example also shows 1731 ENDED transactions, yet the EXCTD field shows that only 1086 have been executed.

```
REPORT BY: POLICY=HPTSPOL1  WORKLOAD=PRODWKLD  SERVICE CLASS=CICSPROD  RESOURCE GROUP=*NONE  PERIOD=1 IMPORTANCE=1

TRANSACTIONS  TRANS.-TIME  HHH.MM.SS.TTT
AVG          0.00  ACTUAL                91
MPL          0.00  EXECUTION             113
ENDED       1731  QUEUED                0
END/S        1.92  R/S AFFINITY         0
#SWAPS        0   INELIGIBLE           0
EXCTD       1086  CONVERSION            0
AVG ENC      0.00  STD DEV               92
REM ENC      0.00
MS ENC       0.00
```

Figure 93. Execution Time Greater than Response Time

Possible explanation

Mixing routed and non-routed CICS transactions: The AORs may have recorded states which account for more time than the average response time of all the transactions. The non-routed transactions do not show up in the EXE phase. In addition, most non-routed transactions end very quickly, and decrease the actual response time. The response time (ACTUAL field) shows 0.091 seconds as the average of all 1731 transactions, while the AORs can only describe the execution of the 1086 transactions they participated in.

Possible actions

Classify routed and non-routed transactions to different service classes: This would keep the numbers consistent with the expectation.

Problem: Large SWITCH percentage in CICS execution phase

Figure 94 on page 155 shows a work manager state data section for a CICSPROD service class. The LOCAL value in the SWITCHED SAMPL(%) section shows a value of 7092.

```

REPORT BY: POLICY=HPTSPOL1  WORKLOAD=PRODWKLD  SERVICE CLASS=CICSPROD  RESOURCE GROUP=*NONE  PERIOD=1 IMPORTANCE=1

TRANSACTIONS  TRANS.-TIME  HHH.MM.SS.TTT
AVG  0.00  ACTUAL  150
MPL  0.00  EXECUTION  134
ENDED  3599  QUEUED  16
END/S  4.00  R/S AFFINITY  0
#SWAPS  0  INELIGIBLE  0
EXCTD  2961  CONVERSION  0
AVG ENC  0.00  STD DEV  446
REM ENC  0.00
MS ENC  0.00

      RESP ----- STATE SAMPLES BREAKDOWN (%) -----
SUB  P  TIME  --ACTIVE--  READY  IDLE  -----WAITING FOR-----  -----STATE-----
TYPE (%)  SUB  APPL  MISC  PROD  CONV  I/O  LOCAL  SYSPL  REMOT
CICS BTE 26.8K 0.3  0.0  0.4  2.5  96.3  0.0  0.6  0.0  0.6  0.0  0.0
CICS EXE 93.7 41.2  0.0  6.0  0.0  0.0  52.7  0.0  0.1  7092  0.0  0.0
    
```

Figure 94. Large SWITCH Percentage in a CICS Execution Environment

Possible explanations

Distributed transaction processing: If a program initiates distributed transaction processing to multiple back-end sessions, there can be many AORs all associated with the original transaction. Each of these multiple back-end regions can indicate they are switching control back to the front-end region (SWITCH to another region on the LOCAL MVS image, or to a region on another MVS image in the sysplex). Thus, with a one-to-many mapping like this, there are many samples of the execution phase of requests switched — long enough to exceed 100% of the response time of other work completing in the service class.

Distributed program link (DPL): The distributed program link function from CICS/ESA® 3.3 builds on the distributed transaction functions available in CICS by enabling a CICS program (the client program) to call another program (the server program) in another CICS region.

While the server program is running, the client program will reflect that it is switched to another CICS region.

Possible Actions

- None -

Problem: Decreased number of ended transaction with increased response times

The Workload Activity report shows increased response times, and a decrease in the number of ended transactions over a few days.

Possible explanation

Conversion from ISC link to MRO: When two CICS regions are connected via a VTAM inter-system communication (ISC) link, they behave differently than when they are connected via multi-region (MRO) option. One key difference is that, with ISC, both the TOR and the AOR are receiving a request from VTAM, so each believes it is starting and ending a given transaction.

So for a given user request routed from the TOR via ISC to an AOR, there would be 2 completed transactions. Let us assume they have response times of 1 second and 0.75 seconds, resulting in an average of 0.875 seconds.

Understanding WLM

When the TOR routes via MRO, the TOR will describe a single completed transaction taking 1 second, and the AOR will report its 0.75 seconds as execution time. Therefore, converting from an ISC link to an MRO connection, for the same workload, as shown in this example, could result in half the number of ended transactions and an increase in the response time reported by RMF. The difference could be much more significant if the AOR to FOR link was converted.

Migration of an AOR to CICS 4.1: CICS 4.1 has extended the information it passes between regions involved in processing a given transaction. If all the regions are at the 4.1 level, this allows RMF to report the number of ended transactions as just those that have completed in the TOR. But if there is a mixture of release levels involved, this is not guaranteed.

Let us assume that the TOR and FOR have been upgraded to release 4.1, but the AOR has not yet been upgraded. When the TOR receives a transaction, it will determine the service class. Then when it calls the AOR, it will pass that service class information.

But since the AOR is not yet at a 4.1 level, the AOR has no code that passes this service class along to the FOR. So when the FOR is invoked, it believes it is starting a new transaction and must classify the request to a service class. When the FOR finishes processing, it states that it has completed the transaction. Later the TOR also states that it has completed the transaction. This results in RMF showing multiple completions for a given end-user transaction (one completion from the TOR, and one completion for each invocation of the FOR), with an average response time less than the value the TOR alone is reporting.

When the AOR is migrated to CICS 4.1, it will pass the service class to the FOR. The FOR recognizes that it is participating in executing some portion of the original end-user transaction. When the FOR returns to the AOR, and the AOR returns to the TOR, there will be just one ended transaction recorded for RMF. Its response time will be the time from reception in the TOR until the TOR has completed.

So RMF data will show a reduced number of transactions executed and a larger response time after the AOR is migrated to CICS 4.1 — even when the end-user request is processed in exactly the same elapsed time.

Possible action

Increase CICS transaction goals: Do this prior to your conversion to an MRO connection, or prior to migrating your AOR to CICS 4.1, if the FOR transactions are classified to the same service class as your end-user transactions.

Analyzing coupling facility activity

This information unit presents the reports offered by RMF that assist you in monitoring coupling facility activity. For this purpose, you can use the Postprocessor Coupling Facility Activity report, a spreadsheet macro, or the various interactive Monitor III reports.

Using the Postprocessor Coupling Facility Activity report

This topic gives guidelines for some indicators when monitoring the coupling facility.

Coupling facility performance

Usage Summary

COUPLING FACILITY ACTIVITY												
z/OS V1R10		SYSPLEX UTCPLXJ8		DATE 11/28/08		INTERVAL 030.00.000						PAGE 1
		RPT VERSION V1R10 RMF		TIME 12.00.00		CYCLE 01.000 SECONDS						

COUPLING FACILITY NAME = CF1												
TOTAL SAMPLES(AVG) = 1781 (MAX) = 1799 (MIN) = 1671												

COUPLING FACILITY USAGE SUMMARY												

STRUCTURE SUMMARY												

TYPE	STRUCTURE NAME	STATUS CHG	ALLOC SIZE	% OF CF STOR	# REQ	% OF ALL REQ	% OF CF UTIL	AVG REQ/ SEC	LST/DIR ENTRIES TOT/CUR	DATA ELEMENTS TOT/CUR	LOCK ENTRIES TOT/CUR	DIR REC/ XI'S
LIST	DB2D_SCA	ACTIVE PRIM	16M	2.1	1766	0.9	2.0	1.96	41K 168	81K 483	N/A	N/A
	DB2E_SCA	ACTIVE PRIM	16M	2.1	1765	0.9	3.0	1.96	9915 154	20K 288	N/A	N/A
	DB2H_SCA	ACTIVE PRIM	16M	2.1	0	0.0	1.0	0.00	1024 1	1000 17	N/A	N/A
	IXCPLEX_PATH1	ACTIVE PRIM	5M	0.6	64993	34.8	25.0	72.21	7730 11	16K 96	N/A	N/A
LOCK	DB2D_LOCK1	ACTIVE PRIM	31M	4.0	1770	0.9	5.0	1.97	2032 0	0 0	131K 27	N/A
	DB2E_LOCK1	ACTIVE PRIM	31M	4.0	1770	0.9	5.0	1.97	15K 1	0 0	1049K 650	N/A
	DB2T_LOCK1	ACTIVE PRIM	125M	16.2	82436	44.1	30.0	91.60	0 0	0 0	66K 3215	N/A
CACHE	IRRXCF00_B001	ACTIVE	2M	0.2	221	0.1	1.0	0.25	390 0	75 0	N/A	0
	SYSIGGCAS_ECS	ACTIVE	5M	0.6	27907	14.9	16.0	31.01	28 0	17 0	N/A	0
STRUCTURE TOTALS			420M	54.4	186917	100	88.0	207.69				

STORAGE SUMMARY												

			ALLOC SIZE	% OF CF STORAGE	----- DUMP SPACE -----							
					% IN USE	MAX % REQUESTED						
TOTAL CF STORAGE USED BY STRUCTURES			263M	12.9								
TOTAL CF DUMP STORAGE			6M	0.3	0.0	0.0						
TOTAL CF STORAGE AVAILABLE			2G	86.8								
TOTAL CF STORAGE SIZE			2G									
			ALLOC SIZE	% ALLOCATED								
TOTAL CONTROL STORAGE DEFINED			2G	13.2								
TOTAL DATA STORAGE DEFINED			0K	0.0								

PROCESSOR SUMMARY												

COUPLING FACILITY		2084	MODEL B16	CFLEVEL 15	DYNDISP OFF							
AVERAGE CF UTILIZATION (% BUSY)		0.9	LOGICAL PROCESSORS:	DEFINED 5	EFFECTIVE 0.9							
				SHARED 2	AVG WEIGHT 50.0							

Figure 95. Coupling Facility Activity Report - Usage Summary

The Usage Summary part of the report gives summary data for the structures, the storage and the processors.

Indicator**Fields:**

LST/DIR	DATA	LOCK
ENTRIES	ELEMENTS	ENTRIES
TOT/CUR	TOT/CUR	TOT/CUR

Description: The first column reports on list entries defined and in use OR directory entries defined and in use, dependent on the type of structure. Similarly, the second column reports on the data elements defined and in use in the structure. The third column reports on lock entries.

The information presented on this report can be used to validate the size of various types of structures. In most cases, some knowledge of the structure is required to interpret the data.

- Cache structures

The CUR in-use values will indicate if the structure is overlarge (if it never fills up) or misproportioned in terms of entry to element ratio (one never fills up but the other does). It will not help in knowing if the structure is too small (i.e. 100% full is normal). To know if the structure is too small you have to look at reclaiming and perhaps hit ratios.

- List structures

Some list structures, like the JES2 checkpoint structure, are very static in size if not content/static data. For these structures, the TOT can be only slightly larger than CUR. Other functions, like the MVS logger, off-load data as the structure fills up. The threshold which triggers the Offload is based on parameters associated with the function. The ratio of CUR to TOT should not be higher than this threshold. Other list structures, such as the XCF structure and the VTAM Generic Resource structure, must handle large peaks of data. For these structures, the TOT should be much larger than the CUR in-use to prevent backups of signals during spikes.

- Lock structures

Lock structures, like the IRLM lock structures, are divided in two parts: The actual lock table contains the number of entries in the LOCK ENTRIES TOT column. The number of LOCK ENTRIES CUR is a sampled value of the number in use at the end of the interval. The best way to evaluate the size of this table is to look at the FALSE CONTENTION data for the lock structure on the Coupling Facility Structure Activity report, as described later in this document.

The second part of the lock structure contains RECORD DATA entries. The number of these entries is reported in the LIST ENTRIES column. If there are not enough of these entries, you will not be able to obtain locks, and transactions will start failing at that point.

IRLM provides no external means of apportioning the structure storage between record data and locks; you can only obtain adequate entries for both parts by judiciously choosing the size of the lock structure. The easiest way to do this is to use the lookup tables in INFOSYS Q662530.

Indicator

Field: DIR RECLAIMS - DIR RECLAIMS XI 'S

Description: All cache structures have directory entries and may or may not have data entries (for example, IMS OSAM and VSAM cache structures have only directory entries; while RACF and DB2 cache structures have both directory and data entries). N/A, for not applicable, will be displayed in the directory reclaim column for list and lock structures.

A cache structure can be over committed by the data base managers. This occurs when the total number of unique entities cached by the data base managers exceeds the number of directory entries in the cache structure. Whenever a shortage of directory entries occurs, the coupling facility will reclaim in-use directory entries associated with unchanged data. These reclaimed items will be used to satisfy new requests for directory entries by the data base manager. (Directory entries are used by the data base manager to ensure data validity). For the coupling facility to reclaim a directory entry, all users of the data item represented by the directory entry must be notified their copy of the data item is invalid. As a consequence, when the data base manager needs access to the now invalidated data item, the item must be re-read from DASD and registered with the coupling facility. When there are insufficient directory entries, the directory reclaim activity can lead to "thrashing". (The situation is analogous to real storage shortages and page stealing in MVS.) Directory reclaim activity can result in the following:

- Increased read I/O activity to the data base to re-acquire a referenced data item.
- Increased CPU utilization caused by registered interest in data items having to be re-acquired.
- Elongated transaction response times whenever a read miss occurs in the local buffer pool.

Directory reclaim activity can be managed by increasing the number of directory entries for a particular structure. This can be accomplished by:

- Increasing the size of the structure. For directory only structures, only the number of directory entries is affected (i.e. IMS OSAM and VSAM structures). For structures with data elements and directory entries, both will increase in the ratio specified by the structure user (i.e. RACF structures).
- Changing the proportion of the structure space used for directory entries and data elements. This action is dependent on the structure users' implementation. Some cache structure exploiters allow the installation to specify the ratio of directory entries to data entries, which it internally maps to a ratio of directory entries to data elements (a data entry may be composed of multiple data elements). DB2 provides this capability. On the other hand, RACF is hard coded to organize the structure in a 1:1 ratio of directory entries to data elements.

If the cache structure directory to data element ratio is installation specifiable AND there are ample data elements, then one can increase the number of directory entries at the expense of data elements without a performance impact, and without increasing the structure size. Determining the impact of decreasing the number of data items is at best an inexact science. Unless the Structure Summary report indicates a consistent difference between the total and current number of data elements, it is difficult to estimate the impact of changing the

directory entry to data element ratio. The CUR value reported by RMF is not a peak value or average value during the interval, it is the number in use at the end of the interval.

The usual causes for directory reclaim activity are:

1. Initial structure size is insufficient to support the amount of local buffering.
2. Increasing the size of the local buffer pools without an accompanying increase in structure size.
3. Increasing the number of users of the structure without an accompanying increase in structure size.

The advice here is to monitor trends. If directory reclaims are increasing and response time is impacted, take appropriate action.

The second value in this column is the number of reclaims that caused an XI (see 168). A high value is an indicator for a performance problem in this structure.

Indicator

Field: STRUCTURE SUMMARY - TOTAL CONTROL/DATA STORAGE DEFINED

Description: The amount of coupling facility storage that is allowed to be occupied by control information (CONTROL STORAGE) or data (DATA STORAGE).

Guideline: Each structure plus the dump area is allocated some control storage and some data storage. The coupling facility defines an area called control storage; structure control information is restricted to that area. The remaining storage is called data storage and is used for structure data. If the data storage area becomes full, structure data can then be allocated from the control storage area. If TOTAL DATA STORAGE DEFINED is zero, it means control information can reside anywhere on the coupling facility and there are no allocation restrictions.

If % ALLOC field for control storage shows a percentage approaching 100, it means the control storage is close to being completely allocated even though the CF SPACE AVAILABLE field may still show an amount of total free space. Possible 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

Indicator

Fields: AVERAGE CF UTILIZATION (%BUSY) LOGICAL PROCESSORS DEFINED / EFFECTIVE

Description: These fields report on the average CPU utilization of the CPs supporting the coupling facility partition, and in addition, on the number of logical processors assigned to the partition and the effective number of logical processors active during the RMF interval.

Coupling facility performance

For example, if a CEC contains six CPs, and a test coupling facility LPAR has two logical CPs, but is capped at 5% of the CEC, then the LOGICAL PROCESSORS DEFINED will be 2 and the EFFECTIVE LOGICAL PROCESSORS number will be 0.3 (if the partition is uncapped the EFFECTIVE LOGICAL PROCESSORS number will be 2).

CFCC does not enter a wait state when there are no commands to be processed. CFCC continues to search for work until LPAR management removes processor resources or work is found (the "active wait" phenomenon).

Note: The above statement is true for LPARs providing real coupling facility functionality. It is NOT true for LPARs providing integrated coupling facility (ICMF) function. LPAR and CFCC have been enhanced to allow an ICMF LPAR with nothing to do to give up its CP resources. In this environment, use the Monitor I Partition Data Report, rather than the Coupling Facility Usage Summary report, to determine the utilization of the ICMF LPAR.

Guideline: These fields will assist you in diagnosing the following types of problems which have occurred in customer environments.

- "High" coupling facility processor utilization and excessively high service times when only one CP out of six in the CEC was assigned to the coupling facility LPAR.
- "High" coupling facility processor utilization and excessively high service times caused by two LPARs in the same CEC with active CFs. (Production coupling facility weight of 80, test coupling facility LPAR weight of 50, all CPs shared). The test coupling facility consumed 5/13 of the available CP resources even though it was unused.

Coupling facility processor utilization does influence structure access service times and thus the responsiveness and CPU utilization of the coupling facility exploiter.

Please, see "Factors affecting Coupling Facility performance" on page 174 for additional information.

Structure Activity

COUPLING FACILITY ACTIVITY														
z/OS VIR10			SYSPLEX PLEXPERF			DATE 04/11/2008			INTERVAL 030.00.000			PAGE 3		
			RPT VERSION VIR10 RMF			TIME 13.39.00			CYCLE 1.000 SECONDS					
COUPLING FACILITY NAME = CF5														
COUPLING FACILITY STRUCTURE ACTIVITY														
STRUCTURE NAME = COUPLE_CKPT1 TYPE = LIST STATUS = ACTIVE														
SYSTEM NAME	# REQ TOTAL	AVG/SEC	# REQ	% OF ALL	-SERV TIME(MIC)- AVG	STD_DEV	REASON	# REQ	% OF REQ	DELATED	REQUESTS /DEL	AVG TIME(MIC)	STD_DEV	EXTERNAL REQUEST CONTENTIONS
J80	8463	4.70	2927	5.8	306.6	127.8	NO SCH	1240	22.4	508.4	597.7	113.9	REQ TOTAL	8339
			5535	11.0	1502.1	1263.2	PR WT	0	0.0	0.0	0.0	0.0	REQ DEFERRED	123
			1	0.0	INCLUDED	IN ASYNC	PR CMP	0	0.0	0.0	0.0	0.0		
							DUMP	0	0.0	0.0	0.0	0.0		
...														
TOTAL	50370	27.98	17K	34.3	342.2	124.5	NO SCH	8797	26.6	1047	1245	278.4	REQ TOTAL	50K
			33K	65.5	557551G	0.0	PR WT	0	0.0	0.0	0.0	0.0	REQ DEFERRED	497
			69	0.1			PR CMP	0	0.0	0.0	0.0	0.0		
							DUMP	0	0.0	0.0	0.0	0.0		
...														
STRUCTURE NAME = IRLMLOCK1 TYPE = LOCK STATUS = ACTIVE														
SYSTEM NAME	# REQ TOTAL	AVG/SEC	# REQ	% OF ALL	-SERV TIME(MIC)- AVG	STD_DEV	REASON	# REQ	% OF REQ	DELATED	REQUESTS /DEL	AVG TIME(MIC)	STD_DEV	EXTERNAL REQUEST CONTENTIONS
JF0	128K	71.26	128K	7.7	221.9	37.5	NO SCH	0	0.0	0.0	0.0	0.0	REQ TOTAL	128K
			0	0.0	0.0	0.0	PR WT	0	0.0	0.0	0.0	0.0	REQ DEFERRED	80
			0	0.0	INCLUDED	IN ASYNC	PR CMP	0	0.0	0.0	0.0	0.0	-CONT	80
							DUMP	0	0.0	0.0	0.0	0.0	-FALSE CONT	25
...														
TOTAL	1676K	930.9	1676K	100	198.1	44.7	NO SCH	0	0.0	0.0	0.0	0.0	REQ TOTAL	1674K
			0	0.0	0.0	0.0	PR WT	0	0.0	0.0	0.0	0.0	REQ DEFERRED	906
			0	0.0			PR CMP	0	0.0	0.0	0.0	0.0	-CONT	906
							DUMP	0	0.0	0.0	0.0	0.0	-FALSE CONT	225
...														
STRUCTURE NAME = DSNDBIG_GBP3 TYPE = CACHE STATUS = ACTIVE PRIMARY														
SYSTEM NAME	# REQ TOTAL	AVG/SEC	# REQ	% OF ALL	-SERV TIME(MIC)- AVG	STD_DEV	REASON	# REQ	% OF REQ	DELATED	REQUESTS /DEL	AVG TIME(MIC)	STD_DEV	EXTERNAL REQUEST CONTENTIONS
JA0	1032	0.57	1025	19.1	448.8	84.3	NO SCH	7	100	2060	1671	2060		
			0	0.0	4508.9	1409.1	PR WT	0	0.0	0.0	0.0	0.0		
			7	0.1	INCLUDED	IN ASYNC	PR CMP	0	0.0	0.0	0.0	0.0		
							DUMP	0	0.0	0.0	0.0	0.0		
...														
TOTAL	5356	2.98	5342	100	432.8	89.0	NO SCH	14	100	1510	1374	1510	-- DATA ACCESS --	
			0	0.0	3685.2	1634.2	PR WT	0	0.0	0.0	0.0	0.0	READS	1331
			14	0.3			PR CMP	0	0.0	0.0	0.0	0.0	WRITES	1709
							DUMP	0	0.0	0.0	0.0	0.0	CASTOUTS	44
													XI'S	1296

Figure 96. Coupling Facility Activity Report - Structure Activity

This section of the report provides detailed information on the frequency of structure operations and the associated service/queue times.

Coupling facility performance

Indicator						
Fields:						
----- DELAYED REQUESTS -----						
REASON	#	% OF	----	AVG TIME(MIC)	-----	
	REQ	REQ	/DEL	STD_DEV	/ALL	

Coupling Facility exploiters make structure manipulating requests through the APIs provided by Cross-system extended services (XES). In some instances, the API requests can specify that the operation be performed synchronously or asynchronously relative to requesting task. XES will execute the API request through some number of commands to the Coupling Facility. The coupling facility commands are in turn executed synchronously or asynchronously, based on the type of operation and the API specification. RMF reports on the operations to the coupling facility and the structures within. The only data on API requests is reported in CONTENTION data for lock and serialized list structures.

For an operation to be initiated against a coupling facility structure, one of the subchannels associated with the Coupling Facility must be available. When no subchannel is available, XES must either queue the operation for later scheduling, or "spin" waiting for a subchannel to become available. XES takes one action or the other, depending on the type of structure and operation.

Most SYNC operations which cannot be started due to unavailable subchannels are changed to ASYNC operations (RMF reports them as CHNGD). Certain SYNC operations are not changed from SYNC to ASYNC when no subchannels are available. Most notably, lock requests are not changed to ASYNC. For these operations, the processor "spins" until a subchannel becomes available.

RMF reports on ASYNC operations to a structure which are delayed due to subchannel unavailability. RMF does NOT report on SYNCH operations to a structure which are delayed ("spinning") due to subchannel unavailability. The amount of delay for SYNC operations must be estimated from the subchannel activity report.

Subchannel unavailable condition will manifest itself in one of the following ways:

1. A non-zero percentage of CHNGD requests relative to the total number of ASYNC requests.
2. A non-zero percentage of delayed ASYNC requests.

Depending on the structure either one or both of the above phenomena may be observed, not necessarily both.

In evaluating response times for ASYNC requests, the following should be noted:

- ASYNC command processing is performed primarily by the I/O processor (IOP or SAP).
- ASYNC requests generally take longer (elapsed time) than SYNC requests.
- The reported service time (not shown in the above example, but present on the ASYNC line) does not include delay time. To determine the average elapsed time for an ASYNC operation add the average service time to the average amortized delay time (/ALL column) for all ASYNC operations during the interval.

Coupling facility performance

- ASYNC delay times (/DEL column) can be an indication of inadequate processing capacity in the coupling facility as well as insufficient number of subchannels.

In evaluating response times for SYNC requests, the following should be noted:

- SYNC command processing is performed by the CP.
- The reported service time (not shown in the above example, but present on the SYNC line) does not include the delay spent "spinning" awaiting subchannel busy if OW12415 is installed.
- A SYNC request delayed due to unavailable subchannel is using CPU resources.

Subchannel unavailability has the following negative consequences:

- ASYNC requests will be queued, resulting in delayed access to a coupling facility structure and corresponding diminished responsiveness by the exploiter.
- Non-immediate SYNC requests will be transformed to ASYNC requests, resulting in delayed access to a coupling facility structure and corresponding diminished responsiveness by the exploiter.
- Immediate SYNC requests will "spin" waiting for an available subchannel, resulting in increased CPU usage, delayed access to a coupling facility structure and corresponding diminished responsiveness by the exploiter.

Normally, none of these effects are desirable. The following activities can minimize the detrimental effects:

- Make certain CPU resource is sufficient in the coupling facility LPAR.
- Redistribute structures among the CFs to minimize the delay.
- Add additional coupling facility links between the MVS processor and the coupling facility. Each path will contribute two subchannels.
- Influence the coupling facility activity rates (downward) of coupling facility exploiters if possible. For example, you can reduce the number of XCF signals by reducing lock contention or tuning XCF to eliminate signals related the expansion of transport classes sizes.

Indicator	
Fields:	
EXTERNAL REQUEST CONTENTIONS	EXTERNAL REQUEST CONTENTIONS
REQ TOTAL REQ DEFERRED	REQ TOTAL REQ DEFERRED -CONT -FALSE CONT
Description: These fields provide information on serialized list contention.	

A serialized list structure is a list structure which has an associated lock structure within. For example, JES2 and VTAM structures are serialized list structures; the XCF structure is simply a list structure. A serialized list structure can provide the structure user with increased control (exclusive usage) over the entire structure, a portion of the entries or even a single list entry. The use of the lock structure to control access is completely under the control of the exploiter.

Coupling facility performance

There are a number of operations against a serialized list structure. Some operations against a serialized list structure will succeed if the designated lock entry is not held by any other user. Additionally, these requests will be queued for FIFO processing by XES whenever the command cannot be processed due to lock contention. (Specifically these are UNCONDITIONAL SET and NOTHELD requests). XES will queue the request in a sysplex-wide FIFO queue until the lock is available. Whenever the lock is released, the first queued request will be restarted on the appropriate processor. XES is providing multi-system FIFO queuing on behalf of the users to the serialized list structure.

RMF will report on the number of accesses to the serialized list structure and the number of accesses deferred due to lock contention. This value is reported from the perspective of the indicated connector, the value is not a "global" value across all connectors to the structure. Note: A serialized list structure will not incur false contention due to the manner in which XES manages the contention.

Having XES queue and manage deferred requests requires some amount of CPU usage and XCF signalling usage. Excess contention is most likely a reflection on the usage of the structure itself by the exploiter or poor tuning specifications by the installation. For example, consider JES2s use of a checkpoint structure. Each JES2 member of the MAS will acquire and hold the checkpoint for some period of time (HOLD) and restrain itself from attempting to acquire control for some period of time upon release (DORMANCY). If all members of the MAS specify short dormancy values, then each member (except the current owner) will attempt to acquire the checkpoint lock. XES will happily manage this queue of requestors in a FIFO manner at some cost. It may be worth the effort to adjust these values so members of the MAS receive the necessary service without the overhead of XES managing the queue of requestors.

The average time a serialized list request with lock contention is queued by XES is not tracked and thus not reported by RMF. APAR OW11789 corrected a reporting problem which added the queue time into the AYSNC service time. This resulted in inflated and widely varying service times being reported.

Indicator

Fields:

```
-- DATA ACCESS ---  
READS  
WRITES  
CASTOUTS  
XI'S
```

Description: These fields provide information on cache structure accesses.

Data access counts for a cache structure have been added to the report. Prior to describing what READS, WRITES, etc. counts actually count, the following should be noted:

- The information is acquired from counters in the coupling facility and is global in nature. The information cannot be broken down to give an individual connection contributions to this global count. Thus, the information is reported in the "TOTAL section" of the RMF report.
- Depending on the cache structure implementation one or more of these counter values may be zero.

- Most importantly, only the cache structure user (i.e. data base manager) knows how efficiently the local buffers and cache structure are being used. If the data base manager does not provide usage information, then tuning is a best guess proposition.

The coupling facility maintains a number of structure level counters accessible to RMF. RMF reports on a subset of these counters (to be described momentarily) and stores all the counters in the SMF74 subtype 4 records.

Additional information on these counters can be obtained in *z/OS MVS Programming: Sysplex Services Guide* in the section labelled IXLCACHE REQUEST=READ_STGSTATS.

RMF provides four cache structure usage counters.

READS

This is a count of the number of times the coupling facility returned data on a read request by any connector.

Directory only caches will always have a zero value reported since the structure contains no "data".

WRITES

This is a count of the number of times a connector placed changed or unchanged data into the coupling facility structure. (Changed/unchanged is an attribute assigned to the data when written by the connector. From a performance/capacity view point, the importance of the attribute is: changed data cannot be reclaimed from the structure should directory or data elements become scarce.)

Just as directory only cache structures always have READS counts of zero, they have WRITES counts of zero as the structure contains no "data".

The condition to be concerned with is a large number of WRITES and a small number of READS. This condition may indicate:

- Insufficient structure space allocated, and data entries (and perhaps directory entries) are being discarded by the coupling facility space management routines.
- Inappropriate allocation of the ratio of directory entry to data elements is causing the data entries to be discarded by the coupling facility space management routines.

Directory entry reclaim activity is reported by RMF as discussed earlier; however, data element reclaim activity is not reported.

CASTOUTS

This is a count of the number of times a connector retrieved a changed data entry, wrote the data to DASD (discarded it, whatever), and caused the changed attribute to be reset to unchanged. This process is known as casting out changed data from the structure.

This counter is of interest for store-in cache structures (i.e. DB2 global buffer pool structures) in determining the volume of changed data being removed from the structure. This counter is not an indicator of the number times castout processing was performed during the RMF interval.

A large amount of castout activity on a single structure may warrant additional cache structures and redirecting locally buffered data to different cache structure. Castout processing by the connectors must keep pace with the rate at which changed data is placed in the structure. When all

Coupling facility performance

directory or data elements are associated with changed data, no new data items can be registered or written to the structure. This condition is not desirable and will adversely effect the data base manager and/or the user of the data base manager.

XI'S This is the number of times a data item residing in a local buffer pool was marked invalid by the coupling facility during the RMF interval. To the cache structure user, this means the data item must be re-acquired from DASD or perhaps the coupling facility structure, and interest in the item must be re-registered in the coupling facility structure.

There are several "XI counts" obtained from the coupling facility which are consolidated into this value. They are:

- XI for Directory Reclaim
- XI for Write
- XI for Name Invalidation
- XI for Complement Invalidation
- XI for Local Cache Vector Entry Replacement

Consult section IXLCACHE REQUEST=READ_STGSTATS in the *z/OS MVS Programming: Sysplex Services Guide* for additional information on these XI categories.

It is important to understand what the XI'S counter really counts. This is most easily explained by example. Suppose there are 5 connectors (i.e. #1, #2, #3, #4, and #5). to a cache structure. Further suppose data item "A" is locally cached in connectors #1, #2, #4, and #5 and registered in the coupling facility structure. Now suppose connector number #5, changes data item "A" and issues the command to the coupling facility to have all other copies invalidated (i.e. XI for complement invalidation). In response to the command, the coupling facility would cause the copies of data item "A" in the local buffer pools of connectors #1, #2 and #4 to be marked invalid. At the completion of the XI operation, the XI'S count value would be incremented by 3!

XI'S count values are seen for directory, store-in and store-thru caches. This count reflects both the amount of data sharing among the users of the cache and the amount of write/update activity against the data bases.

Subchannel Activity

COUPLING FACILITY ACTIVITY																
z/OS V1R10		SYSPLEX PLEXPERF			DATE 04/11/2008			INTERVAL 030.00.000							PAGE 6	
		RPT VERSION V1R10 RMF			TIME 13.39.00			CYCLE 1.000 SECONDS								

COUPLING FACILITY NAME = CF1																

SUBCHANNEL ACTIVITY																
SYSTEM NAME	# REQ TOTAL	CF TYPE	LINKS GEN	PTH USE	BUSY	REQUESTS			DELATED		REQUESTS		TIME(MIC)			
						# REQ	-SERVICE AVG	TIME(MIC)- STD_DEV	# REQ	% OF REQ	AVG /DEL	STD_DEV	/ALL			
*W04	30740	ICP	3	3	0	SYNC	957	26.8	95.5	LIST/CACHE	2	0.0	174.0	50.9	348.0	
	51.2	SUBCH	42	21		ASYNC	22791	736.7	2811	LOCK	0	0.0	0.0	0.0	0.0	
						CHANGED	0	INCLUDED	IN ASYNC	TOTAL	2	0.0				
						UNSUCC	0	0.0	0.0							
*W05	9793	CFS	2	2	354	SYNC	790	43.3	116.5	LIST/CACHE	1537	19.1	1797	5626	2761K	
	9.6	SUBCH	4	4		ASYNC	8100	937.8	3543	LOCK	4	0.5	69.5	15.6	278.0	
						CHANGED	1	INCLUDED	IN ASYNC	TOTAL	1541	17.3				
						UNSUCC	0	0.0	0.0							

Figure 97. Coupling Facility Activity Report - Subchannel Activity

The following fields 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.

Indicator

Field: CF LINKS

Description: The subchannel configuration is described by these fields:

TYPE This column describes the channel path type.

GEN USE The number of subchannels that are defined, and that MVS is currently using for coupling facility requests.

The description of the channel path type can help in better analyzing the performance of the different channel path types. There are differences between two types of coupling facility channels:

CFS /CFR - CBS /CBR - ICS / ICR
Two subchannels per path.

CFP - CBP - ICP
Seven subchannels per path.

Coupling facility performance

Indicator

Field: PTH BUSY

Description: Path busy - the number of times a coupling facility request was rejected because all paths to the coupling facility were busy.

Guideline: A high count combined with lengthy 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 the path busy count is low. In fact, in a non-PR/SM environment where the subchannels are properly configured, # DELAYED REQUESTS, not PTH BUSY, is the indicator for path contention. If PTH BUSY is low but # DELAYED REQUESTS is high, it means MVS is delaying the coupling facility requests and in effect gating the workload before it reaches the physical paths.

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 underconfigured), subchannel busy is low, but path busy is high. This means the contention is due to activity from a remote partition.

If a coupling facility capacity constraint is suspected, the first action to consider is adding more paths. If the coupling facility is already fully configured for paths, you need to determine whether some structures can be off-loaded to another coupling facility or whether another coupling facility should be added to the configuration. Use the Coupling Facility Usage Summary reports to balance workloads across coupling facilities based on activity rates and storage usage.

In a PR/SM environment, you may need to increase or adjust the configuration of shared coupling facility channels to reduce path contention.

Indicator

Fields:

```

----- DELAYED REQUESTS -----
      #      % OF  -----AVG TIME(MIC)-----
      REQ    REQ   /DEL  STD_DEV  /ALL

LIST/CACHE
LOCK
TOTAL
    
```

Description: Field TOTAL # DELAYED REQUESTS is the same as BUSY COUNT - SCH that was shown in previous versions of this report. It is the number of times an immediate request was delayed because subchannel resources were not available.

Immediate requests (such as locking operations) could not be completed because all the subchannels were busy. These requests are not queued. They are processed before the non-immediate requests which are reported in the Structure Activity report under QUEUED requests. To get a complete picture of subchannel activity, look at both fields.

Guideline: If this count is high, you should ensure that sufficient subchannels are defined.

See PTH BUSY above for suggested actions to relieve a coupling facility path constraint.

In a data sharing environment, lock structure requests account for a large percentage of coupling facility accesses. By estimating the amortized delay for SYNC operations on lock structure accesses, one can determine the time delay due to subchannel unavailable on lock structure operations and the impact on CPU busy.

Note: Requests which are delayed because a structure was being dumped were removed from this report because they are not related to subchannel contention.

CF to CF Activity

COUPLING FACILITY ACTIVITY													
z/OS V1R10		SYSPLEX PLEXPERF			DATE 04/11/2008			INTERVAL 030.00.000			PAGE 6		
		RPT VERSION V1R10 RMF			TIME 13.39.00			CYCLE 1.000 SECONDS					

COUPLING FACILITY NAME = CF1													

CF TO CF ACTIVITY													

PEER CF	# REQ TOTAL	-- CF LINKS -- AVG/SEC TYPE USE	--	REQUESTS				DELAYED REQUESTS					
				# REQ	-SERVICE TIME(MIC)- AVG	STD_DEV	SYNC	# REQ	% OF REQ	---- AVG TIME(MIC) --- /DEL	STD_DEV	/ALL	
CF2	86830	CBP	1	86830	53.0	33.8	SYNC	19	0.2%	14.5	2.7	0.0	
	86.2	ICP	1										

Figure 98. Coupling Facility Activity Report - CF to CF Activity

This section of the Coupling Facility Activity report provides information about the activities between coupling facilities.

Coupling facility performance

Indicator Fields:

```

----- DELAYED REQUESTS -----
#      % OF      -----AVG TIME(MIC)-----
REQ   REQ   /DEL  STD_DEV  /ALL
  
```

Description: Field # DELAYED REQUESTS is the number of signals of all types which have experienced a delay in being sent from the subject CF to this remote CF. It is a contention indicator for the peer-to-peer communication between both coupling facilities.

Spreadsheet report

If you are interested to display the data of the Coupling Facility Activity report graphically, you can use the Spreadsheet Reporter with an interval report as input. One of the several charts that you can get is the following one:

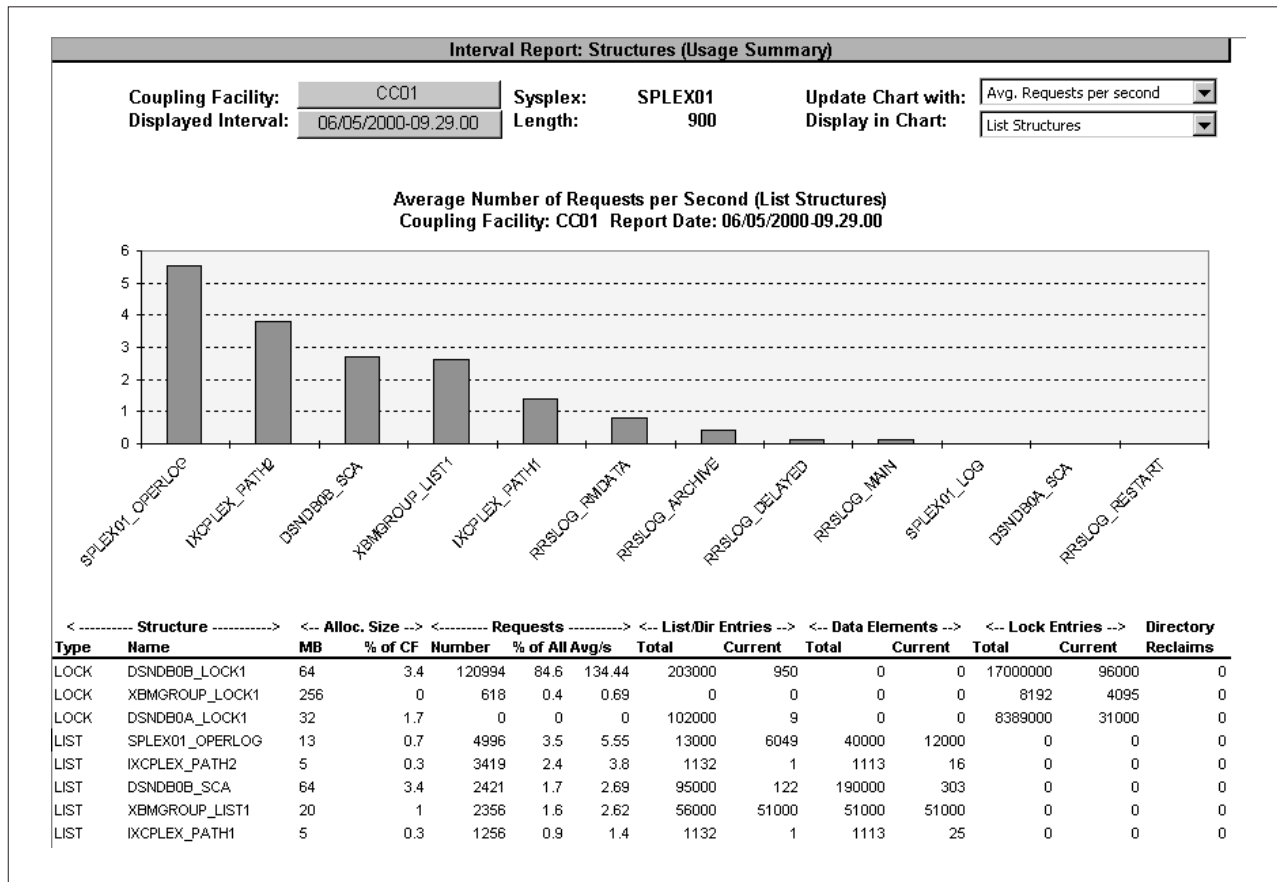


Figure 99. Coupling Facility Structure Activity. Spreadsheet Reporter macro RMFR9CF.XLS (CF Report)

The graphic shows the total number of requests for structure ISGLOCK, and you can modify the graphic by selecting a specific type of data in the drop-down list.

Using the Monitor III online reports

Three Monitor III reports enable you to analyze bottlenecks or problems in the coupling facility area that might result in a performance degradation in the parallel

sysplex. To avoid critical situations for production data bases and transaction processing systems, customers can see the immediate state of the coupling facility and the structures, and they have all relevant data at a glance in a granularity they can choose. In addition, they can see the results of tuning actions in this area, for example after having removed a coupling structure from an overloaded coupling facility to another one.

The Coupling Facility Overview report - CFOVER

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

RMF V1R10 CF Overview - TRXPLEX Line 1 of 2												
Samples: 120			Systems: 2		Date: 11/28/08		Time: 14.40.00		Range: 120		Sec	
CF Policy: SYSPOL1				Activated at: 11/08/06 09.16.04								
----- Coupling Facility -----					----- Processor -----					Request - Storage --		
Name	Type	Model	Lvl	Dyn	Util%	Def	Shr	Wgt	Effect	Rate	Size	Avail
CF01	2094	S18	15	OFF	9.6	1	1	100	1.0	9585	423M	171M
CF02	2094	S18	15	ON	91.3	2	1	50	1.0	4140	423M	232M

Figure 100. CFOVER Report

The Coupling Facility Activity report - CFACT

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

Coupling facility performance

RMF V1R10		CF Activity		- TRXPLEX				Line 1 of 78		
Samples: 120	Systems: 2	Date: 11/28/08	Time: 14.40.00	Range: 120	Sec					
CF: ALL	Type	ST	System	CF Util %	--- Sync Rate	--- Avg Serv	----- Rate	----- Avg Serv	----- Chng %	----- Del %
Structure Name										
BIGONE	LOCK	AP	*ALL	1.0	0.4	330	30.1	492	0.0	0.0
	LOCK		TRX1		0.4	330	30.1	492	0.0	0.0
	LOCK		TRX2		0.0	0	0.0	0	0.0	0.0
IGWLOCK00	LOCK	AP	*ALL	1.2	0.0	0	0.0	0	0.0	0.0
	LOCK		TRX1		0.0	0	0.0	0	0.0	0.0
	LOCK		TRX2		0.0	0	0.0	0	0.0	0.0
	LOCK	AS	*ALL	0.1	0.0	0	0.0	0	0.0	0.0
	LOCK		TRX1		0.0	0	0.0	0	0.0	0.0
	LOCK		TRX2		0.0	0	0.0	0	0.0	0.0
IXCPLEX_PATH1	LIST	A	*ALL	0.1	0.0	0	0.9	406	0.0	0.0
	LIST		TRX1		0.0	0	0.5	98	0.0	0.0
	LIST		TRX2		0.0	0	0.4	786	0.0	0.0
JES2CKPT1	LIST	AP	*ALL	0.2	0.2	30	2.0	570	0.0	0.0
	LIST		TRX1		0.1	7	1.1	374	0.0	0.0
	LIST		TRX2		0.1	53	0.9	802	0.0	0.0
	LIST	AS	*ALL	0.0	<0.1	595	1.5	494	0.0	0.0
	LIST		TRX1		0.0	0	0.8	432	0.0	0.0
	LIST		TRX2		<0.1	595	0.7	566	0.0	0.0
SYSZWM_6AAD2094	CACHE	AP	*ALL	0.2	0.0	0	14.5	406	0.0	0.0
	CACHE		TRX1		0.0	0	7.8	293	0.0	0.0
	CACHE		TRX2		0.0	0	6.7	537	0.0	0.0
	CACHE	AS	*ALL	0.0	0.0	0	7.4	753	0.0	0.0
	CACHE		TRX1		0.0	0	4.0	566	0.0	0.0
	CACHE		TRX2		0.0	0	3.4	974	0.0	0.0

Figure 101. CFACT Report

The Coupling Facility Systems report - CFSYS

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

RMF V1R10		CF Systems		- TRXPLEX				Line 1 of 4			
Samples: 120	Systems: 2	Date: 11/28/08	Time: 14.40.00	Range: 120	Sec						
CF Name	System	Subchannel	-- Paths --	-- Sync --	--- Rate	--- Avg Serv	----- Rate	----- Avg Serv	----- Chng %	----- Del %	
		Delay %	Avail %	Rate %							
		Busy %	Delay %								
CF01	TRX1	0.0	2.3	4	0.0	6359	34	2778	147	0.0	0.0
	TRX2	0.0	0.0	4	0.0	2.0	9	9.1	606	0.0	0.0
CF02	TRX1	0.0	5.9	4	0.0	1116	83	2984	519	0.0	0.0
	TRX2	0.0	0.0	4	0.0	0.0	0	4.0	928	0.0	0.0

Figure 102. CFSYS Report

Factors affecting Coupling Facility performance

This section is intended to assist one in detecting and assessing processing delays which will impact applications utilizing the coupling facility. Processing delays

may manifest themselves as increased in CPU time, increased in elapsed time due to waiting, or some combination of both. There are several approaches to achieving better application or system performance:

- Perform "work" as quick as possible.
- Don't create additional "work".
- Don't do unnecessary work.

Perform "work" as quick as possible

Conditions elongating coupling facility request response (service + delay) times include the following:

- Insufficient processor power allocated to the coupling facility LPAR.
The most common cause of this condition is the sharing of CPs among coupling facility LPARs. The changes to RMF and CFCC will allow the installation to accurately determine the amount of available processing power being used. Elongation of SYNC service times against directory only cache structures is a good indicator of this condition. Alternatively, use the lock structure SYNC service times after removing the no subchannel delay component.
- Insufficient number of paths between the CECs containing multiple MVS images and a coupling facility
RMF does not report on the length of the delay due to path busy conditions, only the number of requests encountering the condition are reported. Requests encountering path delay will have elongated service times.
The most common condition to cause path busy condition is EMIFing the sender ISC links among the MVS LPARs. The current rule of thumb is that the number of requests delayed for path busy (as reported on the Subchannel Activity report in the PTH field under the BUSY COUNTS column) should be less than 10% of total requests (on the same report under the # REQ TOTAL column).
Note: In the case of EMIFed ISC links, be certain to take into account all the MVS LPARS utilizing a common physical connection(s).
The solution to reducing the incidence of path busy conditions is to add ISC links between the processors or refrain from using EMIF, which will probably require additional links as well.
- Insufficient number of subchannels between the MVS image and the coupling facility
RMF reports on delay times and service times at the structure level and the coupling facility level (as previously discussed). Consequences of insufficient subchannels are:
 - Increased MVS CPU consumption when the SYNC operation cannot be converted to an ASYNC operation.
 - Increased service time for the SYNC operation which cannot be converted to an ASYNC operation.
 - Requests being converted from SYNC to ASYNC.
 - Increased queue time for ASYNC requests.The effect of the last two items is difficult to quantify. The effect depends on what the exploiter can do while a request is queued (i.e. was the delay anticipated and the request made asynchronously). Either of these two events will result in an elapsed time increase for some function. Whether this affects only one "user" or many users is dependent on the exploiter and its use of the structure. For example, if all cache structure castout operations are delayed, then it is possible for the cache to become full of changed data and is not a desirable condition.

Coupling facility performance

There are two indicators to be monitored.

1. Increased SYNC operation service times (particularly against the lock structure) and the amount of "spin time" associated with the service time.
2. The percentage of operations delayed.

Use the same rules and values previously discussed.

The only way to increase the number of subchannels is to add more physical paths (i.e. ISC links). There are two subchannels per ISC link. Additional ISC links will help alleviate path busy conditions in an EMIF environment as well.

- Insufficient IOP/SAP capacity on the CEC MVS image and the coupling facility
The IOP/SAP on the CEC housing the MVS image handles IO to DASD, CTC traffic and ASYNC operations. As the number of ASYNC operations increases, the workload on the IOP/SAP increases. If the IOP/SAP becomes overloaded, the service time for ASYNC operations will elongate. ASYNC operations can be quite long and have a large standard deviation, thus changes are somewhat difficult to detect. Perhaps the easiest indicator to track is the Avg Q Length on the I/O Queuing Activity Report; if this value gets large it is an indicator that ASYNC requests as well as I/O requests are being delayed in IOP/SAP processing.
- Hardware problems causing ISC link degradation.
Monitor the hardware incidents in LOGREC and relevant console messages. Get the problem fixed.

Do not create additional "work"

XES and the coupling facility must work within the constraints imposed by the people that setup the parallel sysplex. When structures are not of sufficient size, XES, the coupling facility, and the exploiter will cope as best they can.

For example if a directory only cache structure is being utilized and insufficient structure space is defined, directory entries will be reclaimed according to a least recently used algorithm. This will result in locally cached data items being marked invalid. Whenever these locally cached, but invalid, items are referenced, they will have to be read from DASD and registered with the coupling facility. In essence, unnecessary work has been created due to an improper definition. CPU cycles are required to perform the I/O and to register the data item with the coupling facility. Furthermore, the entity accessing the data item was delayed (elapsed time increased) while these two events occurred.

RMF now provides information allowing the installation to detect, and perhaps alleviate, these conditions on a structure by structure basis.

- Lock Structures

Lock structure request service times should be short. Lock contention should be avoided, as contention must be resolved with the assistance of XES lock services requiring time and CPU cycles. Lock contention also generates additional XCF signalling. False lock contention should be avoided, it is expensive to resolve and is avoidable. RMF reports on the number of lock requests, the number delayed due to contention and the number delayed due to false contention.

False lock contention can be alleviated by increasing the size of the lock table. To reduce true lock contention may require a change to the application, the data base, or the workload mix during the period of time the conflict occurs. One should attempt to manage true contention to less than 1.0 percent and false contention to less than 0.1 percent of the requests.

- Cache Structures

Cache structures are not meant to replace local caches, they are intended to augment their use by:

1. Providing a mechanism for local buffer coherency
2. Providing additional high speed caching for frequently changed data.

Use of the cache structure should not negatively impact the use of local cache.

Prevent invalidation of a local cache entity unless the entity has been changed. This means avoid doing directory entry reclaims. This activity may be avoided by specifying a larger structure. RMF now reports on this activity to assist the installation in verifying the cache structure contains sufficient directory entries to keep track of locally cached data.

If the cache structure user keeps data in the structure, attempt to keep frequently referenced data in the structure. Caching data in the coupling facility is only a performance boost when the data resides in the structure. It is important to have the structure sufficiently large and of the correct entry/data item configuration to support coupling facility caching of data without hindering performance.

The RMF changes discussed previously will assist the installation in tracking anomalies or changes in data access patterns. RMF data is not sufficient in determining optimal read hit ratios.

Avoid unnecessary work

This activity is difficult to assess. The servers using coupling facility structures for message passing, locking, caching, state saving, etc. normally do very little coupling facility accesses unless prompted by some user of the server. Normally applications are causing the servers to perform some task involving coupling facility accesses.

In extensive data sharing environments, the data base manager and the lock manager together will account for over half of accesses to the coupling facility. The applications using the data base manager are normally high profile and critical to the business. Thus the first two activities are most often pursued.

When all coupling facility tuning efforts have failed to meet the applications response time objectives, be it CPU time and/or elapsed time, examination of the application is in order. Unfortunately, this involves real work among different organizations. The causes of contention, excessive locking, cache disruption, etc. must be determined and reduced. This reduction will involve some type of application change. This is obviously not a popular option and is presented last as it is most begrudgingly done.

Coupling facility performance

Appendix A. PR/SM LPAR considerations

Considerations in PR/SM Environments

This appendix presents performance considerations for systems running in LPAR mode, including:

- Interpretation of RMF reports
- Tuning of an LPAR configuration

Understanding the Partition Data report

Using the Postprocessor report option `REPORTS(CPU)`, you get the CPU Activity report. This report is divided into these sections:

CPU ACTIVITY — Information about processors and address spaces

PARTITION DATA REPORT — Information about all configured partitions

LPAR CLUSTER REPORT — Information about LPAR clusters

Although this is an inconsistency, the first section is typically referred to as CPU Activity report, too. The third section is new, it is available with the Intelligent Resource Director (IRD) of the zSeries processors and is described in Appendix B. The Intelligent Resource Director.

CPU ACTIVITY										PAGE 1		
z/OS V1R10			SYSTEM ID NP1			DATE 11/28/08		INTERVAL 14.59.678				
			RPT VERSION V1R10 RMF			TIME 09.30.00		CYCLE 1.000 SECONDS				
CPU	2097	MODEL 720	H/W MODEL	E26	SEQUENCE CODE	00000000000473BF		HIPERDISPATCH=NO				
---CPU---		ONLINE		LPAR BUSY		MVS BUSY		PARKED		LOG PROC	--I/O INTERRUPTS--	
NUM	TYPE	TIME %	TIME %	TIME %	TIME %	TIME %	TIME %	TIME %	TIME %	SHARE %	RATE	% VIA TPI
2	CP	19.33	16.77	100.0	0.00	22.0	0.00	0.00	0.00			
3	CP	100.00	26.49	99.88	0.00	22.0	7.98	2.48				
CP TOTAL/AVERAGE		24.92		99.90		44.0	7.98	2.48				

Figure 103. CPU Activity Report

The first part of the CPU Activity report provides information about the processors assigned to that partition which is gathering the data.

Report Analysis

- ONLINE TIME PERCENTAGE

The total online time percentage for all processors is 119.33% which is equivalent to 1.2 processors. This value is shown in the Partition Data report in the field `PROCESSOR NUM`.

- LPAR BUSY TIME PERC

The average utilization of the processors is 24.92%, this value can be found in the Partition Data report in the field `LOGICAL PROCESSORS TOTAL` for partition NP1. The LPAR BUSY time is (with Wait Completion = NO) the dispatch time of the processors that are assigned to the partition, and the percentage is based on their online time.

- MVS BUSY TIME PERC

This field shows the MVS view of the CPU utilization. The MVS BUSY time is the difference between online time and wait time. If the MVS system is busy when the partition is losing control, it stays in busy mode until the partition will be dispatched again, then the current task can continue, and wait state will be reached later. Therefore, the MVS BUSY time can be higher than the LPAR BUSY time, and the difference between both values is an indicator for CPU constraints in the system.

The Partition Data report shows measurement data for all configured partitions. The line `*PHYSICAL*` is for reporting purposes only; it does not reflect a real partition.

PHYSICAL PROCESSORS TOTAL								
DISPATCH TIME DATA TOTAL (Partition 1)			...	DISPATCH TIME DATA TOTAL (Partition n)			...	*PHYSICAL*
DISPATCH TIME DATA EFFECTIVE			...	DISPATCH TIME DATA EFFECTIVE			...	
MVS Captured Time	MVS Uncaptured Time	Partition LPAR Mgmt Time	...	MVS Captured Time	MVS Uncaptured Time	Partition LPAR Mgmt Time	...	LPAR Time not attributed

PARTITION DATA REPORT																
z/OS V1R10		SYSTEM ID NP1				DATE 11/28/08				INTERVAL 14.59.678						
		RPT VERSION V1R10 RMF				TIME 09.30.00				CYCLE 1.000 SECONDS						
MVS PARTITION NAME		NP1				NUMBER OF PHYSICAL PROCESSORS				9		GROUP NAME		N/A		
IMAGE CAPACITY		100				CP				7		LIMIT		N/A		
NUMBER OF CONFIGURED PARTITIONS		9				ICF				2						
WAIT COMPLETION		NO														
DISPATCH INTERVAL		DYNAMIC														
----- PARTITION DATA ----- -- LOGICAL PARTITION PROCESSOR DATA -- -- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --																
NAME	S	WGT	DEF	ACT	DEF	WLM%	PROCESSOR- NUM	TYPE	DISPATCH TIME DATA- EFFECTIVE	TOTAL	LOGICAL PROCESSORS EFFECTIVE	TOTAL	PHYSICAL PROCESSORS LPAR MGMT	EFFECTIVE	TOTAL	
NP1	A	20	100	10	NO	62.2	1.2	CP	00.04.27.302	00.04.27.519	24.86	24.92	0.01	4.24	4.25	
NP2	A	1	0	1	YES	0.0	4	CP	00.00.21.680	00.00.22.083	0.60	0.61	0.01	0.34	0.35	
NP3	A	10	5	8	NO	3.3	1.0	CP	00.03.35.761	00.03.35.859	23.97	23.98	0.02	3.41	3.43	
NP4	A	300	95	155	NO	0.0	6.0	CP	01.12.05.405	01.12.06.405	80.10	80.15	0.01	68.68	68.69	
NP5	A	200	50	52	NO	0.0	4	CP	00.24.11.147	00.24.11.311	40.30	40.31	0.01	23.02	23.03	
PHYSICAL									00.00.03.103				0.05		0.05	
TOTAL									01.44.41.295 01.44.46.280				0.11		99.69 99.80	
CFC1	A	DED					1	ICF	00.14.59.611	00.14.59.625	99.98	99.99	0.01	99.95	99.96	
CFC2	A	DED					1	ICF	00.00.00.000	00.00.00.000	0.00	0.00	0.00	0.00	0.00	
PHYSICAL									00.00.00.065				0.03		0.03	
TOTAL													0.04		99.95 99.99	
CB88	D															
CB89	D															

Figure 104. Partition Data Report

Report Analysis

- PROCESSOR NUM

The number of online processors for partition NP1 is **1.2**, this reflects the online time percentages of 19.33% + 100% = 119.33% in the CPU report.

- LOGICAL PROCESSORS - TOTAL DISPATCH TIME

DISPATCH TIME **4.27.519** = 268 sec
 ONLINE TIME 119.33% of **14.59.678** = 1074 sec
 LP TOTAL UTIL (268/1074)*100 = **24.92** %

The value of 24.92% is shown as LPAR BUSY TIME PERC in the CPU report.

Report Analysis

- PHYSICAL PROCESSORS - TOTAL DISPATCH TIME
DISPATCH TIME 4.27.519 = 268 sec
INTERVAL TIME (7 CPSs) 7 * 14.59.678 = 6298 sec
LP TOTAL UTIL (268/6298)*100 = 4.25 %

Partition NP1 has a processor utilization of 4.25% of the total 2064-109 system.

- PHYSICAL PROCESSORS - LPAR MGMT

Each partition's CPU consumption for LPAR management is calculated as the difference between total and effective dispatch 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. In this case, field LPAR MGMT is filled with '****'.

- *PHYSICAL*

The *Physical Management Time* is collected and reported by RMF in this line. The partition named *PHYSICAL* does not exist, this line is created for reporting purposes.

Report Analysis

- LOGICAL PROCESSORS - UTILIZATION CF PARTITION
PARTITION CFC1 99.99% UTILIZATION

CF partitions are actually always busy 100% of the time. For purposes of reporting CF utilization, the CPU time is accumulated in two buckets: *busy* and *idle*. *Idle* is the CPU time spent looking for work (the polling loop), *busy* is all other time (time spent processing a command or background work). From these times, the CF Utilization reported in the Coupling Facility Activity report is 100*busy/(busy+idle).

Defining logical processors

Obviously, there is a cost to operate in LPAR Mode. One of the key aspects for defining an LPAR configuration is the number of logical partitions.

Defining more logical partitions than are required to handle a partition's workload affects LPAR management time. Define as few LPARs as are needed to support your business needs.

Rule-of-Thumb

Number of Logical Processors

Number of Logical Processors ≤ 2 * Number of Physical Processors

The following example demonstrates the impact of the definition of logical processors on the total utilization of a system. In this sample, the processor is an 8-way system.

For a system with 8 physical processors, 9 partitions have been defined with a total of 70 logical processors which is far away from the rule-of-thumb that has been mentioned above.

LPAR #	LPAR Name	# LP	Weights
1	PROD1	8	80
2	PROD2	8	50
3	PROD3	8	40
4	PROD4	6	12
5	PROD5	8	50
6	TEST	8	7
7	PROD6	8	106
8	PROD7	8	40
9	PROD8	8	5

With this configuration, the LPAR management time percentage was 22.83% which is nearly the capacity of 2 processors in this 8-way system. Due to significant performance problems, the configuration has been changed with defining 20 LPs for the same number of partitions, which is close to the 2:1 rule. The weighting factor was taken to estimate the number of LPs for each partition.

LPAR #	LPAR Name	# LP	Weights
1	PROD1	4	80
2	PROD2	2	50
3	PROD3	1	40
4	PROD4	2	12
5	PROD5	3	50
6	TEST	1	7
7	PROD6	4	126
8	PROD7	2	40
9	PROD8	1	5

By changing the number of LPs from 70 to 20, the LPAR management time percentage could be reduced by 20% of the total capacity (from 22.83% to 2.87%) for an 8-way processor which is equivalent to the capacity of 1.6 processors.

Intelligent Resource Director

Considerations about the number of logical processors are relevant and required only for partitions which are not under control of WLM LPAR management. This is a function of the Intelligent Resource Director which is available on z900 servers, and it is described in the following chapter.

PR/SM environment

Appendix B. The Intelligent Resource Director

Learning about the Intelligent Resource Director

This appendix describes the functions of the Intelligent Resource Director and its reporting in RMF, including:

- Dynamic Channel Path Management
- Channel Subsystem Priority Queuing
- LPAR CPU Management

Overview

The Intelligent Resource Director (IRD) is a functional enhancement exclusive to IBM's zSeries family. It is provided in z900 servers and is an extension of IBM's industry leading clustering architecture, the Parallel Sysplex. IRD uses LPAR clusters (see "LPAR cluster" on page 190 for a description) that will efficiently balance processor and I/O resources between multiple applications based on quality of service goals defined by the customer. These enhancements will ensure that the unpredictable needs of new e-transaction processing workloads can be managed dynamically according to business requirements.

The current implementation addresses three separate but mutually supportive functions:

- **Dynamic Channel Path Management (DCM)**
This feature enables customers to have channel paths that dynamically and automatically move to those I/O devices that have a need for additional bandwidth. The benefits are enhanced by the use of clustered LPARS.
- **Channel Subsystem Priority Queuing**
Channel subsystem priority queuing on the z900 allows the priority queuing of I/O requests within the channel subsystem and the specification of relative priority among LPARs.
- **LPAR CPU Management**
WLM dynamically adjusts the number of logical processors online to an LPAR and the processor weight based on the WLM policy. The ability to move the CPU resources across an LPAR cluster provides processing power to where it is most needed based on WLM service policy.

Dynamic channel path management

There is no typical workload. The requirements for processor capacity, I/O capacity and other resources vary throughout the day, week, month and year. Traditional methods of monitoring and tuning result in configurations that attempt to address the majority of your workload characteristics, leaving resources under or over utilized at various times. Even a well tuned DASD subsystem can quickly become inappropriately configured as data is migrated to and from different subsystems. A better approach would involve tuning decisions being made more responsively, using the capabilities of the operating system to self-monitor and self-adjust.

DCM provides the ability to have the system automatically manage the number of paths available to disk subsystems. By making additional paths available where they are needed, this increases the effectiveness of your installed channels, and potentially reduces the number of channels required to deliver a given level of service. DCM also provides availability benefits by attempting to ensure that the paths it adds to a control unit have as few points of failure in common with existing paths as possible, and configuration management benefits by removing unused paths from over-configured control units. Where paths can be shared by Multiple Image Facility (MIF), DCM will coordinate its activities within an LPAR cluster.

Where several channels are attached from a central processor complex (CPC) to a switch, the channels can be considered as a resource pool for accessing any of the control units attached to the same switch. To achieve this without DCM would

require deactivating paths, performing a dynamic I/O reconfiguration and activating new paths. DCM achieves the equivalent process automatically, using those same mechanisms.

Channels managed by DCM are referred to as *managed* channels.

With HCD, it has to be defined how many channel paths of a control unit are to be managed. At least one channel path must be defined as static.

Value of dynamic channel path management

- **Improved system performance**

Improved system performance is achieved by automatic path balancing and service policy. Reassignment is done in real-time or on WLM adjustment interval. A customer may reassign once a year, WLM does it every few seconds.

- **More effective utilization of installed hardware**

Channels are automatically balanced, providing opportunities to use fewer I/O paths to service the same workload.

- **Simplified I/O definition**

The connection between managed channels and control units does not have to be explicitly defined, but at least one static path is required for each control unit.

- **Reduced skills required to manage z/OS**

Managed channels and control units are automatically monitored, balanced, tuned and reconfigured.

Reporting of dynamic channel path management

All RMF reports about channel activities and I/O queuing activities provide information about DCM. The following examples show Monitor III reports.

Channel Path Activity report

```

RMF V1R10 Channel Path Activity
Command ===>
Line 1 of 12
Scroll ===> HALF

Samples: 100 System: SYS1 Date: 11/28/08 Time: 15.55.00 Range: 120 Sec

Channel Path Utilization(%) Read(B/s) Write(B/s) MSG MSG Send Recv
ID No G Type S Part Tot Bus Part Tot Part Tot Rate Size Fail Fail
  8 *FC_SM 21.7 38.9 40.2 10M 24M 11M 18M
 12 *FCV_M 13.6 26.8 33.3 9M 16M 7M 12M
  1 1 *FCV_M 12.4 31.9 43.3 8M 22M 7M 19M
  1 *CNC_M 17.2 34.4
80 FC Y 22.0 43.9 44.0 11M 24M 11M 20M
81 FC_SM Y 21.7 38.9 40.2 10M 24M 11M 18M
82 FCV Y 14.7 28.5 31.1 8M 17M 6M 11M
83 FCV_M Y 13.6 26.8 33.3 9M 16M 7M 12M
85 CTC_S Y 0.1 3.9
89 CNC_S 0.7 5.8
8C CNC Y 1.7 13.5
8D CNC_S Y 1.8 13.7
DA CNC_M Y 17.2 34.4
AE IQD Y 705M 4G 923K 764 0 0
AF IQD Y 1G 1G 1M 838 0 0
B0 1 FCV_M Y 12.4 31.9 43.3 8M 22M 7M 19M

```

Figure 105. CHANNEL Report

For all channels that are managed by DCM, additional information is available. These channels are not assigned to a specific control unit permanently, 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 channel paths 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.

I/O Queuing Activity report

RMF V1R10 I/O Queuing Activity										Line 1 of 54			
Command ==>										Scroll ==> HALF			
Samples: 30		System: S5C		Date: 11/28/08		Time: 03.23.30		Range: 30 Sec					
Path	DCM	CTL	Units	DCM Group MN MX DEF	LCU	Cont Rate	De1 Lngh	Q CSS	AVG CHPID Taken	%DP Busy	%CU Busy	AVG CUB	AVG CMR
D7		5F00			0048				1.13	0.0	0.0	0.0	0.2
D6		5F00			0048				0.97	0.0	0.0	0.0	0.2
					0048	0.0	0.00	0.3	2.10	0.0	0.0	0.0	0.2
B0	PF	8000			0069				82.17	0.0	0.0	0.0	0.2
B1	PF	8000			0069				83.83	0.0	0.0	0.0	0.2
B2	NP	8000			0069				0.00	0.0	0.0	---	---
B3	NP	8000			0069				0.00	0.0	0.0	---	---
95	PF	8000			0069				83.17	0.0	0.0	0.0	0.2
					0069	0.0	0.00	0.4	249.17	0.0	0.0	0.0	0.2
B0	NP	8100			006A				0.00	0.0	0.0	---	---
B1	NP	8100			006A				0.00	0.0	0.0	---	---
B2	PF	8100			006A				124.53	0.0	0.0	0.0	0.2
B3	PF	8100			006A				124.87	0.0	0.0	0.0	0.2

Figure 106. IOQUEUE Report

The values in columns MIN MAX report the minimum and maximum number of DCM managed channels for one LCU (in this interval). DEF is the maximum number of managed channels for each LCU as it has been defined with HCD.

The line with these values is available only for LCUs with DCM managed channels. It also contains the I/O activity rate, director port contention, and control unit contention of all DCM managed channels. These values may include also measurements of managed channels which were partially online.

Channel subsystem priority queuing

Channel subsystem (CSS) priority queuing is a new function available on z900 processors either in basic or LPAR mode that allows the operating system to specify a priority value when starting an I/O request. When there is contention causing queuing in the channel subsystem, the request will be prioritized by this value.

z/OS sets the priority based on the WLM service policy. This complements the priority management that sets I/O priority for IOS UCB queues, and for queuing in the storage subsystem.

CSS priority queuing uses different priorities calculated in a different way than the I/O priorities used for UCB and control unit queuing.

Value of channel subsystem priority queuing

- **Improved performance**

I/O from work that is not meeting its goals may be prioritized ahead of I/O from work that is meeting its goals, providing the Workload Manager with an additional method of adjusting I/O performance. Channel subsystem priority queuing is complementary to UCB priority queuing and control unit priority queuing, each addressing a different mechanism that may affect I/O performance.

- **More effective use of providing I/O bandwidth**

Because interference from other I/O is reduced, the total I/O bandwidth required to meet the goals of high importance work may possibly be reduced. Previously, bandwidth had to be acquired to service the whole workload in order to meet the performance requirements of important work.

- **Reduced skills required to manage z/OS**

Monitoring and tuning requirements are reduced because of the self tuning abilities of the channel subsystem.

Reporting of channel subsystem priority queuing

There is no specific reporting of this function in RMF, but you may see a reduced number of I/O delay samples in the Workload Activity report.

LPAR CPU management

LPAR CPU management allows WLM to manage the processor weighting and number of logical processors across an LPAR cluster. CPU resources are automatically moved toward LPARs with the most need, by adjusting the partition's weight. The sum of the weights for the participants of an LPAR cluster is viewed as a pooled resource which can be apportioned among the participants to meet the goals defined in the WLM service policy. The installation can place limits on the processor weight range.

WLM also manages the available processors by varying unneeded CPs offline (more logical CPs implies more parallelism, and less weight per CP). LPAR overheads are decreased by reducing the number of logical processors.

LPAR cluster

An LPAR cluster is a group of logical partitions that are resident on the same physical server and in the same sysplex.

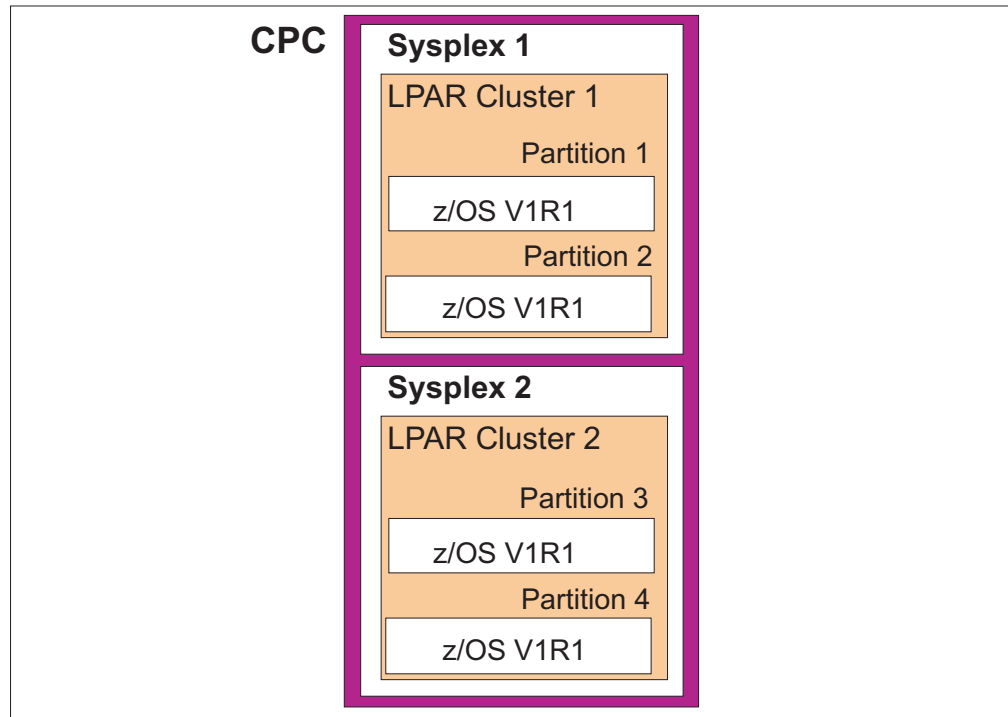


Figure 107. LPAR Cluster - Example 1

The next figure is a more complex configuration, with four LPAR clusters grouped into two sysplexes across two CPCs:

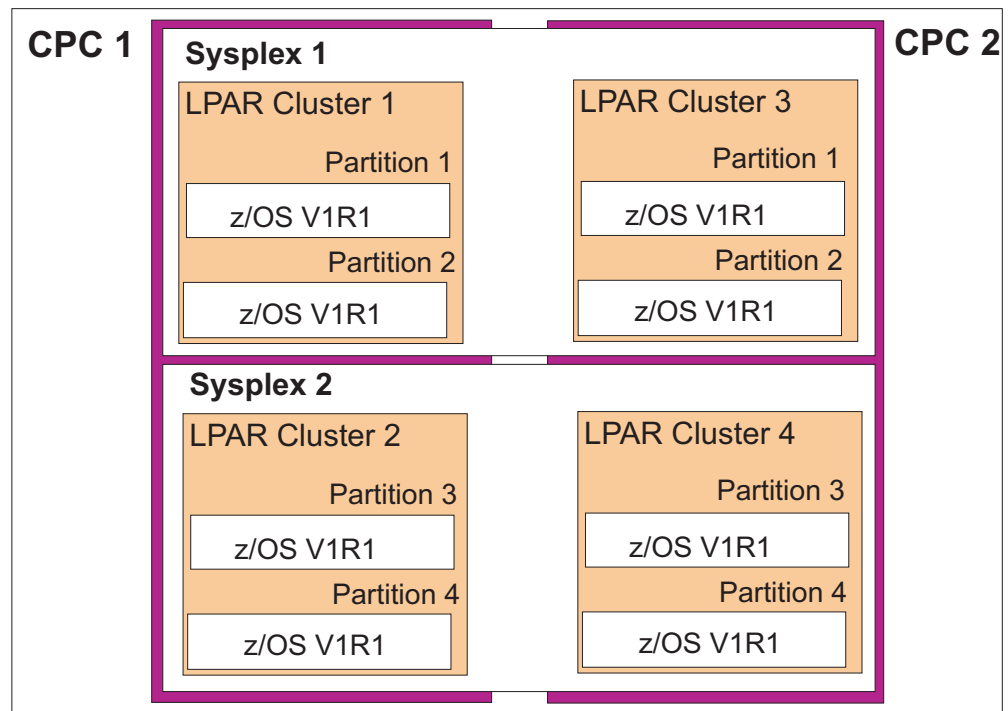


Figure 108. LPAR Cluster - Example 2

Value of LPAR CPU management

- **Logical CPs are performing at the fastest uniprocessor speed available.**

This results in the number of logical CPs tuned to the number of physical CPs of service being delivered by the LPARs current weight. If the LPAR is getting 4 equivalent physical CPs of service and has 8 logical CPs online to z/OS, then each logical CP only gets half of an equivalent physical CP. For example, if a CP delivers 200 MIPS, half of it will deliver 100 MIPS. This occurs because each logical CP gets fewer time slices.
- **Reduces LPAR overhead.**

There is an LPAR overhead for managing a logical CP. The higher the number of logical CPs in relation to the number of physical CPs, the higher the LPAR overhead. This is because LPAR has to do more processing to manage the number of logical CPs which exceeds the number of physical CPs.
- **Gives z/OS more control over how CP resources are distributed.**

z/OS is able to change the assigned CP resources (LPAR weights) and place them where they are required for the work. Making both of these adjustments is simple from an operator perspective, but what is difficult is identifying when the changes are required and whether the changes have had a positive effect.

LPAR CPU management

- identifies what changes are needed and when
- projects the likely results on both the work it is trying to help and the work that it will be taking the resources from
- performs the changes
- analyzes the results to ensure the changes have been effective.

LPAR CPU management decision controls

There are two parts of LPAR CPU management:

- WLM LPAR Weight Management
- WLM LPAR Vary CPU Management

WLM LPAR weight management

WLM LPAR weight management involves extending the receiver/donor logic for CPU delays to CP resource being used by other LPARs in the LPAR cluster. Until now, WLM could only move CP resource from one service class period to another on the same system by adjusting the dispatching priority of the receiver and/or the donor. When using WLM LPAR weight management, WLM can now alter the weights of LPARs in the LPAR cluster, effectively moving CP resource from one LPAR to another.

In the past, an operator was required to change LPAR current weights manually. However, an operator would first have to determine that there was a CPU delay and then choose which other LPAR the weight units should be taken from. This approach was not practical.

WLM LPAR vary CPU management

WLM LPAR vary CPU management does not look at service class periods when making adjustments. It ensures the number of online logical CPs is approximately equal to the number of physical CPs required to do the work. It is more like a tuning action for the CP resources available to the system. Prior to LPAR CPU management, there was no function in WLM to tune the available CP resources in relation to the physical CP resources being used by the system. Once again if this

function is to be performed manually, an operator first needs to identify this is a problem before adjusting the number of logical CPs.

LPAR CPU management controls

The IBM z900 CPC and CPU management introduce new terms for LPAR weights. The following are used when discussing LPAR weights.

- Initial processing weight
This weight becomes the LPAR’s actual weight immediately following an IPL.
- Minimum and maximum processing weight
This is the minimum and maximum weight that WLM LPAR weight management will assign as an LPAR’s actual weight.
- Actual weight
An LPAR’s weight at a point in time when the systems is running. This is the weight that WLM LPAR weight management and WLM LPAR vary CPU management use when performing their primary functions.
- Enable work load management must be checked to enable LPAR CPU management.
Recommended usage is to use large (2 or 3 digits) values for CPU weights to allow WLM a better granularity.

Reporting of LPAR CPU management

The Postprocessor CPU Activity report has been extended by a new section called LPAR Cluster report.

LPAR Cluster report

L P A R C L U S T E R R E P O R T														PAGE 3			
z/OS VIR10			SYSTEM ID NP1				DATE 11/28/08		INTERVAL 14.59.678								
			RPT VERSION VIR10 RMF				TIME 09.30.00		CYCLE 1.000 SECONDS								
----- WEIGHTING STATISTICS -----														---- PROCESSOR STATISTICS ----		----- STORAGE STATISTICS -----	
CLUSTER	PARTITION	SYSTEM	--- DEFINED ---			---- ACTUAL ----			---- NUMBER ---		-- TOTAL% --		--- CENTRAL ---	-- EXPANDED --			
			INIT	MIN	MAX	AVG	MIN %	MAX %	DEFINED	ACTUAL	LBUSY	PBUSY					
SVPLEX1	NP1	NP1	20	20	40	20	100	0.0	4	1.2	24.92	3.30	1024	N/A			
	NP2	NP2	1			1			4	4	0.73	0.27	1024	N/A			
	NP3	NP3	10	10	50	10	100	0.0	4	1.0	23.99	2.67	1024	N/A			
	TOTAL			31						12		49.64 6.24		3072 N/A			
SVPLEX2	NP4	NP4	300	100	500	300	0.0	0.0	6	4.2	80.15	53.43	1024	N/A			
	NP5	NP5	200			200			4	4	40.33	17.92	1024	N/A			
TOTAL			500						10		120.5 71.36		2048 N/A				

Figure 109. LPAR Cluster Report

This new report, as well as the Partition Data report (also part of the CPU Activity report), provides information about LPAR CPU management.

Report Analysis

- **CLUSTER PARTITION**
The three partitions NP1, NP2, and NP3 belong to the LPAR cluster SVPLEX1, but only NP1 and NP3 are under control of LPAR CPU management.
- **DEFINED INIT MIN MAX**
NP1 has an initial weight of 20 which also is the defined minimum, and it has a defined maximum weight of 40.
- **ACTUAL AVG**
The average actual weight in the interval for NP1 is 20.
- **ACTUAL MIN % MAX %**
These values provide the percentage of the interval when the partition's weight was within 10% of the defined minimum or maximum weight. NP1 was within 10% of the specified minimum (20) for the entire interval (MIN%=100).
- **NUMBER DEFINED ACTUAL**
NP1 has been defined with 4 logical processors. In the current interval, WLM LPAR CPU management has assigned 1.2 processors to this partition in average. This value can be seen also in the Partition Data report (PROCESSOR NUM) and reflects the total ONLINE TIME PERCENTAGE in the CPU Activity report.
- **TOTAL% LBUSY PBUSY**
These values give the utilization of the logical processors assigned to the partition (based on its online time) and of all physical processors of the CPC (based on the interval time). They are also shown in the Partition Data report.
- **STORAGE STATISTICS**
The amount of central and expanded storage assigned to each partition is given in these columns.

Partition Data report

PARTITION DATA REPORT																
z/OS V1R10		SYSTEM ID NP1				DATE 11/28/08				INTERVAL 14.59.678				PAGE 2		
		RPT VERSION V1R10 RMF				TIME 09.30.00				CYCLE 1.000 SECONDS						
MVS PARTITION NAME		NP1				NUMBER OF PHYSICAL PROCESSORS				9		GROUP NAME		N/A		
IMAGE CAPACITY		100				CP				7		LIMIT		N/A		
NUMBER OF CONFIGURED PARTITIONS		9				ICF				2						
WAIT COMPLETION		NO														
DISPATCH INTERVAL		DYNAMIC														
----- PARTITION DATA ----- -- LOGICAL PARTITION PROCESSOR DATA -- -- AVERAGE PROCESSOR UTILIZATION PERCENTAGES --																
NAME	S	WGT	DEF	ACT	DEF	WLM%	PROCESSOR- NUM	TYPE	DISPATCH EFFECTIVE	TIME DATA TOTAL	LOGICAL PROCESSORS EFFECTIVE	TOTAL	PHYSICAL PROCESSORS LPAR MGMT	EFFECTIVE	TOTAL	
NP1	A	20	100	10	NO	62.2	1.2	CP	00.04.27.302	00.04.27.519	24.86	24.92	0.01	4.24	4.25	
NP2	A	1	0	1	YES	0.0	4	CP	00.00.21.680	00.00.22.083	0.60	0.61	0.01	0.34	0.35	
NP3	A	10	5	8	NO	3.3	1.0	CP	00.03.35.761	00.03.35.859	23.97	23.98	0.02	3.41	3.43	
NP4	A	300	95	155	NO	0.0	6.0	CP	01.12.05.405	01.12.06.405	80.10	80.15	0.01	68.68	68.69	
NP5	A	200	50	52	NO	0.0	4	CP	00.24.11.147	00.24.11.311	40.30	40.31	0.01	23.02	23.03	
PHYSICAL										00.00.03.103		0.05		0.05		
TOTAL									01.44.41.295		01.44.46.280		0.11		99.69 99.80	
CFC1	A	DED					1	ICF	00.14.59.611	00.14.59.625	99.98	99.99	0.01	99.95	99.96	
CFC2	A	DED					1	ICF	00.00.00.000	00.00.00.000	0.00	0.00	0.00	0.00	0.00	
PHYSICAL										00.00.00.065		0.03		0.03		
TOTAL													0.04		99.95 99.99	
CB88	D															
CB89	D															

Figure 110. Partition Data Report

The Partition Data report contains in the fields WGT and PROCESSOR NUM the actual values for weighting and number of logical processors, either as dynamically assigned by LPAR CPU management (for example, NP1 or NP3) or as statically defined by the customer (for example, NP2).

Workload License Charges (WLC)

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

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

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

Report Analysis

The following values in the report are related to Workload License Charges (WLC):

- IMAGE CAPACITY

A new value IMAGE CAPACITY in this report gives information about the CPU capacity which is available to the partition. It is measured in MSUs (millions of CPU service units) per hour. There are these alternatives:

- The partition's defined capacity set via the Hardware Management Console, if any. The report shows that NP1 has a defined capacity of 100 MSUs.
- The capacity based on logical processors defined for the partition, if the partition is uncapped and has no defined capacity.

Example:

MSU of CPC	290
# physical processors	9
# logical processors	4
Image capacity	$(4/9) * 290 = 129$

- The capacity at the partition's weight, if the partition is capped via the Hardware Management Console.

- MSU DEF ACT

For each partition, the defined and actual MSU values are given. For the partition which is gathering the data, here partition NP1, DEF MSU (**100**) is identical to IMAGE CAPACITY while the actual MSU consumption in the interval is **10** MSUs (calculated as MSUs per hour). A value of zero in field DEF indicates there is no defined capacity set for this partition.

ACT MSU contains the actual consumption of service units in MSUs per hour. This value does not correlate to the value of service units given in the Workload Activity report.

- Service Units in Workload Activity Report

In the Service Policy page you find for each processor the value SU/SEC. This value is based on the number of logical processors assigned to the partition. If you have 4 logical processors in a z900 server, then this value relates to a 2064-104 (with 148 MSUs).

1 LP	= 10289 SU/SEC
4 LPs	= 41156 SU/SEC = 148 MSU/H

- Service Units in Partition Data Report

Service unit consumption in a WLC context is based on the number of physical processors in the CPC. In a 9-way z900 server 2064-109 with 290 MSUs, this value is taken:

1 CP	= 8964 SU/SEC
9 CPs	= 80676 SU/SEC = 290 MSU/H

- CAPPING WLM%

The concept of WLC is to allow a partition a limited resource consumption - the defined capacity limit. If the partition is using more service units than defined, WLM will cap this partition. This is not done directly on the current consumption but it is based on a long-term average. So peaks exceeding the defined limit are possible. The field CAPPING WLM% shows the percentage of the reporting interval when a partition has been capped by WLM.

Appendix C. Planning considerations for HiperDispatch mode

With z/OS V1R10, workload management and dispatching have been enhanced to take advantage of the IBM System z10™ hardware design. A dispatching mode called HiperDispatch has been implemented to provide additional processing efficiencies. With HiperDispatch the intent is to align work to a smaller subset of processors in order to maximize the benefits of the processor cache structures, and thereby, reduce the amount of CPU time required to execute work.

Introduction to HiperDispatch mode

For all levels of z/OS, a TCB or SRB may be dispatched on any logical processor of the type required (general purpose processor, zAAP zIIP). A unit of work starts on one logical processor and subsequently may be dispatched on any other logical processor. The logical processors for one LPAR image receive an equal share for equal access to the physical processors under PR/SM LPAR control. For example, if the weight of a logical partition with four logical processors results in a share of two physical processors, or 200%, the LPAR hypervisor manages each of the four logical processors with a 50% share of a physical processor. All logical processors are used if there is work available, and they typically have similar processing utilizations.

With HiperDispatch mode, the intention is to manage work across fewer logical processors. A new concept of maintaining a working set of processors required to handle the workload is introduced. In the previous example of a logical partition with a 200% processor share and four logical processors, two logical processors are sufficient to obtain the two physical processors worth of capacity specified by the weight; the other two logical processors allow the partition to access capacity available from other partitions with insufficient workload to consume their share.

z/OS limits the number of active logical processors to the number needed based on partition weight settings, workload demand and available capacity. z/OS also takes into account the processor topology when dispatching work, and it works with enhanced PR/SM microcode to build a strong affinity between logical processors and physical processors in the processor configuration. The logical processors for a partition in HiperDispatch mode fall into one of the following categories:

- Some of the logical processors for a partition may receive a 100% processor share, meaning this logical processor receives an LPAR target for 100% share of a physical processor. These would be viewed as having a *high* processor share. Typically, if a partition is large enough, most of the logical partition's share is allocated among logical processors with a 100% share. PR/SM LPAR establishes a strong affinity between the logical processor and a physical processor, and these processors provide optimal efficiencies in HiperDispatch mode.
- Other logical processors may have a *medium* amount of physical processor share. The logical processors would have a processor share greater than 0% and up to 100%. These medium logical processors have the remainder of the partition's share after the allocation of the logical processors with the high share. LPAR reserves at least a 50% physical processor share for the medium processor assignments, assuming the logical partition is entitled to at least that amount of service.

- Some logical processors may have a low amount, or 0%, of physical processor share. These *discretionary* logical processors are not needed to allow the partition to consume the physical processor resource associated with its weight. These logical processors may be parked. In a parked state, discretionary processors do not dispatch work; they are in a long term wait state. These logical processors are parked when they are not needed to handle the partition's workload (not enough load) or are not useful because physical capacity does not exist for PR/SM to dispatch (no available time from other logical partitions).

How RMF supports HiperDispatch mode

When examining a *CPU Activity* report for a system running in HiperDispatch mode, you may see very different processor utilizations across different logical processors of a logical partition. Figure 111 shows a report with an example of this behavior, and it also shows two columns related to HiperDispatch: TIME % PARKED and LOG PROC SHARE %.

CPU ACTIVITY									
z/OS V1R10			SYSTEM ID S59			DATE 11/28/2008		PAGE 1	
			RPT VERSION V1R10 RMF			TIME 16.45.00		INTERVAL 14.59.997	
			SEQUENCE CODE 00000000000DC6CE					CYCLE 1.000 SECONDS	
HIPERDISPATCH=YES									
---CPU---	MODEL	732	H/W MODEL	E40	SEQUENCE CODE	LOG PROC SHARE %	--I/O INTERRUPTS--		
NUM	TYPE	ONLINE	LPAR BUSY	TIME %	MVS BUSY	PARKED	RATE	% VIA TPI	
0	CP	100.00	96.33	97.34	0.00	100.0	5.80	48.75	
1	CP	100.00	95.96	97.07	0.00	100.0	4.59	55.30	
2	CP	100.00	95.79	96.84	0.00	100.0	5.10	55.18	
3	CP	100.00	95.46	96.68	0.00	100.0	2.40	53.75	
4	CP	100.00	95.08	96.41	0.00	100.0	8435	10.05	
5	CP	100.00	73.92	96.86	0.00	70.0	20.74	4.95	
6	CP	100.00	74.33	97.13	0.00	70.0	14.15	19.39	
7	CP	100.00	13.84	98.89	85.78	0.0	0.00	0.00	
TOTAL/AVERAGE			80.09	96.94		640.0	8488	10.14	
...									
...									
...									

Figure 111. CPU Activity report

In this example, the logical processor share for the partition of **640.0%** was allocated across five logical processors with a high share of 100%, two logical processors with a medium share of 70%, and one discretionary logical processor, CP 7, with a low share of 0% which was parked **85.78%** and thus unparked 14.22% of the online interval time. During this same interval, CP 7 was busy processing **13.84%** of the time.

With HIPERDISPATCH=NO, the logical processor share would be 80% for each of the 8 logical processors.

The TIME % MVS BUSY field in the CPU Activity report reflects the effective used capacity for the logical processors and the entire logical partition. The figures are calculated from the difference between online time and MVS wait time to provide an operating system perspective of busy time. Parked processors in HiperDispatch mode will generally reflect unavailable capacity at high physical processor utilizations. The formula for MVS BUSY TIME % has been changed with HiperDispatch mode to exclude Parked Time to show how busy the logical processor was when not parked. The formula is now:

$$\text{MVS BUSY TIME \%} = \frac{\text{Online Time} - (\text{Wait Time} + \text{Parked Time})}{\text{Online Time} - \text{Parked Time}} * 100$$

HiperDispatch mode changes the effect of the CPU management provided by the Intelligent Resource Director (IRD). The WLM LPAR Weight Management functions

as before. However, the Vary CPU Management is replaced by the parked / not parked aspect of discretionary processors. As with the IRD function, the initial specification for the number of logical processors in HiperDispatch mode is simple: define as many as there will probably be required. The MSU calculation associated with a logical partition does not change with the presence of parked processors.

Processing benefits

HiperDispatch can improve efficiency in both hardware and software functionality:

- Work may be dispatched across fewer logical processors, therefore reducing the “multi-processor (MP) effects” and lowering the interference among multiple partitions.
- Specific z/OS tasks may be dispatched to a small subset of logical processors which PR/SM will tie to the same physical processors thus improving the hardware cache re-use and reference performance such as reducing the rate of cross-book communication.

Therefore, the potential improvement from HiperDispatch depends from:

- the number of physical processors
- the size of the z/OS images in the configuration
- the logical to physical processor ratio
- the memory reference pattern or storage hierarchy characteristics of the workload.

Generally, a configuration where the largest z/OS image fits within a book will see minimal improvement. Workloads which are fairly CPU-intensive (like batch applications) will see only small improvements even for configurations with larger z/OS images since they typically have long-running tasks that tend to stick on a logical engine anyway. Workloads that tend to have common tasks and high dispatch rates, typical for transactional applications, may see larger improvements, again depending on the size of the z/OS images involved. LPAR configurations that are over-committed, that is, they have higher logical to physical ratios, may see some improvement although the benefit of dispatching to a reduced number of logical processors overlaps with benefits already available with IRD and various automation techniques that tend to reduce the number of online logical processors to match capacity needs.

The range of benefit may vary from 0% to 10% depending on the circumstances as described above; specifically, configurations with z/OS images small enough to fit in a book or running batch-like workloads will tend to achieve only small benefit. Multibook configurations with z/OS images in the 16way to 32way range and running transactional workloads will tend to achieve medium benefit, and very large multi-book configurations with very large z/OS images and running workloads with intense memory reference patterns will tend to achieve benefit at the high end of the range.

For further topics related to HiperDispatch (for example, an introduction to HiperDispatch or WLM policy considerations) refer to the white paper:

Planning Considerations for HiperDispatch Mode

Appendix D. Other delays

Some Other Delays

This information unit discusses the delays shown by Monitor III that have not been covered yet:

- Enqueue delays
- HSM delays
- JES delays
- OPER delays
- Unknown delays

Enqueue delays

Here is an approach to investigating ENQ delays using Monitor III.

ENQ report

```

RMF V1R10  ENQ Delays                                     Line 1 of 2
Command ==>>>                                           Scroll ==>> HALF
Samples: 100      System: PRD1  Date: 04/07/08  Time: 10.32.00  Range: 100  Sec
Jobname  DLY  ----- Resource Waiting -----  ---- Holding ----
          %   %  STAT Major/Minor Names (Scope)          % Name/SYS STAT
BHOLEQB  49   49  EW  SYSDSN (SYS)                      49 BTEUPRT
                   SYS3.BB.DATA

```

Figure 112. Monitor III Enqueue Delays Report

Look at the ENQ report for the user or user group.

Use the cursor to select the resource name with most delay and press ENTER to get the ENQR report:

- Use this report to view all the contentions.
- What is the major name contributing to most of the delay?
- Find the largest "holding" user for this resource, then go through the dialog with this jobname.

The HELP function (PF1) on the ENQJ report provides description of some different MAJOR/MINOR names. Also check the *Debugging Handbook - Volume 1* for a description of the MAJOR/MINOR names.

If the name does not exist there, it may be a local name. Knowing the resource or the application provides a greater understanding of this problem.

Major names

SYSDSN Enqueue, look for the following:

- Submitted job needs a data set that the TSO user has
- Jobs actually have to run sequentially

SYSIGGV2 / SYSCTLG Enqueue, look for the following:

- Large LISTCATs running concurrently with a DELETE or HDELETE
- Massive deletes occurring from another address space
- Allocating and deleting but not using the ISPF log data set

Try to find out what the largest "holder" of SYSIGGV2 was doing at the time of the problem. If this is not possible, and the user is a TSO user, he may have been doing massive deletes or a LISTCAT.

If the problem occurred while exiting from ISPF, the delete of the Log data set could have caused the catalog access. If so, consider making the primary and secondary allocation of the log data set zero. This causes the log data set never to be allocated when entering ISPF.

SYSVTOC Enqueue, look for the following:

Try to determine who is allocating space and why it is taking so long. Possibly a job or user is frequently acquiring and freeing the space. If this is the case, acquire it once and do not free it. If the space is for a temporary data set, try using VIO instead.

SYSZRACF Enqueue

SYSZRACF is enqueued exclusively when updating the RACF data set. Find out why it is taking so long and reduce the time.

SYSIEFSD Enqueue

SYSIEFSD has several minor names: CHNGDEVS, DDRTPUR, ALLOCTP, DDRDA, Q4, Q6, QIO, RPL, STCQUE, VARYDEV, and TSOQUE.

If the minor name is CHNGDEVS, look for a large DEVICE USING or DEVICE DELAY for the HOLDING address space. For example DEVICE DELAY may be near 100%. Obtain the dominant volumeID involved. If the time duration spans several intervals look at the volumeIDs for each interval within the problem time range. Browse the SYSLOG for mount or vary activity at or about this time. Determine if anything can be done to prevent this in the future.

HSM delays

RMF V1R10 HSM Delays				Line 1 of 3
Command ==>				Scroll ==> HALF
Samples: 100	System: MVS1	Date: 11/28/08	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 113. HSM Report

Here are some general problems to look for:

- HSM is down
- HSM is partially down (HSM device not responding)
- HSM backed up because:
 - Not enough primary space
 - Not enough level one space
 - HSM doing housekeeping

Display the HSM address space with the DELAYJ command. Is the HSM data normal? Ask yourself why, when HSM data is unusual. You can compare this data to times when HSM delays were considered acceptable. Typical things to look for are:

- PROC TCB + SRB at 0.0% (HSM may be dead)
- PROC DELAY very high (HSM may have a poor dispatching queue position)
- DEVICE DELAY or DEVICE USING excessive (see DEVR report and determine why)
- A non-typical volume is in heavy use (modules being fetched?)
- Excessive ENQ DELAY time because of CATALOG or VTOC contention

Do a HLIST PVOL TERMINAL for volumes under HSM or ISMF control. The AGE column under SPACE-MGMT is the number of days a data set on this volume must be inactive before it is eligible for the type of space management indicated under TYPE. Different volumes may have different ages. MIN AGE shows the inactive age of the data set that most recently migrated from the volume.

Pick the one of the following that seems to be most probable:

- Something is broken (HSM in loop, HSM in infinite wait, control unit down, HSM address space terminated, STOR LOCL is 100%).
- HSM address space is running slowly because of contention with others (high PROC DELAY, high STOR LOCL, high DEVICE DELAY, high ENQ DELAY).
- HSM is in house cleaning mode (HSM is reorganizing data sets to reclaim space. There will be a high amount of PROC TCB + SRB and DEVICE USING for the amount of work in progress).
- There is a shortage of primary or level-one space (for example migration age less than 7 days).
- There is excessive traffic (migration age greater than 7 but many recall requests outstanding).

JES delays

Here are some general problems to look for:

- JES is down
- JES is partially down (JES device not responding)
- JES backed up

Check the SYSLOG for more information. Look for excessive amount of I/O in the JES address space.

OPER delays

In RMF terms, the operator is another resource hence can appear as a delay.

The Monitor III Delay report has a field of OPER%. This means the job is delayed by a mount request or is waiting for a reply to a WTOR message.

Where practicable, batch jobs or applications should not issue WTORs.

Unknown delays

The Monitor III Delay report has a field of UKN%. RMF considers jobs that are not delayed for a monitored resource, or not in an idling state to be in an unknown state.

Examples of unknown state delays are:

- AS waiting for I/O other than DASD or tape
- Idle address spaces which use an unmonitored mechanism for determining when they are active. Most STCs show as unknown when they are idle.
- AS waiting for a request from another AS to be satisfied.

Other delays

Accessibility

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully. The major accessibility features in z/OS enable users to:

- Use assistive technologies such as screen readers and screen magnifier software
- Operate specific or equivalent features using only the keyboard
- Customize display attributes such as color, contrast, and font size

Using assistive technologies

Assistive technology products, such as screen readers, function with the user interfaces found in z/OS. Consult the assistive technology documentation for specific information when using such products to access z/OS interfaces.

Keyboard navigation of the user interface

Users can access z/OS user interfaces using TSO/E or ISPF. Refer to *z/OS TSO/E Primer*, *z/OS TSO/E User's Guide*, and *z/OS ISPF User's Guide Vol I* for information about accessing TSO/E and ISPF interfaces. These guides describe how to use TSO/E and ISPF, including the use of keyboard shortcuts or function keys (PF keys). Each guide includes the default settings for the PF keys and explains how to modify their functions.

z/OS information

z/OS information is accessible using screen readers with the BookServer/Library Server versions of z/OS books in the Internet library at:

<http://www.ibm.com/systems/z/os/zos/bkserv/>

Glossary

A

AS. address space

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

B

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.

BLSR. Batch LSR Subsystem. See also *LSR*

C

cache fast write. A storage control capability in which the data is written directly to cache without using nonvolatile storage. This 3990 Model 3 Storage Control extended function should be used for data of a temporary nature, or data which is readily recreated, such as the sort work files created by DFSORT. Contrast with *DASD fast write*.

cache hit. Finding a record already in the storage (cache) of the DASD control unit (3990).

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

captured storage. See shared page group.

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.

CFWHIT. Cache fast write and read hits

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

CICS. Customer Information Control System

contention. Two or more incompatible requests for the same resource. For example, contention occurs if a user requests a resource and specifies exclusive use, and another user requests the same resource, but specifies shared use.

CP. Central processor Synonymous with *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).

CRR. Cache RMF Reporter

CS. Central storage

CPU speed. Measurement of how much work your CPU can do in a certain amount of time.

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. The time at the end of which one sample is taken. Varies between 50 ms and 9999 ms. See also *sample*

D

DASD fast write. An extended function of the 3990 Model 3 Storage Control in which data is written concurrently to cache and nonvolatile storage and automatically scheduled for destaging to DASD. Both copies are retained in the storage control until the data is completely written to the DASD, providing data integrity equivalent to writing directly to the DASD. Use of DASD fast write for system-managed data sets is controlled by storage class attributes to improve performance. Contrast with *cache fast write*.

data sample. See *sample*

DCM. See *Dynamic Channel Path Management*

DCME. Dynamic cache management extended

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.

DFWHIT. DASD fast write hit

DFWRETRY. DASD fast write retry

DIM. Data-In-Memory

direct access storage device (DASD). A device in which the access time is effectively independent of the location of the data.

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.

E

EMIF. ESCON multiple image facility

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

EPDM. Enterprise Performance Data Manager/MVS, new name of the product is Performance Manager for MVS

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

ESCON multiple image facility (EMIF). A facility that allows channels to be shared among PR/SM logical partitions in an ESCON environment.

ETR. External throughput rate, see also ITR

exception reporting. So that you don't have to watch your monitor all the time to find out if there is a performance problem, you can define performance criteria. RMF will then send a report when the criteria are not being met (an exception has occurred).

G

generalized trace facility (GTF). An optional OS/VS 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 this mode, 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*.

GTF. generalized trace facility

H

high-speed buffer (HSB). A cache or a set of logically partitioned blocks that provides significantly faster access to instructions and data than provided by central storage.

HS. hiperspace

HSB. High-speed buffer

I

IFA. Integrated Facility for Applications, see zSeries Application Assist Processor (zAAP).

Integrated Facility for Applications (IFA). see zSeries Application Assist Processor (zAAP).

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.

Intelligent Resource Director (IRD). The Intelligent Resource Director (IRD) is a functional enhancement exclusive to IBM's zSeries family that can continuously and automatically reallocate resources throughout the system in response to user requirements and is based on a customer's business priorities.

ITR. Internal throughput rate, see also *ETR*.

L

Latency. The time waiting for a particular record on a track to rotate to the read/write head. This usually averages one half the rotation time.

LCU. Logical control unit. Logical control units are also called 'Control Unit Headers' (CUH). For details about LCU/CUH please refer to the applicable *System z Input/Output Configuration Program User's Guide for ICP IOCP* (SB10-7037).

local shared resources. An option for sharing I/O buffers, I/O-related control blocks, and channel programs among VSAM data sets in a resource pool that serves one partition or address space.

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.

LSPR methodology. Recommended methodology for assessing processor power.

LSR. Local shared resources

LUE. Low utilization effect

M

migration rate. The rate (pages/second) of pages being moved from expanded storage through central storage to auxiliary storage.

mintime. The smallest unit of sampling in Monitor III. Specifies a time interval during which the system is sampled. The data gatherer combines all samples gathered into a set of samples. The set of samples can be summarized and reported by the reporter.

mode. Monitor III can run in various modes: The GO mode displays a chosen report, and updates it according to the interval you have chosen, the STOP mode, which is the default mode. The GRAPHIC mode presents the reports in graphic format using the GDDM[®]* product. You can also use TABULAR mode by setting GRAPHIC OFF.

MPL. Multiprogramming level

MTW. Mean time to wait

N

NVS. Nonvolatile storage

nonvolatile storage (NVS). Additional random access electronic storage with a backup battery power source, available with a 3990 Model 3/6 Storage Control, used to retain data during a power failure. Nonvolatile storage, accessible from all storage directors, stores data during DASD fast write and dual copy operations.

P

partitioned data set (PDS). A data set in direct access storage that is divided into partitions, called members, each of which can contain a program, part of a program, or data. Synonymous with program library.

PDS. partitioned data set

peak-to-average ratio. The ratio between the highest CPU utilization and the average utilization (peak hour busy/prime shift average busy).

PER. Program event recording

performance management. (1) The activity which monitors and allocates data processing resources to applications according to goals defined in a service level agreement or other objectives. (2) The discipline that encompasses collection of performance data and tuning of resources.

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*.

program event recording (PER). A hardware feature used to assist in debugging programs by detecting and recording program events.

R

range. The time interval you choose for your report.

RMA. The 3990 RPS miss avoidance feature.

RHIT. Read hit

S

sample. Once in every cycle, the number of jobs waiting for a resource, and what job is using the resource at that moment, are gathered for all resources of a system by Monitor III. These numbers constitute one sample.

SCP. System control program

seek. The DASD arm movement to a cylinder. A seek can range from the minimum to the maximum seek time of a device. In addition, some I/O operations involve multiple imbedded seeks where the total seek time can be more than the maximum device seek time.

service 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

speed. See *workflow*

SRB. Service request block

SRM. System resource manager

SSCH. Start subchannel

sysplex. A complex consisting of a number of coupled z/OS systems.

system control program (SCP). Programming that is fundamental to the operation of the system. SCPs include z/OS, z/VM and z/VSE operating systems and any other programming that is used to operate and maintain the system. Synonymous with *operating system*.

T

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.

THRU. These are DASD writes to devices behind 3990s that are not enabled for DFW.

TP. Teleprocessing

TPNS. Teleprocessing network simulator

TSO. Time Sharing Option, see *Time Sharing Option/Extensions*

Time Sharing Option Extensions (TSO/E). In z/OS, a time-sharing system accessed from a terminal that allows user access to z/OS system services and interactive facilities.

U

UIC. Unreferenced interval count

uncaptured time. CPU time not allocated to a specific address space.

using. Jobs getting service from hardware resources (PROC or DEV) are *using* these resources.

V

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

VLF. Virtual Lookaside Facility

VTOC. Volume table of contents

W

WLM. Workload Manager

workflow. (1) The workflow of an address space represents how a job uses system resources and the speed at which the jobs moves through the system in relation to the maximum average speed at which the job could move through the system. (2) The workflow of resources indicates how efficiently users are being served.

workload. A logical group of work to be tracked, managed, and reported as a unit. Also, a logical group of service classes.

WSM. Working Set Manager

Z

zAAP. see zSeries Application Assist Processor.

zSeries Application Assist Processor (zAAP). A specialized processing assist unit configured for running Java programming on selected zSeries machines.

Notices

This information was developed for products and services offered in the U.S.A.

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

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

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

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

IBM Corporation
Mail Station P300
2455 South Road
Poughkeepsie New York 12601-5400
U.S.A.

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

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

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

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS"

WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Policy for unsupported hardware

Various z/OS elements, such as JES, DFSMS, HCD 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 support' 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.

Programming Interface Information

This book is intended to help you work with the reports RMF provides to understand your system's performance and to plan for appropriate tuning activities. The book documents information that is Diagnosis, Modification or Tuning Information provided by RMF.

Warning: Do not use this Diagnosis, Modification or Tuning Information as a programming interface.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States and/or other countries:

IBM	Parallel Sysplex
The IBM logo	PR/SM
ibm.com	Processor Resource/Systems Manager
CICS	RACF
CICS/ESA	RAMAC
DB2	RMF
DFSMS	Seascape
DFSMS/MVS	System z
DFSORT	System z9
Enterprise Storage Server	System z10
ESCON	S/390
eServer	S/390 Parallel Enterprise Server
FICON	Versatile Storage Server

GDDM	VTAM
Hiperbatch	z9
IMS	z10
MVS	z10 EC
NetView	z/OS
OS/390	zSeries

IBM, the IBM logo, ibm.com and DB2 are registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Java and all Java-based trademarks are trademarks 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.

Index

Special characters

@server zSeries Application Assist
Processor xiv

Numerics

64-bit common virtual storage
measurements xi

A

access to processor storage,
 prioritize 144
accessibility 207
analyzing workload characteristics 9
APPL% CP, rule-of-thumb 55
AUX 132
auxiliary storage 132
auxiliary storage space 141
auxiliary storage tuning 144
AVG CONN TIME, rule-of-thumb 113
AVG DISC TIME, rule-of-thumb 112
AVG HIGH UIC, rule-of-thumb 141
AVG IOSQ TIME, rule-of-thumb 110
AVG PEND TIME, rule-of-thumb 111
AVG RESP TIME, rule-of-thumb 109,
 143
AVG SLOTS ALLOCATED,
 rule-of-thumb 141
AVG SLOTS USED, rule-of-thumb 141

B

balanced systems 3
batch 37
blocked workload promotion xii
BUSY TIME PERCENTAGE,
 rule-of-thumb 45, 69
BYPASS mode 90

C

cache hit ratio, rule-of-thumb 101
Cache Subsystem Activity report 97
CACHE, what to measure 89
capacity groups xiii
capacity planning
 approaches for 4
 benefits 3
 common mistakes 3
 definition 3
 input to 55
 questions to ask 3
 resources 3
capture ratio
 computing 55
 definition 13
central storage 132
CF (coupling facility activity) 157

CFWHIT 89
Channel Path Activity report, Monitor
 I 48
Channel Subsystem Priority
 Queuing 189
CICS 55
 workload activity 146
collection, long term data 5
collection, short term data 5
components of response time 7
computing CPU time per workload
 type 14
constraint, single-CP 55
contention
 device, report 120
 for resources, Monitor III report 35
continuous monitoring of system
 status 5
control unit, logical 48
coupling facility activity report 157
CPU
 Activity report, Monitor I 45
 decreasing delay 73
 idle due to bottleneck 72
 indication of bottleneck 72
 indicators of problem 62
 measuring utilization 13
 Monitor I indicators of problem 68
 problem with 61
 problem, delay for 73
 speed, calculating 20
 time per transaction 15
 time per workload type,
 computing 14
 time, calculating 20
CPU Activity report, Monitor I 69
CPU Activity report, partition data,
 Monitor I 71, 79
CR 55
CS 132
cycle, sampling 6

D

daily monitoring, system indicators 41
DASD
 guidelines, general 107
 response time components 104
 subsystem health check 84
 system resource indicator 24
 what to measure 89
DASD Activity report, Monitor I 143
DASD guidelines 107
data collection, long term 5
data collection, short term 5
DCM 186
DCME 101
delay
 causes of 8
 enqueue 202
 finding major 35

delay (*continued*)
 HSM 204
 JES 205
 OPER 205
 report, Monitor III 38
 tuning approach 26
 unknown 205
 using Monitor III to find 35
delay due to paging 139
delay for CPU, problem 73
Delay report, Monitor III 64
demand, latent 3
DEV 41
device contention, report 120
DEVICE UTILIZATION, %,
 rule-of-thumb 114
DFWHIT 89
DFWRETRY 90
diagnosing a problem 23
DIM 78
disability 207
dispatching priority 74
distributing uncaptured time 15
DLY %, rule-of-thumb 137
DP 74
Dynamic Cache Management
 Extended 101
Dynamic Channel Path
 Management 186

E

ENQ 72
enqueue delay 202
Enterprise Storage Server 85
ESS (Enterprise Storage Server) 85
exception reporting 40
 guideline values 40
 speed 40
 Workflow/Exceptions report, Monitor
 III 66

F

FICON performance monitoring xii

G

GO mode 35
group capacity limit xiii
Group Response Time report,
 Monitor III 37
GTF 79

H

HiperDispatch xi
HiperDispatch mode
 introduction 197

HiperDispatch mode (*continued*)
 processing benefits 199
 support in RMF 198
hit ratio, low, cache 101
HSM delay 204
HyperPAV support xii

I

I/O

I/O Queuing Activity report, Monitor I 49
 indicators for contention 82
 rate of workload 16
 resolving specific problems 118
IBM System z9 Integrated Information Processor xiii
IFA xiv, 210
IMS
 workload activity 146
indicator
 sysplex monitoring, Monitor III 28
indicators of problems
 CPU problem 62
 for storage problems 37
 of overcommitted tape buffers 50
 processor storage, Monitor I 132
 processor storage, Monitor III 132, 133
 speed 40
 swap 53
 workflow 40
INHIBIT mode 90
Integrated Facility for Applications xiv
Intelligent Resource Director
 Channel Subsystem Priority Queuing 189
 Dynamic Channel Path Management 185
 LPAR CPU Management 190
 Overview 185
interval 35
 reporting 6
IRD (Intelligent Resource Director) 185
ITR 21

J

JES delay 205
Job Delays report, Monitor III 39, 135

K

keyboard 207

L

latent demand 3
LCU 48
LCU AVG RESP TIME, rule-of-thumb 50
LOG SWAP/EXP STOR EFFECTIVE, rule-of-thumb 53
logical control unit 48
logical partition 47
long term data collection 5

LookAt message retrieval tool viii
loops 76
low hit ratio, cache 101
LP 47
LPAR 47
 constraints, what to do 75
 CPU constraints 47
 Monitor I Channel Path Activity report 48
LPAR Cluster 190
LPAR CPU Management 190

M

measuring
 CPU utilization 13
 online 5
 resource consumption 11
message retrieval tool, LookAt viii
MFLOPS 21
MIPS 20
Monitor III processor storage indicators 133
Monitor I reports
 Channel Path Activity report 48
 CPU Activity report 45, 69
 CPU Activity report, partition data 47, 71, 79
 DASD Activity report 50, 108, 143
 I/O Queuing Activity report 49
 Magnetic Tape Device Activity report 127
 Page/Swap Data Set Activity report 54, 141
 Paging Activity report 53, 138, 141
Monitor II reports
 Address Space State Data report 74
Monitor III reports
 Delay report 38, 64, 119
 Device Delays report 120
 Device Resource Delays report 121
 Group Response Time report 37, 63
 Job Delays report 39, 73, 77, 135
 Processor Delays report 66, 76
 Storage Delays report 38, 137
 System Information report 118, 134
 Workflow/Exceptions report 40, 66, 78
monitoring of system status, continuous 5

O

online measurements 5
OPER delay 205
OUT READY, rule-of-thumb 45
out-ready delay, indicator 134

P

PAGE MOVEMENT WITHIN CENTRAL STORAGE, rule-of-thumb 138
PAGE-IN EVENTS, rule-of-thumb 139
PAGE-IN RATES, rule-of-thumb 55
page/swap 143

Page/Swap Data Set Activity report, Monitor I 54, 141
PAGES XFER'D, rule-of-thumb 54
paging 55
 paging activity 19
 Paging Activity report 141
 Paging Activity report, Monitor I 53, 138
 paging delay 134, 139
 parallel access volumes 87
 partition data report section, Monitor I CPU Activity report 47
 partition, logical 47
PAV 87
peak-to-average ratio 3, 45
PER 78
performance management
 approaches for 4
 definition 2
performance problem
 auxiliary storage tuning 144
 CPU 61
 definition 24
 delay for CPU 73
 diagnosing 23
 finding major delay 35
 I/O, resolving specific problems 118
 indicators of CPU problem 62
 investigating 25
 Monitor I indicators of CPU problem 68
 processor storage 131
 resolving 25
PR/SM LPAR considerations 179
prioritize access to processor storage 144
problem indicators
 CPU problem 62
 for storage problems 37
 of overcommitted tape buffers 50
 processor storage, Monitor I 132
 processor storage, Monitor III 132, 133
 speed 40
 swap 53
 workflow 40
PROC 41
Processor Delays report, Monitor III 66
processor storage
 indicators of 131
 Monitor I 132, 138
 Monitor III 132, 133
 prioritizing access to 144
 problem 131
 use by workload 18
Program Event Recording 78
promoting blocked workload xii

R

range 35
rate, SSCH 17
ratio
 capture 13
 capture, computing 55
 ratio, low hit, cache 101
 ratio, peak-to-average 3
 recommendations 144
reconfigurable storage 138

- Relative Processor Power 21
- report on device contention 120
- reporting interval 6
- reports, using Monitor III 36
- resolving specific I/O problems 118
- resource consumption, measuring 11
- response time
 - batch 37
 - components of 7
 - computing 7
 - end-to-end 35
 - external response time 35
 - finding major delay 35
 - internal response time 35
 - Monitor I Workload Activity report 7
 - Monitor III Group Response Time report 7
 - TSO 37
 - where to start 35
- Response Time Distribution report
 - description 26
- RHIT 89
- RMF
 - monitors 5
 - tasks 5
- RPP 21
- RSU 138
- rule-of-thumb
 - % Delayed for STR 135
 - % IN USE 54, 142
 - APPL% CP 55
 - AVG CONN TIME 113
 - AVG DISC TIME 112
 - AVG HIGH UIC 141
 - AVG IOSQ TIME 110
 - AVG PEND TIME 111
 - AVG RESP TIME 109, 143
 - AVG SLOTS ALLOCATED 141
 - AVG SLOTS USED 141
 - BUSY TIME PERCENTAGE 45
 - BUSY TIME PERCENTAGE, CPU 69
 - cache hit ratio 101
 - DASD guidelines 107
 - DEVICE UTILIZATION, % 114
 - DLY % 137
 - LCU AVG RESP TIME 50
 - LOG SWAP/EXP STOR
 - EFFECTIVE 53
 - OUT READY 45
 - PAGE MOVEMENT WITHIN
 - CENTRAL STORAGE 138
 - PAGE-IN EVENTS 139
 - PAGE-IN RATES 55
 - PAGES XFER'D 54
 - STOR 134

S

- sampling cycle 6
- service level agreement
 - contents of 2
 - definition 2
- Service Level Agreement 2
- short term data collection 5
- shortcut keys 207
- single-CP constraint 55
- SLA 2, 41

- SLIP 78
- slip trap 78
- space, auxiliary storage 141
- speed 40
- SRB 55
- SSCH rate 17
- STAGE 89
- STOR 41
- STOR, rule-of-thumb 134
- storage
 - See also* storage problem
 - auxiliary 132
 - central 132
 - delay report, Monitor III 137
 - reconfigurable 138
 - space, auxiliary 141
 - tuning, auxiliary 144
- Storage Delays report, Monitor III 38
- storage problem
 - indicators for 37
 - processor 131
- storage use by workload, processor 18
- swapping delay, indicator 134
- swaps, TSO 53
- Sysplex Summary report
 - description 26
- system indicators 26
- system indicators, daily monitoring 41
- System Information report, Monitor III 134
- system status, continuous monitoring
 - of 5
- systems, balanced 3

T

- TCB 79
- throughput
 - external throughput rate 21
 - internal throughput rate 21
- THRU 90
- time per transaction, CPU 15
- time, distributing uncaptured 15
- tracing 79
- transaction
 - computing CPU time 15
 - definition 2
 - response time, components of 7
- TSO 37
- TSO swaps 53
- tuning 24
 - top-down approach 25
 - where to begin 5
- types of workload 9

U

- uncaptured time, distributing 15
- unknown delay 205
- use by workload, processor storage 18
- using Monitor III reports 36
- utilization, measuring CPU 13

V

- VIO, delay, indicator 134

W

- WLC 195
- WLM blocked workload promotion xii
- workflow 40
- Workflow/Exceptions report, Monitor III 40, 66
- workload
 - analyzing 9
 - computing CPU time per type 14
 - grouping of 9
 - I/O rate of 16
 - measuring CPU utilization 13
 - measuring resource utilization 10
 - paging rate of 19
 - processor storage use by 18
 - types of 9
- Workload Activity report 27, 55
- Workload License Charges 195
- workload type, compute CPU time per 14

X

- XES structure activity 158

Z

- z9 EC xiii
- zAAP xiv, 212
- zIIP xiii
- zSeries Application Assist Processor 212

Readers' Comments — We'd Like to Hear from You

z/OS
Resource Measurement Facility
Performance Management Guide
Version 1 Release 10

Publication No. SC33-7992-08

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Comments:

Thank you for your support.

Submit your comments using one of these channels:

- Send your comments to the address on the reverse side of this form.
- Send a fax to the following number: FAX (Germany): 07031+16-3456
FAX (Other Countries): (+49)+7031-16-3456
- Send your comments via e-mail to: s390id@de.ibm.com
- Send a note from the web page: WORLD WIDE WEB: <http://www.ibm.com/servers/eserver/zseries/zos/rmf>

If you would like a response from IBM, please fill in the following information:

Name

Address

Company or Organization

Phone No.

E-mail address



Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Deutschland Research & Development GmbH
Department 3248
Schoenaicher Strasse 220
D-71032 Boeblingen
Federal Republic of Germany



Fold and Tape

Please do not staple

Fold and Tape



Program Number: 5694-A01

Printed in USA

SC33-7992-08



Spine information:



z/OS

**z/OS V1R10.0 RMF Performance Management
Guide**

Version 1
Release 10

SC33-7992-08