

z/OS



# MVS Planning: Workload Management



z/OS



# MVS Planning: Workload Management

**Note**

Before using this information and the product it supports, be sure to read the general information under Appendix C, "Notices" on page 233.

**Fifth Edition, October 2002**

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

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

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

IBM Deutschland Entwicklung GmbH

Department 3248

Schoenaicher Strasse 220

D-71032 Boeblingen

Federal Republic of Germany

FAX (Germany): 07031+16-3456

FAX (Other Countries): (+49)+7031-16-3456

Internet e-mail: [s390id@de.ibm.com](mailto:s390id@de.ibm.com)

World Wide Web: <http://www.ibm.com/servers/eserver/zseries/zos/webqs.html>

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 1994, 2002. All rights reserved.

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

---

# Contents

<b>Figures</b> . . . . .	<b>vii</b>
--------------------------	------------

<b>Tables</b> . . . . .	<b>ix</b>
-------------------------	-----------

## About This Document . . . . . **xi**

Who Should Use This Document . . . . .	xi
Where to Find More Information . . . . .	xi
Other Referenced Documents . . . . .	xii
The WLM/SRM Web Page . . . . .	xiii
Accessing z/OS™ licensed documents on the Internet . . . . .	xiii
Using LookAt to look up message explanations	xiii

## Summary of Changes . . . . . **xv**

## Chapter 1. Why MVS Workload Management? . . . . . **1**

Problems . . . . .	1
MVS Workload Management Solution for Today and Tomorrow . . . . .	2

## Chapter 2. What is MVS Workload Management? . . . . . **3**

Performance Administration . . . . .	4
Performance Management . . . . .	5
Workload Balancing . . . . .	5
Workload Management Concepts . . . . .	6
What is a Service Definition? . . . . .	6
Why Use Service Policies? . . . . .	7
Organizing Work into Workloads and Service Classes . . . . .	7
Why Use Resource Groups? . . . . .	9
Assigning Work to a Service Class . . . . .	10
Why Use Application Environments? . . . . .	11
Why Use Scheduling Environments? . . . . .	13
Summary of Service Definition and Service Policy Concepts . . . . .	14

## Chapter 3. Workload Management Participants . . . . . **17**

Workload Management Work Environments . . . . .	17
Subsystem Support for Goal Types and Multiple Periods . . . . .	18
Subsystem-Specific Performance Hints . . . . .	19
Workload Balancing . . . . .	19
Workload Management in a CPSM Environment	20
Workload Management in a DB2 Distributed Data Facility (DDF) Environment . . . . .	20
Batch Workload Management . . . . .	21
Multisystem Enclave Support . . . . .	22
Intelligent Resource Director . . . . .	22
Non z/OS Partition CPU Management . . . . .	22

Workload Management, IBM License Manager, and Workload License Charges . . . . .	23
Workload Management with Other Products . . . . .	24

## Chapter 4. Setting up a Service Definition . . . . . **27**

Specifying a Service Definition . . . . .	27
Storing Service Definitions . . . . .	28
Defining the Parts of a Service Definition . . . . .	29

## Chapter 5. Defining Service Policies . . . . . **31**

Using Policy Overrides . . . . .	32
----------------------------------	----

## Chapter 6. Defining Workloads . . . . . **35**

Defining a Departmental Workload . . . . .	35
--	----

## Chapter 7. Defining Resource Groups **37**

## Chapter 8. Defining Service Classes and Performance Goals . . . . . **39**

Velocity Formula . . . . .	42
Defining Performance Goals . . . . .	42
Determining System Response Time Goals . . . . .	43
Examples . . . . .	44
Defining Velocity Goals . . . . .	44
Adjusting Velocity Goals Based on Samples Included in Velocity Calculation . . . . .	45
Using Velocity Goals for Started Tasks . . . . .	45
Using Discretionary Goals . . . . .	46
Using Performance Periods . . . . .	47
Defining Goals Appropriate for Performance Periods . . . . .	47
Using Importance Levels in Performance Periods	47

## Chapter 9. Defining Classification Rules . . . . . **49**

Defining Classification Rules for Each Subsystem . . . . .	51
Defining Work Qualifiers . . . . .	57
Defining the Order of Classification Rules . . . . .	68
Defining a Subsystem Service Class Default . . . . .	68
Organizing Work for Classification . . . . .	73
Using the Start Position . . . . .	73
Using Masking Notation . . . . .	74
Using Wildcarding Notation . . . . .	74
Using Groups . . . . .	75
Using the System Supplied Service Classes . . . . .	76

## Chapter 10. Defining Report Classes **81**

## Chapter 11. Defining Service Coefficients and Options . . . . . **83**

Service Definition Coefficients . . . . .	83
Changing Your Coefficient Values . . . . .	83

Specifying I/O Priority Management . . . . .	84
Considerations for I/O Priority Management . . . . .	84
Specifying Dynamic Alias Management . . . . .	85
Workload Management Considerations . . . . .	85
HCD Considerations . . . . .	85

**Chapter 12. Defining Special Protection Options for Critical Work . . . . . 89**

Long-Term Storage Protection . . . . .	89
Long-Term CPU Protection . . . . .	90
Exemption from Transaction Response Time Management . . . . .	91
Sample Scenarios . . . . .	91
Scenario 1 . . . . .	92
Scenario 2 . . . . .	93
Scenario 3 . . . . .	95
Scenario 4 . . . . .	96
Scenario 5 . . . . .	97
Reporting . . . . .	98
Option Summary . . . . .	99

**Chapter 13. Defining Application Environments . . . . . 101**

Getting Started with Application Environments . . . . .	101
Specifying Application Environments to Workload Management . . . . .	102
Selecting Server Limits for Application Environments . . . . .	104
Using Application Environments . . . . .	106
Managing Application Environments . . . . .	107
Using Operator Commands for Application Environments . . . . .	107
Making Changes to the Application Environment Servers . . . . .	108
Changing the Definition of an Application Environment . . . . .	108
Handling Error Conditions in Application Environments . . . . .	109
Authorizing Application Environment Servers . . . . .	109

**Chapter 14. Defining Scheduling Environments . . . . . 113**

Getting Started with Scheduling Environments . . . . .	113
Specifying Scheduling Environments to Workload Management . . . . .	114
Managing Resource States . . . . .	117
Associating Scheduling Environments with Incoming Work . . . . .	121
Displaying Information about Scheduling Environments and Resource States . . . . .	121
MVS Operator Commands . . . . .	121
JES2/JES3 Operator Commands . . . . .	123
SDSF Commands . . . . .	123

**Chapter 15. Worksheets for Defining a Service Definition . . . . . 125**

**Chapter 16. Workload Management Migration . . . . . 131**

Creating a Service Definition for the First Time . . . . .	131
Migrating to a New z/OS Release With an Existing Service Definition . . . . .	135
Migration Activities . . . . .	138
Restricting Access to the WLM Service Definition . . . . .	138
Start the Application and Enter/Edit the Service Definition . . . . .	139
Calculate the Size of the WLM Couple Data Set . . . . .	145
Allocate a WLM Couple Data Set . . . . .	145
Make a WLM Couple Data Set Available to the Sysplex for the First Time . . . . .	149
Make a Newly Formatted Couple Data Set Available to the Sysplex . . . . .	150
Migration Considerations for Velocity . . . . .	151
Migration Considerations for Discretionary Goal Management . . . . .	152
Migration Considerations for Dynamic Alias Management . . . . .	152
Migration Considerations for Multisystem Enclaves . . . . .	152
Migration Considerations for Protection of Critical Work . . . . .	153

**Chapter 17. Defining a Coupling Facility Structure for Multisystem Enclave Support . . . . . 155**

Defining the Coupling Facility . . . . .	155
Shutting Down the Coupling Facility . . . . .	158
Coupling Facility Failures . . . . .	158

**Chapter 18. The Intelligent Resource Director . . . . . 159**

LPAR CPU Management . . . . .	161
Dynamic Channel Path Management . . . . .	161
Channel Subsystem Priority Queuing . . . . .	162
Example . . . . .	163
Making the Intelligent Resource Director Work . . . . .	164
Defining the SYSZWLM_xxxxxxxx Coupling Facility Structure . . . . .	165
Enabling LPAR CPU Management . . . . .	166
Enabling Non z/OS CPU Management . . . . .	167
Enabling Dynamic Channel Path Management . . . . .	167
Enabling Channel Subsystem Priority Queuing . . . . .	169
Handling Downlevel Releases . . . . .	169
For More Information . . . . .	169

**Chapter 19. Using the WLM ISPF Application . . . . . 171**

Before You Start . . . . .	171
Panel Areas and How to Use Them . . . . .	171
Using the Menu Bar . . . . .	172
Using the Status Line . . . . .	173
Using the Scrollable Area . . . . .	173
Using the Action Field . . . . .	174
Using the Command Line . . . . .	174
Using the Function Keys . . . . .	175
Starting the WLM Application . . . . .	175
Now You're Started . . . . .	176

Using the Definition Menu . . . . .	176
Using the Menu Bar on the Definition Menu	178
Working with Service Policies . . . . .	181
Working with Workloads . . . . .	182
Working with Resource Groups . . . . .	183
Working with Service Classes . . . . .	183
Defining Goals . . . . .	184
Using Action Codes on Service Class Panels . . . . .	185
Defining Service Policy Overrides . . . . .	186
Working with Classification Rules . . . . .	188
Using Action Codes on the Modify Rules Panel	190
Using Selection Lists for Classification Rules	192
Creating a Subsystem Type for Rules . . . . .	192
Deleting a Subsystem Type for Rules . . . . .	193
Working with Classification Groups . . . . .	193
Working with Report Classes . . . . .	194
Working with Service Coefficients and Options . . . . .	195
Working with Application Environments . . . . .	196
Working with Scheduling Environments . . . . .	196
Creating a New Scheduling Environment . . . . .	197
Modifying a Scheduling Environment . . . . .	199
Copying a Scheduling Environment . . . . .	200
Browsing a Scheduling Environment . . . . .	200
Printing a Scheduling Environment . . . . .	201
Deleting a Scheduling Environment . . . . .	201
Creating a New Resource . . . . .	201
Showing All Cross-References for a Resource Definition . . . . .	205
Deleting a Resource . . . . .	206
Coordinating Updates to a Service Definition. . . . .	207
Using the WLM Couple Data Set . . . . .	207
Using MVS PDSs . . . . .	210
Restricting Access to Your Service Definition . . . . .	210
Activating a Service Policy . . . . .	210
Printing in the Application . . . . .	211
Browsing Definitions . . . . .	211
Using XREF Function to View Service Definition Relationships . . . . .	212

WLM Application Messages . . . . .	213
------------------------------------	-----

**Appendix A. Customizing the WLM ISPF Application . . . . . 217**

Introduction . . . . .	217
Specifying the Exits . . . . .	217
Coding the WLM Exits . . . . .	218
IWMARIN1 . . . . .	218
Customizing the WLM Application Libraries -	
IWMAREX1 . . . . .	218
Processing . . . . .	219
Parameters . . . . .	219
Example . . . . .	220
Customizing the WLM Application Data Sets -	
IWMAREX2 . . . . .	220
Processing . . . . .	220
Parameters . . . . .	220
Examples . . . . .	222
Adding WLM as an ISPF Menu Option. . . . .	222
Moving Pop-up Windows . . . . .	223
Customizing the Keylists . . . . .	224

**Appendix B. CPU Capacity Table . . . 225**

Using SMF Task Time . . . . .	231
-------------------------------	-----

**Appendix C. Notices . . . . . 233**

Trademarks . . . . .	235
----------------------	-----

**Appendix D. Accessibility . . . . . 237**

Using assistive technologies . . . . .	237
Keyboard navigation of the user interface . . . . .	237

**Workload Management Terms . . . . 239**

**Index . . . . . 243**



## Figures

1. MVS Workload Management Overview . . . . .	4	42. Override Service Class Selection List . . . . .	187
2. Workload Organized by Subsystem . . . . .	8	43. Override Attributes for a Service Class . . . . .	187
3. Workload Organized by Department . . . . .	9	44. Subsystem Type Selection List for Rules . . . . .	188
4. Resource Groups . . . . .	10	45. Modify Rules for the Subsystem Type . . . . .	189
5. Work Classification . . . . .	11	46. Modify Rules for the Subsystem Type, Scrolled Right to Description Fields . . . . .	189
6. Application Environment Example . . . . .	12	47. Modify Rules for the Subsystem Type, Scrolled Right to Storage Critical and Region Management Fields . . . . .	190
7. Scheduling Environment Example . . . . .	13	48. Action Codes for Classification Rules . . . . .	191
8. Service Definition, Including Two Service Policies . . . . .	15	49. Create a Group . . . . .	193
9. Sysplex View of the Management Environment . . . . .	18	50. Modify Rules for STC Subsystem . . . . .	194
10. Workload Management in a CICS Environment . . . . .	20	51. Creating a Report Class on the Modify Rules Panel . . . . .	194
11. Workload Management Enforcing a Defined Capacity Limit . . . . .	24	52. Create a Report Class Confirmation Panel . . . . .	195
12. Using Classification Rules to Assign Work to Service Classes . . . . .	50	53. Service Coefficients Panel . . . . .	196
13. Specifying the "Storage Critical" Option . . . . .	89	54. Application Environments . . . . .	196
14. Specifying the "Storage Critical" Option . . . . .	90	55. Decide What to Create Panel . . . . .	197
15. Specifying the "Manage Region Using Goals Of:" Option . . . . .	91	56. Scheduling Environment Selection List Panel . . . . .	197
16. Scenario 1: Address Spaces . . . . .	92	57. Create a Scheduling Environment Panel . . . . .	198
17. Scenarios 2,3,4,5: CICS/IMS Regions . . . . .	93	58. Resource Definition List Panel . . . . .	198
18. Sample Systems and Scheduling Environments . . . . .	120	59. Create a Scheduling Environment Panel . . . . .	199
19. One LPAR Cluster on One CPC . . . . .	159	60. Scheduling Environment Selection List Panel . . . . .	199
20. Four LPAR Clusters on Two CPCs . . . . .	160	61. Modify a Scheduling Environment Panel . . . . .	200
21. Intelligent Resource Director Example – Day Shift . . . . .	163	62. Copy a Scheduling Environment Panel . . . . .	200
22. Intelligent Resource Director Example – Night Shift . . . . .	164	63. Browse a Scheduling Environment Panel . . . . .	201
23. Menu Bar on the Definition Menu . . . . .	172	64. Delete a Scheduling Environment Panel . . . . .	201
24. Definition Menu File Choices . . . . .	172	65. Scheduling Environment Selection List Panel . . . . .	202
25. Service Class Selection List . . . . .	173	66. Resource Definition List Panel . . . . .	202
26. Service Class Selection List . . . . .	174	67. Define Resource Panel . . . . .	202
27. Action field on the Subsystem Type Selection List . . . . .	174	68. Resource Definition List Panel . . . . .	203
28. Function Key Area . . . . .	175	69. Modify a Scheduling Environment Panel . . . . .	203
29. Choose Service Definition Pop-Up . . . . .	176	70. Resource Definition List Panel . . . . .	204
30. Definition Menu . . . . .	177	71. Resource Definition List Panel . . . . .	204
31. Create a Service Policy . . . . .	181	72. Modify a Scheduling Environment Panel . . . . .	204
32. Service Policy Selection List . . . . .	181	73. Resource Definition List Panel . . . . .	206
33. Create a Workload . . . . .	182	74. Resource Cross-Reference Of Scheduling Environments Panel . . . . .	206
34. Workload Selection List . . . . .	182	75. Resource Definition List Panel . . . . .	206
35. Create a Resource Group . . . . .	183	76. Overwrite service definition. . . . .	207
36. Create a Service Class . . . . .	184	77. Allocate Couple Data Set Using CDS Values . . . . .	208
37. Choose a Goal Type Pop-Up . . . . .	184	78. Allocate Couple Data Set Using Service Definition Values . . . . .	209
38. Average Response Time Goal Pop-Up . . . . .	185	79. Activating a Service Policy . . . . .	211
39. Create a Service Class Panel . . . . .	185	80. Browse Function from the Service Class Selection List . . . . .	212
40. Action Codes for Goal . . . . .	185	81. Service Class Subsystem Xref . . . . .	212
41. Service Policy Selection List . . . . .	186	82. Keylist Utility Panel . . . . .	224



---

## Tables

1. IBM-Defined Subsystem Types . . . . .	52	9. Application Environment Server	
2. Non-IBM Subsystems . . . . .	54	Characteristics . . . . .	105
3. Enclave Transactions, Address Space-Oriented		10. Functionality Levels for Service Definition	140
Transactions, and CICS/IMS Transactions . . .	55	11. WLM Couple Data Set Format Levels	143
4. Qualifiers Supported by Each IBM-Defined		12. Workload Management Compatibility	144
Subsystem Type . . . . .	55	13. Mapping of WLM Couple Data Set	
5. Recommended Values for I/O Priority		Subrecords to IXCL1DSU Keywords . . . . .	148
Management . . . . .	85	14. Values to Use in Storage Estimation Formulas	156
6. Effects of WLMPAV Settings on Base and Alias		15. Menu Bar Options on the Definition Menu	179
Devices . . . . .	86	16. Return Codes from IWMARIN0 . . . . .	217
7. Summary of Storage Protection, CPU		17. Processors: zSeries 900 . . . . .	225
Protection, and Exemption from Transaction		18. Processors: S/390 9672, 2003, 3000, and 7060	225
Response Time Management Options . . . . .	99	19. Processors: ES/9000 9021, 9121, and 9221	229
8. IBM-supplied Subsystems Using Application			
Environments . . . . .	102		



---

## About This Document

This manual supports z/OS (5694–A01) and z/OS.e (5655–G52). This document contains information to help you convert to MVS workload management, to use workload management, and to make the most out of workload management.

**Note:** The MVS workload management component is sometimes also called “Workload Manager.”

---

## Who Should Use This Document

This document is intended for the system programmers, system analysts, and systems engineers who are responsible for developing a conversion plan for MVS workload management, and who are responsible for implementing MVS workload management.

---

## Where to Find More Information

Where necessary, this document references information in other documents, using shortened versions of the document title. For complete titles and order numbers of the documents for all products that are part of z/OS, see *z/OS Information Roadmap*.

Short Title Used in This Document	Title	Order Number
<i>z/OS MVS Initialization and Tuning Guide</i>	<i>z/OS MVS Initialization and Tuning Guide</i>	SA22-7591
<i>z/OS MVS Initialization and Tuning Reference</i>	<i>z/OS MVS Initialization and Tuning Reference</i>	SA22-7592
<i>z/OS MVS Programming: Workload Management Services</i>	<i>z/OS MVS Programming: Workload Management Services</i>	SA22-7619
<i>z/OS MVS JCL Reference</i>	<i>z/OS MVS JCL Reference</i>	SA22-7597
<i>z/OS MVS JCL User's Guide</i>	<i>z/OS MVS JCL User's Guide</i>	SA22-7598
<i>z/OS MVS Migration</i>	<i>z/OS MVS Migration</i>	GA22-7580
<i>z/OS and z/OS.e Planning for Installation</i>	<i>z/OS and z/OS.e Planning for Installation</i>	GA22-7504
<i>z/OS MVS System Management Facilities (SMF)</i>	<i>z/OS MVS System Management Facilities (SMF)</i>	SA22-7630
<i>z/OS MVS Setting Up a Sysplex</i>	<i>z/OS MVS Setting Up a Sysplex</i>	SA22-7625
<i>z/OS MVS System Commands</i>	<i>z/OS MVS System Commands</i>	SA22-7627

## Other Referenced Documents

Softcopy Title	Title	Order Number
<i>z/OS RMF Report Analysis</i>	<i>z/OS RMF Report Analysis</i>	SC33-7991
<i>z/OS MVS Planning: APPC/MVS Management</i>	<i>z/OS MVS Planning: APPC/MVS Management</i>	SA22-7599
<i>DB2 Administration Guide</i>	<i>DB2 Administration Guide</i>	SC26-9931
<i>DB2 SQL Reference</i>	<i>DB2 SQL Reference</i>	SC26-9944
<i>DB2 for OS/390 and z/OS Administration Guide</i>	<i>DB2 for OS/390 and z/OS Administration Guide</i>	SC26-9931
<i>DB2 for OS/390 and z/OS Data Sharing: Planning and Administration</i>	<i>DB2 for OS/390 and z/OS Data Sharing: Planning and Administration</i>	SC26-9935
<i>z/OS JES2 Commands</i>	<i>z/OS JES2 Commands</i>	SA22-7526
<i>z/OS JES2 Initialization and Tuning Guide</i>	<i>z/OS JES2 Initialization and Tuning Guide</i>	SA22-7532
<i>z/OS JES2 Initialization and Tuning Reference</i>	<i>z/OS JES2 Initialization and Tuning Reference</i>	SA22-7533
<i>z/OS JES2 Installation Exits</i>	<i>z/OS JES2 Installation Exits</i>	SA22-7534
<i>z/OS JES3 Initialization and Tuning Guide</i>	<i>z/OS JES3 Initialization and Tuning Guide</i>	SA22-7549
<i>z/OS JES3 Initialization and Tuning Reference</i>	<i>z/OS JES3 Initialization and Tuning Reference</i>	SA22-7550
<i>ISPF Dialog Management Guide and Reference</i>	<i>ISPF Dialog Management Guide and Reference</i>	SC34-4266
<i>CICS/ESA Customization Guide</i>	<i>CICS/ESA Customization Guide</i>	SC33-1165
<i>CICS/ESA Dynamic Transaction Routing in a CICSplex</i>	<i>CICS/ESA Dynamic Transaction Routing in a CICSplex</i>	SC33-1012
<i>CICS/ESA Performance Guide</i>	<i>CICS/ESA Performance Guide</i>	SC33-1183
<i>CICS/ESA Resource Definition Guide</i>	<i>CICS/ESA Resource Definition Guide</i>	SC33-1166
<i>CICS/ESA CICS-RACF Security Guide</i>	<i>CICS/ESA CICS-RACF Security Guide</i>	SC33-1185
<i>CICSplex SM Managing Workloads</i>	<i>CICSplex SM Managing Workloads</i>	SC33-1807
<i>IMS/ESA Installation Volume 1 and Volume 2</i>	<i>IMS/ESA Installation Volume 1 and Volume 2</i>	SC26-8736 SC26-8737
<i>IMS/ESA Administration Guide: System</i>	<i>IMS/ESA Administration Guide: System</i>	SC26-8730
<i>z/OS UNIX System Services Programming Tools</i>	<i>z/OS UNIX System Services Programming Tools</i>	SA22-7805
<i>z/OS UNIX System Services Planning</i>	<i>z/OS UNIX System Services Planning</i>	GA22-7800
<i>Internet Connection Server User's Guide</i>	<i>Internet Connection Server User's Guide</i>	SC31-8204
<i>z/OS HTTP Server Planning, Installing, and Using</i>	<i>z/OS HTTP Server Planning, Installing, and Using</i>	SC34-4826
<i>MQSeries Workflow for OS/390 System Management Guide</i>	<i>MQSeries Workflow for OS/390 System Management Guide</i>	SC34-5374
<i>OS/390 SOMobjects Configuration and Administration Guide</i>	<i>OS/390 SOMobjects Configuration and Administration Guide</i>	GC28-1851
<i>z/OS UNIX System Services Parallel Environment: Operation and Use</i>	<i>z/OS UNIX System Services Parallel Environment: Operation and Use</i>	SA22-7810
<i>SAP R/3 on DB2 for OS/390: Planning Guide</i>	<i>SAP R/3 on DB2 for OS/390: Planning Guide</i>	SC33-7964
<i>z/OS TSO/E Customization</i>	<i>z/OS TSO/E Customization</i>	SA22-7783
<i>PR/SM Planning Guide</i>	<i>PR/SM Planning Guide</i>	GA22-7236
<i>Support Element Operations Guide</i>	<i>Support Element Operations Guide</i>	SC28-6802
<i>DCF SCRIPT/VS User's Guide</i>	<i>DCF SCRIPT/VS User's Guide</i>	S544-3191

---

## The WLM/SRM Web Page

For the latest information, see the WLM/SRM web page at:

<http://www.ibm.com/servers/eserver/zseries/zos/wlm/>

## Accessing z/OS™ licensed documents on the Internet

z/OS licensed documentation is available on the Internet in PDF format at the IBM Resource Link™ Web site at:

<http://www.ibm.com/servers/resourceLink>

Licensed documents are available only to customers with a z/OS license. Access to these documents requires an IBM Resource Link user ID and password, and a key code. With your z/OS order you received a Memo to Licensees, (GI10-0671), that includes this key code.<sup>1</sup>

To obtain your IBM Resource Link user ID and password, log on to:

<http://www.ibm.com/servers/resourceLink>

To register for access to the z/OS licensed documents:

1. Sign in to Resource Link using your Resource Link user ID and password.
2. Select **User Profiles** located on the left-hand navigation bar.

**Note:** You cannot access the z/OS licensed documents unless you have registered for access to them and received an e-mail confirmation informing you that your request has been processed.

Printed licensed documents are not available from IBM.

You can use the PDF format on either **z/OS Licensed Product Library CD-ROM** or IBM Resource Link to print licensed documents.

## Using LookAt to look up message explanations

LookAt is an online facility that allows you to look up explanations for most 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 access LookAt from the Internet at:

<http://www.ibm.com/eserver/zseries/zos/bkserv/lookat/>

or from anywhere in z/OS where you can access a TSO/E command line (for example, TSO/E prompt, ISPF, z/OS UNIX System Services running OMVS). You can also download code from the *z/OS Collection* (SK3T-4269) and the LookAt Web site that will allow you to access LookAt from a handheld computer (Palm Pilot VIIx suggested).

To use LookAt as a TSO/E command, you must have LookAt installed on your host system. You can obtain the LookAt code for TSO/E from a disk on your *z/OS Collection* (SK3T-4269) or from the **News** section on the LookAt Web site.

---

1. z/OS.e™ customers received a Memo to Licensees, (GI10-0684) that includes this key code.

Some messages have information in more than one document. For those messages, LookAt displays a list of documents in which the message appears.

---

## Summary of Changes

### Summary of Changes for SA22-7602-04 z/OS Version 1 Release 4

The book contains information previously presented in *z/OS MVS Planning : Workload Management* which supports z/OS Version 1 Release 3..

This book contains terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change. Technical changes include:

- Information is added to indicate this book supports z/OS.e.

Starting with z/OS V1R2, you may notice changes in the style and structure of some content in this book--for example, headings that use uppercase for the first letter of initial words only, and procedures that have a different look and format. The changes are ongoing improvements to the consistency and retrievability of information in our books.

### Summary of Changes for SA22-7602-03 z/OS Version 1 Release 3

The book contains information previously presented in *z/OS MVS Planning : Workload Management* which supports z/OS Version 1 Release 2..

This book contains terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change. Technical changes include:

- Removal of compatibility mode.
- Addition of an Install Definition Utility.
- An example service definition and JCL job used by the Utility, shipped in SYS1.SAMPLIB, as members IWMSSDEF and IWMINSTL. See Chapter 16, "Workload Management Migration" on page 131.

Starting with z/OS V1R2, you may notice changes in the style and structure of some content in this book--for example, headings that use uppercase for the first letter of initial words only, and procedures that have a different look and format. The changes are ongoing improvements to the consistency and retrievability of information in our books.

### Summary of Changes for SA22-7602-02 z/OS Version 1 Release 2

The book contains information previously presented in *z/OS MVS Planning : Workload Management* which supports z/OS Version 1 Release 1.

Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change. Technical changes include:

- Chapter 3, “Workload Management Participants” on page 17 and Chapter 18, “The Intelligent Resource Director” on page 159 are updated to show that WLM can now manage operating systems such as LINUX running in a separate logical partition in an LPAR cluster. This function will be shipped disabled in z/OS V1R2. A PTF for APAR OW50221 is needed to enable this support. This PTF will be provided by 12/01.
- Chapter 10, “Defining Report Classes” on page 81 is enhanced to show the effects of heterogeneous report classes.
- Chapter 16, “Workload Management Migration” on page 131 is updated to include migration to z/OS Version 1 Release 2.

You may notice changes in the style and structure of some content in this book—for example, headings that use uppercase for the first letter of initial words only, and procedures that have a different look and format. The changes are ongoing improvements to the consistency and retrievability of information in our books.

**Summary of Changes  
for SA22-7602-01  
as Updated June, 2001**

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

**Summary of Changes  
for SA22-7602-00  
z/OS Version 1 Release 1**

The book contains information also presented in *OS/390 MVS Planning : Workload Management*.

Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change. Technical changes include:

- Chapter 18, “The Intelligent Resource Director” on page 159 is added.
- Chapter 16, “Workload Management Migration” on page 131 is updated to include migration to z/OS Version 1 Release 1.

---

## Chapter 1. Why MVS Workload Management?

Before the introduction of MVS workload management, MVS required you to translate your data processing goals from high-level objectives about what work needs to be done into the extremely technical terms that the system can understand. This translation requires high skill-level staff, and can be protracted, error-prone, and eventually in conflict with the original business goals. Multi-system, sysplex, parallel processing, and data sharing environments add to the complexity.

MVS workload management provides a solution for managing workload distribution, workload balancing, and distributing resources to competing workloads. MVS workload management is the combined cooperation of various subsystems (CICS, IMS/ESA, JES, APPC, TSO/E, OS/390 UNIX System Services, DDF, DB2, SOM, LSFM, and Internet Connection Server) with the MVS workload management (WLM) component.

This chapter identifies the problems that led to workload management, and describes the high level objectives for the long term.

---

### Problems

The problems that needed to be solved by workload management include:

- ***System Externals:***

Performance and tuning externals are scattered between MVS and various subsystems, as well as throughout monitoring and reporting products. This is true for the OPT parameters, and, prior to z/OS 1.3, for the MVS installation performance specifications (IPS) and the installation control specifications (ICS). MVS and the subsystems each have their own terminology for similar concepts, each have their own controls, and the controls are not coordinated.

- ***Expectations:***

Many of the MVS externals are geared towards implementation. You tell MVS how to process work, not your expectations of how well it should run work. There is no single way to make sure that important work is getting the necessary system resources.

- ***Integration and Feedback:***

The multiple monitoring and reporting products show many different views of how well MVS is doing, or how well individual subsystems are managing work. Since there is no single way to specify performance goals for your installation, it is also difficult to get feedback from monitors and reporters on how well your installation actually achieved what you expected it to. There is little sense of which reports and fields relate to the externals you specified.

- ***Managing towards expectations:***

Some installations configure their systems to handle a peak load. This could result in inefficient use of expensive resources. Today's externals do not allow you to reflect your performance expectations for work. Today, a system programmer must completely understand the implications of changing a parameter before he or she can configure an installation to achieve performance objectives. Since there is no direct path between specification and expectation, it is difficult to predict the effects of changing any one control or parameter.

Given this mix of problems, workload management has some high level objectives to provide a solution.

---

## MVS Workload Management Solution for Today and Tomorrow

Workload management requires a shift of focus from tuning at a system resources level to defining performance expectations. This requires a basic shift in philosophy towards goal-oriented systems management. The complete workload management solution offers a shift in philosophy in the following areas:

- ***Fewer, simpler, and consistent system externals:***

Workload management provides a way to define MVS externals and tune MVS without having to specify low-level parameters. The focus is on setting performance goals for work, and letting the workload manager handle processing to meet the goals.

- ***Externals reflect customer expectations:***

Workload management provides new MVS performance management externals in a *service policy* that reflects goals for work, expressed in terms commonly used in service level agreements (SLA). Because the terms are similar to those commonly used in an SLA, you can communicate with end-users, with business partners, and with MVS using the same terminology.

- ***Expectations-to-feedback correlation:***

With one common terminology, workload management provides feedback to support performance reporting, analysis, and modelling tools. The feedback describes performance achievements in the same terms as those used to express goals for work.

- ***“System Managed” toward expectations:***

Workload management provides automatic work and resource management support that dynamically adapts as needed. It manages the trade-offs between meeting your service goals for work and making efficient use of system resources.

- ***Sysplex externals and management scope:***

Workload management eliminates the need to micro-manage each individual MVS image, providing a way for you to increase the number of systems in your installation without greatly increasing the necessary skill level.

---

## Chapter 2. What is MVS Workload Management?

Installations today process different types of work with different response times. Every installation wants to make the best use of its resources and maintain the highest possible throughput and achieve the best possible system responsiveness.

With workload management, you define performance goals and assign a business importance to each goal. You define the goals for work in business terms, and the system decides how much resource, such as CPU and storage, should be given to the work to meet its goal.

An installation should know what it expects to accomplish in the form of performance goals, as well as how important it is to the business that each performance goal be achieved. With workload management, you define performance goals for work, and the system matches resources to the work to meet those goals, constantly monitoring and adapting processing to meet the goals. Reporting reflects how well the system is doing compared to its goals.

Figure 1 on page 4 shows a high level overview of the workload management philosophy. The following section explains the performance administration and performance management processes.

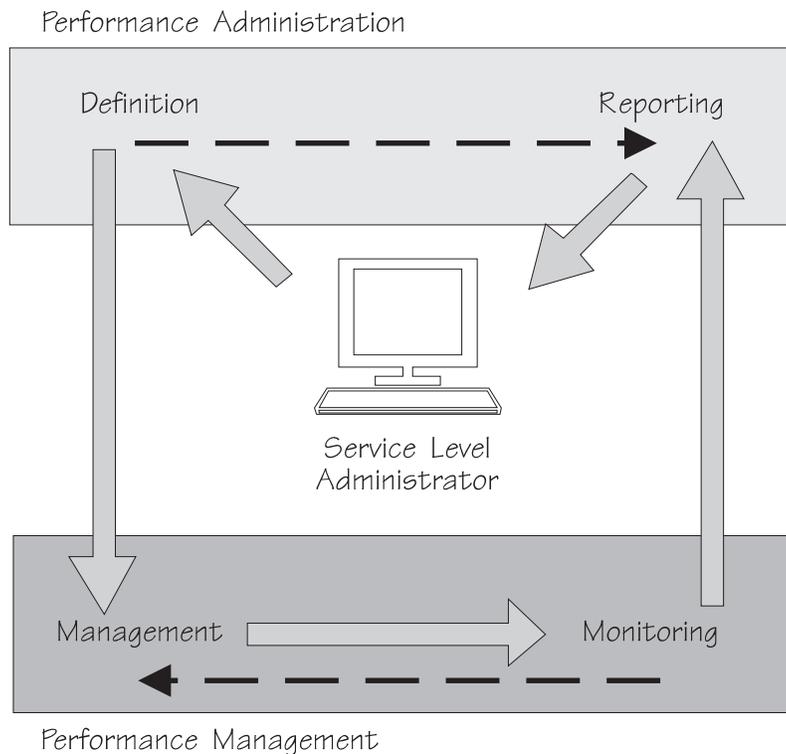


Figure 1. MVS Workload Management Overview

## Performance Administration

Performance administration is the process of defining and adjusting performance goals. Workload management introduces the role of the service level administrator. The service level administrator is responsible for defining the installation's performance goals based on business needs and current performance. This explicit definition of workloads and performance goals is called a *service definition*. Some installations might already have this kind of information in a service level agreement (SLA). The service definition applies to all types of work, CICS, IMS, TSO/E, OS/390 UNIX System Services, JES, APPC/MVS, LSFM, DDF, DB2, SOM, Internet Connection Server (also referred to as IWEB in this book) and others. You can specify goals for all MVS managed work, whether it is online transactions or batch jobs. The goals defined in the service definition apply to all work in the sysplex.

Because the service definition terminology is similar to the terminology found in an SLA, the service level administrator can communicate with the installation user community, with upper level management, and with MVS using the same terminology. When the service level requirements change, the service level administrator can adjust the corresponding workload management terms, without having to convert them into low-level MVS parameters.

Workload management provides an online panel-based application for setting up and adjusting the service definition. You specify the service definition through this ISPF administrative application.

Workload management provides the capability to collect performance and delay data in context of the service definition. The performance and delay data are available to reporting and monitoring products, so that they can use the same terminology.

---

## Performance Management

Performance management is the process workload management uses to decide how to match resources to work according to performance goals. Workload management algorithms use the service definition information and internal monitoring feedback to check how well they are doing in meeting the goals. The algorithms periodically adjust the allocation of resource as the workload level changes.

For each system, workload management handles the system resources. Workload management coordinates and shares performance information across the sysplex. How well it manages one system is based on how well the other systems are also doing in meeting the goals. If there is contention for resources, workload management makes the appropriate trade-offs based on the importance of the work and how well the goals are being met.

Workload management can dynamically start and stop server address spaces to process work from application environments. Workload management starts and stops server address spaces in a single system or across the sysplex to meet the work's goals.

You can turn over management of batch initiators to workload management, allowing workload management to dynamically manage the number of batch initiators for one or more job classes to meet the performance goals of the work.

In addition to internal feedback monitoring, workload management keeps track of what is happening in the sysplex in the form of real time performance data collection, and delay monitoring. All this information is available for performance monitors and reporters for integration into detailed reports.

---

## Workload Balancing

To make the most of workload management, work needs to be properly distributed so that MVS can manage the resources. It is essential that the subsystems distributing work are configured properly for workload distribution in a sysplex. You do this with the controls provided by each subsystem. For example, in a JES2 and JES3 environment, you spread initiator address spaces across each system.

Initial cooperation between MVS and the transaction managers (CICS, IMS, DB2) allows you to define performance goals for all types of MVS-managed work. Workload management dynamically matches resources (access to the processor and storage) to work to meet the goals.

CICS, however, goes further with the CICSplex Systems Manager (CICSplex SM) to dynamically route CICS transactions to meet the performance goals. CPSM monitors the performance of CICS resources and systems and presents the

resources as though they were part of a single system. This type of cooperation greatly improves CICS transaction workload balancing.

Other subsystems also have automatic and dynamic work balancing in a sysplex. For example, DB2 can spread distributed data facility (DDF) work across a sysplex automatically. DB2 can also distribute work in a sysplex through its sysplex query parallelism function. CICS, TSO, and APPC cooperate with VTAM and workload management in a sysplex to balance the placement of sessions. SOMobjects can automatically spread its servers across a sysplex to meet performance goals and to balance the work.

For more detail on workload management in different subsystem environments, see Chapter 3, “Workload Management Participants” on page 17.

---

## Workload Management Concepts

The service definition contains all the information about the installation needed for workload management processing. There is one service definition for the entire sysplex. The service level administrator specifies the service definition through the WLM administrative application. The service level administrator sets up “policies” within the service definition to specify the goals for work. A service level administrator must understand how to organize work, and be able to assign it performance objectives.

### What is a Service Definition?

A service definition consists of:

- One or more *service policies*, which are named sets of overrides to the goals in the service definition. When a policy is activated, the overrides are merged with the service definition. You can have different policies to specify goals intended for different times. Service policies are activated by an operator command, or through the ISPF administrative application utility function.
- *Workloads*, which aggregate a set of service classes together for reporting purposes.
- *Service classes*, which are subdivided into periods, group work with similar performance goals, business importance, and resource requirements for management and reporting purposes. You assign performance goals to the periods within a service class.
- *Report classes*, which group work for reporting purposes. They are commonly used to provide more granular reporting for subsets of work within a single service class.
- *Resource groups*, which define processor capacity boundaries across the sysplex. You can assign a minimum and maximum amount of CPU service units per second to work by assigning a service class to a resource group.
- *Classification rules*, which determine how to assign incoming work to a service class and report class.
- *Application environments*, which are groups of application functions that execute in server address spaces and can be requested by a client. Workload management manages the work according to the defined goal, and automatically starts and stops server address spaces as needed.
- *Scheduling environments*, which are lists of resource names along with their required states. If an MVS image satisfies all of the requirements in a scheduling environment, then units of work associated with that scheduling environment can be assigned to that MVS image.

The following section explains each of the concepts.

## Why Use Service Policies?

A *service policy* is a named set of overrides to the performance goals and processing capacity boundaries in the service definition. A policy applies to all of the work running in a sysplex. Because processing requirements change at different times, service level requirements may change at different times. If you have performance goals that apply to different times or a business need to limit access to processor capacity at certain times, you can define multiple policies. In order to start workload management processing, you must define at least one service policy. You can activate only one policy at a time.

## Organizing Work into Workloads and Service Classes

To workload management, work is a demand for service, such as a batch job, an APPC, CICS, DB2, or IMS transaction, a TSO/E logon, a TSO/E command, or a SOM request. All work running in the installation is divided into *workloads*. Your installations may already have a concept of workload. A workload is a group of work that is meaningful for an installation to monitor. For example, all the work created by a development group could be a workload, or all the work started by an application, or in a subsystem.

Within a workload, you group work with similar performance characteristics into *service classes*. You create a service class for a group of work with similar:

- Performance goals
- Resource requirements
- Business importance to the installation

### Performance Goals

You assign a performance goal to each service class period, such as a response time goal, and you indicate an importance. *Importance* is how important it is to your business that the performance goal be achieved.

There are three kinds of goals: *response-time*, *execution velocity*, and *discretionary*. Response time goals indicate how quickly you want your work to be processed. Since response time goals are not appropriate for all kinds of work, such as long running batch jobs, there are execution velocity goals.

Execution velocity goals define how fast work should run when ready, without being delayed for processor, storage, I/O access, and queue delay. Execution velocity goals are intended for work for which response time goals are not appropriate, such as started tasks, or long running batch work. Discretionary goals are for low priority work for which you do not have any particular performance goal. Workload management then processes the work using resources not required to meet the goals of other service classes.

### Resource Requirements

Because some work may have variable resource requirements, you can define multiple periods for a service class. Periods are a way of defining different goals for work depending on the amount of resources the work consumes. Typically, periods are used to give shorter transactions more aggressive goals and to give longer running work of the same type less aggressive goals. If you have multiple periods, each period except the last has a duration. *Duration* is the amount of resources, in service units, that the work consumes before it is switched to the goals of the next period.

You can also group work into a service class based on resource requirements. If you have a group of batch work that can consume vast amounts of resources, and you want to limit it, you can define a service class and assign it to a resource group with a maximum amount of capacity. If the work exceeds that capacity, workload management slows the execution rate. Also, if a certain group of work needs a minimum amount of processor capacity, you can set up a service class and assign it to a resource group.

### Business Importance

When there is not sufficient capacity for all work in the system to meet its goals, business importance is used to determine which work should give up resources and which work should receive more. You assign an *importance* to a service class period, which indicates how important it is that the goal be met relative to other goals. Importance plays a role only when a service class period is not meeting its goal. There are five levels of importance: lowest (5), low, medium, high, and highest (1).

Figure 2 shows an example of the relationship between workloads and service classes. Work in the figure is represented by different size triangles. The IMS workload in the figure, represents all of the IMS work. There are three single-period service classes set up, each with a different importance, and a different response time goal.

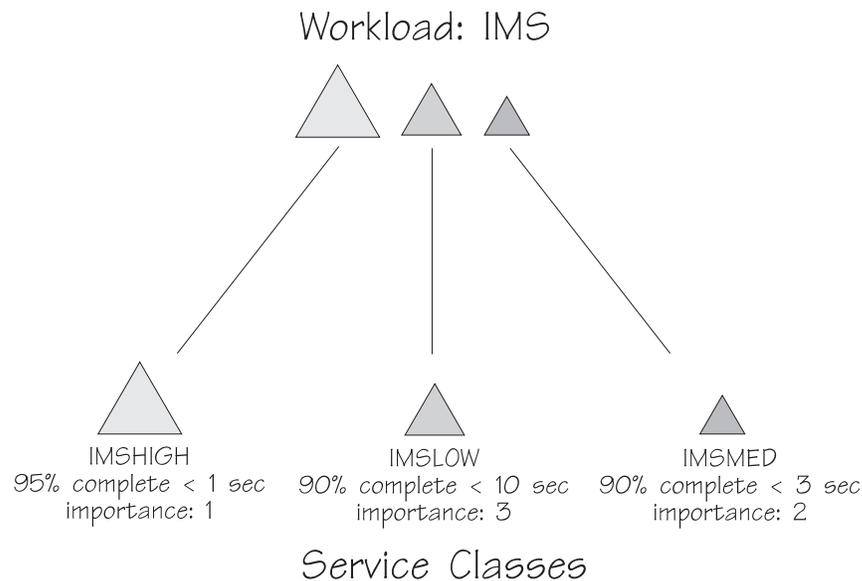


Figure 2. Workload Organized by Subsystem

Figure 3 on page 9 shows an example of a workload that is organized by division. In the figure, work is represented by different shapes: circles, squares, and triangles. The workload, OFFICE, represents all of the work in the office division. There are three service classes set up, each a different kind of work in the OFFICE workload. IMSTEST represents the IMS test work, CICS represents all of the CICS work, and JES represents all of the batch work in the OFFICE workload. Each of the service classes has one period with a different response time goal assigned to it.

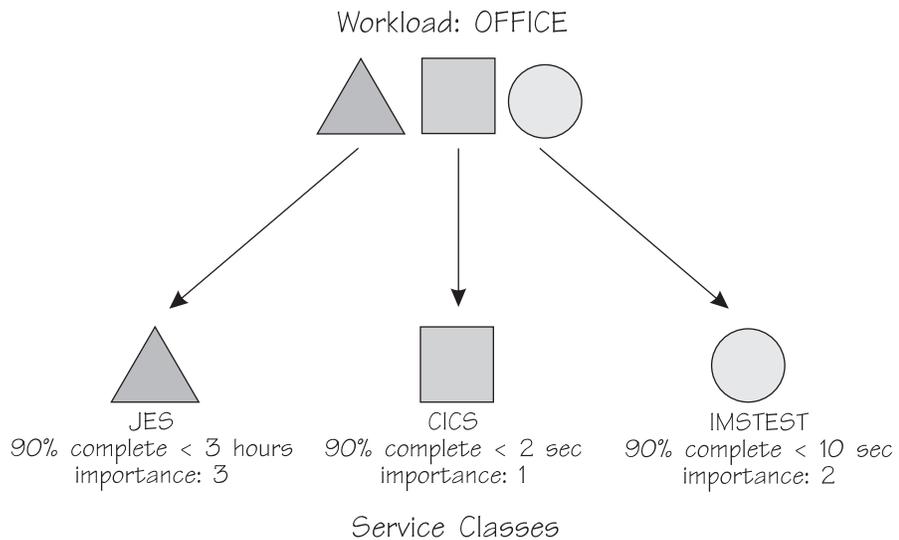


Figure 3. Workload Organized by Department

## Why Use Resource Groups?

Resource groups are a way of limiting or guaranteeing resource capacity. A resource group is a named amount of CPU capacity that you can assign to one or more service classes. For most systems, you can let workload management decide how to manage the resources in the sysplex and not use resource groups. You set performance goals for work and let workload management adjust to meet the goals.

There are some cases, however, where you might want to use a resource group. For example, if you are a service bureau and you “sell” processing capacity, you may need to guarantee or limit the amount of capacity the system delivers to a group of users. Otherwise, if you have any other need to limit capacity, you should use the existing methods: LPAR, initiators, and others. For example, you may have an SLA that charges for a fixed amount of CPU capacity and want to ensure that no more is used. Or, you may want to ensure that some work with discretionary goals receives some minimum amount of capacity. For both cases, a resource group is appropriate. For other cases, it is probably not, so you can skip it!

For a resource group, you specify either a minimum or a maximum amount of sysplex capacity in unweighted CPU service units per second, or both.

Figure 4 on page 10 shows some examples of resource groups.

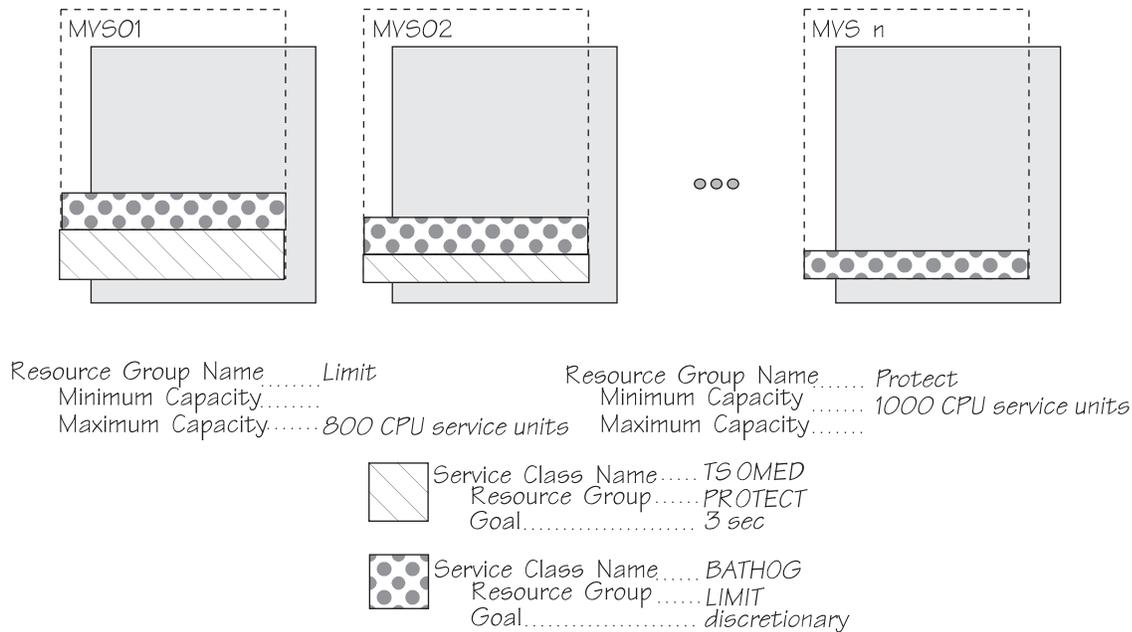


Figure 4. Resource Groups

In the figure, the service class BATHOG is assigned to the resource group LIMIT. BATHOG might include work that consumes processing capacity in huge amounts, so you assign it to a resource group, and limit it to a maximum of 800 CPU service units per second. Also, the goal assigned to BATHOG is discretionary.

The service class TSOMED, on the other hand, is associated with the resource group PROTECT, because according to an SLA, it is contracted to be guaranteed a minimum amount of processing capacity. So the resource group PROTECT is assigned a minimum of 1000 CPU service units per second. Then, when there is sufficient TSOMED work in the system, and TSOMED is underachieving its goal, TSOMED is guaranteed at least 1000 CPU service units per second. Note that if TSOMED is overachieving its goal, then the minimum capacity setting has no effect. Workload management then processes work from both service classes, making sure that TSOMED gets its minimum amount of capacity, and making sure BATHOG does not consume more than its assigned maximum. Keep in mind the service class goal: if BATHOG is assigned a stringent goal, the goal may never be achievable within the LIMIT resource group capacity. You should determine whether the resource group capacity fulfills your purpose, or the goal. Using both controls could introduce some conflict, and the resource group controls will prevail.

## Assigning Work to a Service Class

Classification rules are the rules workload management uses to associate a performance goal and/or a resource group with work by associating incoming work with a service class. Optionally, classification rules can also associate incoming work with a report class, similar to a report performance group.

The classification rules for a subsystem are specified in terms of transaction qualifiers such as job name or transaction class. These qualifiers identify groups of transactions that should have the same performance goals and importance. The

attributes of incoming work are compared to these qualifiers and, if there is a match, the rule is used to assign a service class to the work. A subsystem can also have a default class for work that does not match any of the rules.

Figure 5 shows how work classification rules associate a service class, and optionally a report class with incoming work.

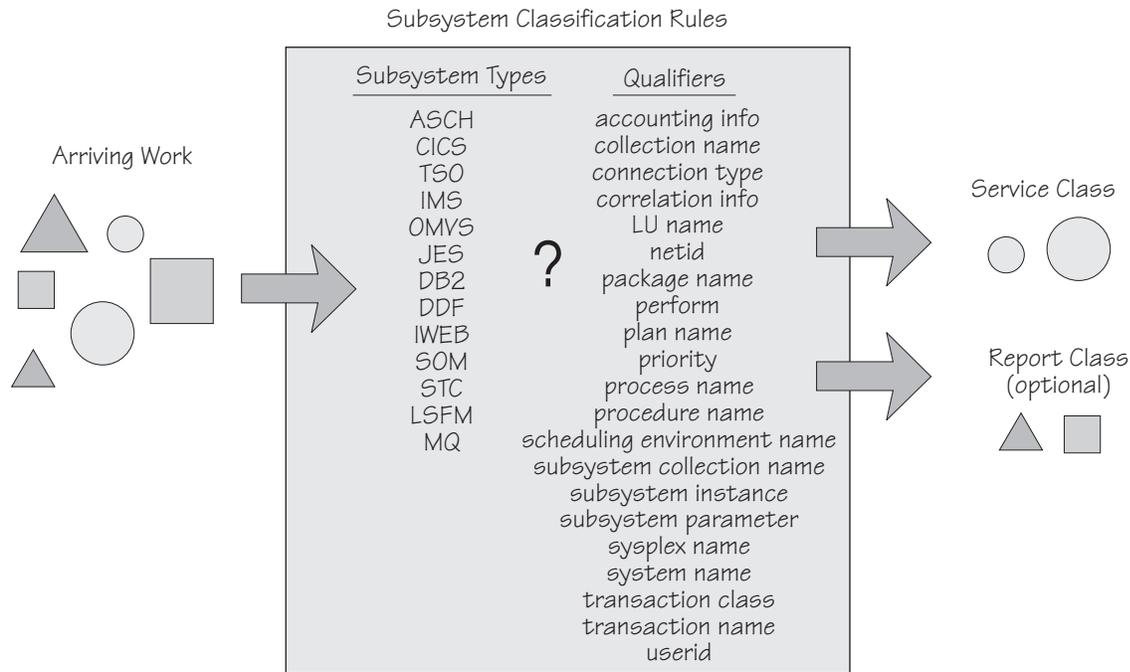


Figure 5. Work Classification

Optionally, classification rules can assign incoming work to a report class. You get reporting information in terms of workloads, service classes, and report classes. Use report classes for obtaining data on a subset of work within a service class, or for rolling up data across workloads.

## Why Use Application Environments?

An *application environment* is a way to group similar server programs together and have workload management dynamically create and delete server address spaces as needed to handle the work. Each application environment typically represents a named group of server functions that require access to the same application libraries. Depending on the subsystem's implementation of application environments, the scope of server address space management is either confined to a single system or is sysplex-wide. There is also an option to manually start and stop the server address spaces for an application environment if a customer has a special or temporary requirement to control the number of servers independently of workload management.

If you are using a subsystem that takes advantage of application environments, you need to refer to the subsystem documentation for guidance on how to use them for that subsystem. For a list of the IBM-supplied subsystems currently using application environments, see Chapter 13, "Defining Application Environments" on page 101.

Figure 6 shows an example of how 2 application environments named AE1 and AE2 can be used to handle work requests from a subsystem. In this example, 3 types of work requests named X, Y, and Z are handled. X and Y might be two different kinds of payroll inquiries, and Z might be a loan inquiry.

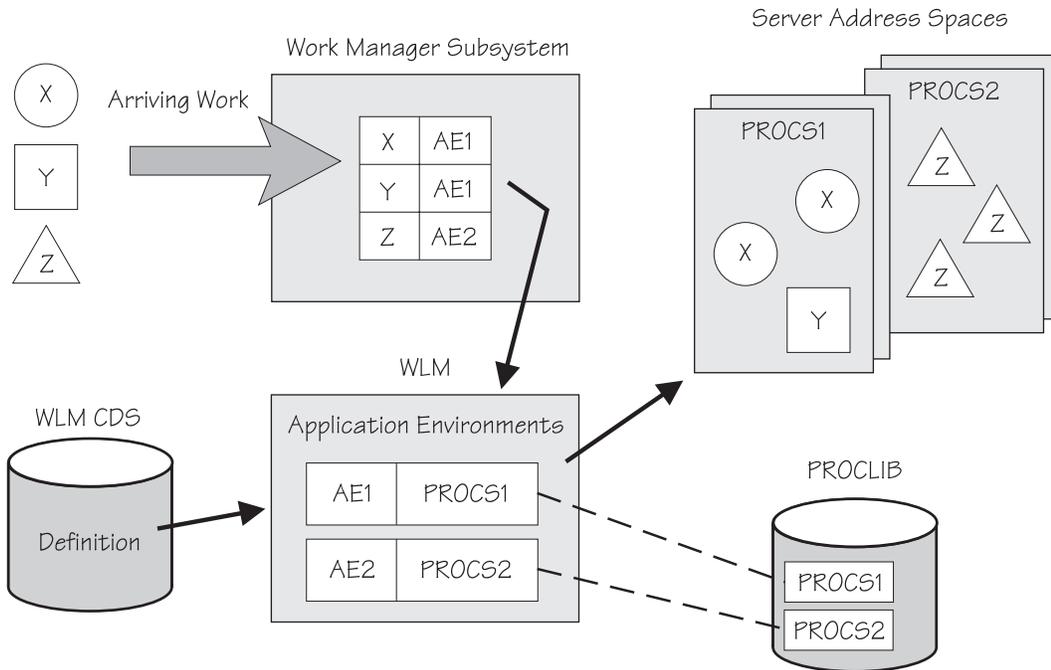


Figure 6. Application Environment Example

The names X, Y, and Z are used by clients when making the work requests. The work manager subsystem contains a table or file that associates the work request names with an application environment name; in this example, X and Y are associated with application environment AE1, and Z with AE2. The application environment names AE1 and AE2 are specified to workload management in the service definition stored in the WLM couple data set, and are included in the active policy when a policy is activated.

Each application environment must be assigned a system procedure library (PROCLIB) member name that contains the JCL required to start server address spaces for the application environment. In the example, PROCS1 and PROCS2 are associated with application environments AE1 and AE2, respectively.

When the work manager subsystem receives a type Y work request from a client, the subsystem requests that workload management associate the request with AE1. Workload management determines if a server address space is available to handle the request, or if an address space needs to be created. If a server address space exists, workload management makes the request available for the server. If a server address space does not exist or if more are required, workload management starts a server address space using the startup JCL procedure named PROCS1 which is defined for AE1 in the active policy.

Workload management dynamically creates new server address spaces if they are needed to handle more incoming work and, for certain subsystems such as DB2, decreases the number of server address spaces if less capacity is needed. Refer to Chapter 13, "Defining Application Environments" on page 101 for a description of how to define application environments to workload management.

## Why Use Scheduling Environments?

*Scheduling environments* help ensure that units of work are sent to systems that have the appropriate resources to handle them. A scheduling environment is a list of resource names along with their required states. Resources can represent actual physical entities, such as a data base or a peripheral device, or they can represent intangible qualities such as a certain period of time (like second shift or weekend).

These resources are listed in the scheduling environment according to whether they must be set to ON or set to OFF. A unit of work can be assigned to a specific system only when all of the required states are satisfied. This function is commonly referred to as *resource affinity scheduling*.

Figure 7 shows a simple scheduling environment example. The arriving work units, X, Y, and Z, could be batch jobs submitted through JES2 or JES3. Each of the jobs had a scheduling environment associated with it at the time it was submitted (in this case the X and Y jobs are each associated with the A scheduling environment, and the Z job is associated with the B scheduling environment).

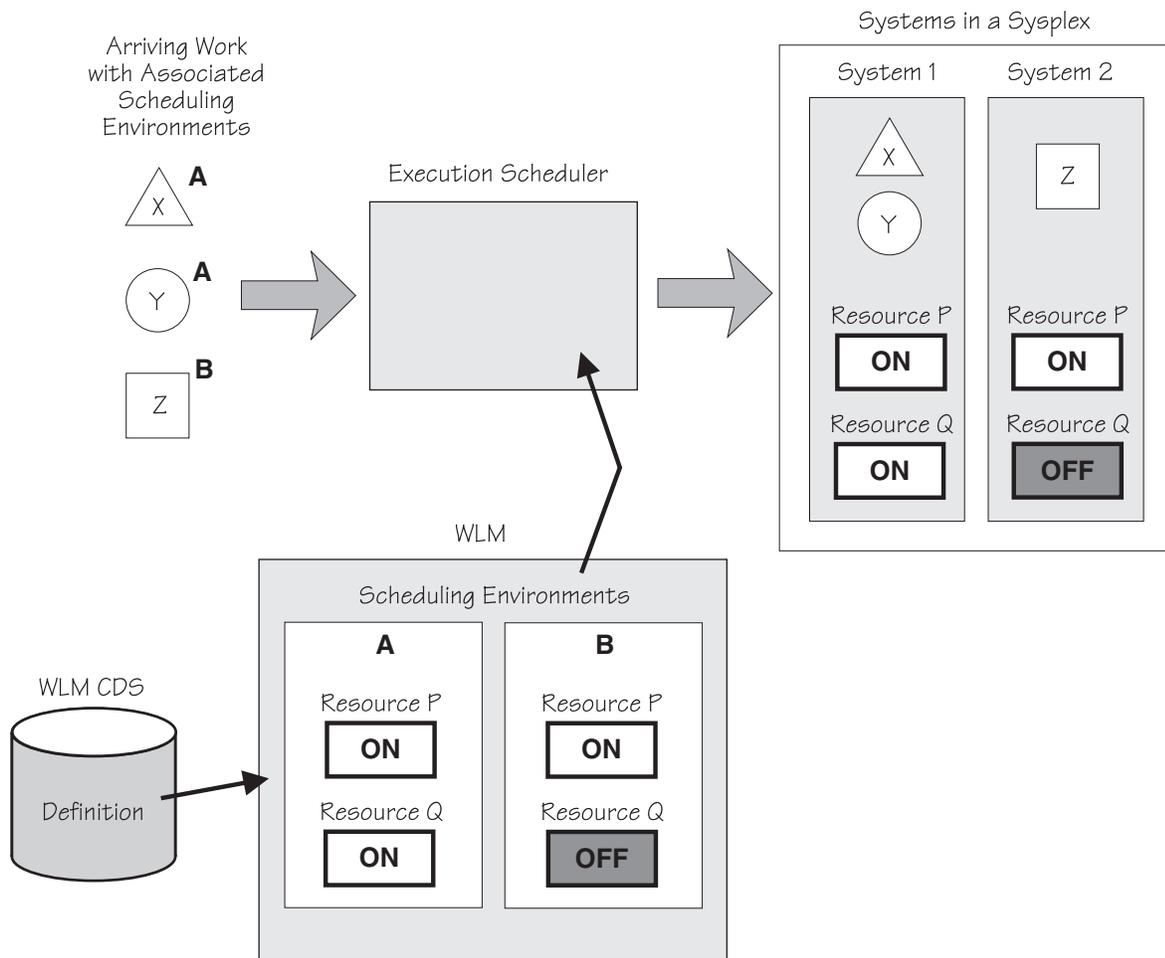


Figure 7. Scheduling Environment Example

JES checks the scheduling environment associated with each arriving batch job and then assigns the work to a system that matches that scheduling environment. In the example, both the X and Y jobs require that both Resource P and Resource Q be set to ON, so those jobs can be initiated only on System 1 in the sysplex. The Z

job requires that Resource P be set to ON and that Resource Q be set to OFF. So that job can be initiated only on System 2.

In a sysplex containing only one system, scheduling environments have some degree of usefulness, as JES will hold batch jobs until the required states become satisfied. In a multi-system sysplex, the full power of scheduling environments becomes apparent, as work is assigned only to those systems that have the correct resource states (the resource affinity) to handle that work.

Presently, JES2 and JES3 are the only participants that use scheduling environments, although the concepts could certainly apply to other types of work in the future.

Refer to Chapter 14, “Defining Scheduling Environments” on page 113 for a description of how to define scheduling environments to workload management.

## Summary of Service Definition and Service Policy Concepts

When you set up your service definition, you identify the workloads, the resource groups, the service classes, the service class periods, and goals based on your performance objectives. Then you define classification rules and one or more service policies. This information makes up the base service definition.

A two-step process is required before the sysplex starts using a new service definition. First, you install the service definition onto the WLM couple data set. Second, you activate one of the service policies from the definition.

With a service policy, you can override specific goals or resource groups in your base service definition. In a typical scenario, you might define a base service definition that is appropriate for your normal business hours. Because you need to have at least one service policy defined, you might create an empty service policy called NORMAL. While the NORMAL service policy is in effect, there would be no overrides to the goals or resource groups in the base service definition. If you have a special business need to change your goals for offshift processing, you might then also create a service policy called OFFSHIFT. If you were to activate this policy at the end of the business day (either by invoking the **VARY WLM,POLICY=*policyname*** command or by using the “Activating a Service Policy” panel in the ISPF application), then the goal overrides in the OFFSHIFT service policy would be in effect until you switched back to NORMAL the next morning.

Chapter 5, “Defining Service Policies” on page 31 tells you more about how to define a policy, and also shows a few examples.

Note that you can override *only* goals, number and duration of periods, resource group assignments and values. All of the workloads, service class names, classification rules, scheduling environments, and application environments defined in the service definition remain the same for any policy. If you need to change any of these, you will need to change the base service definition, re-install the service definition, and then activate a policy from that changed service definition.

Note, also, that you need to define all of your policies at the outset, while you are defining the rest of the service definition. Once the service definition is installed, then you can switch from one defined policy to another. If you need to create a

new policy or change the overrides in an existing policy, you will need to re-install the service definition with the new or changed policy definition before you can activate the new policy.

Figure 8 shows a service definition with two service policies.

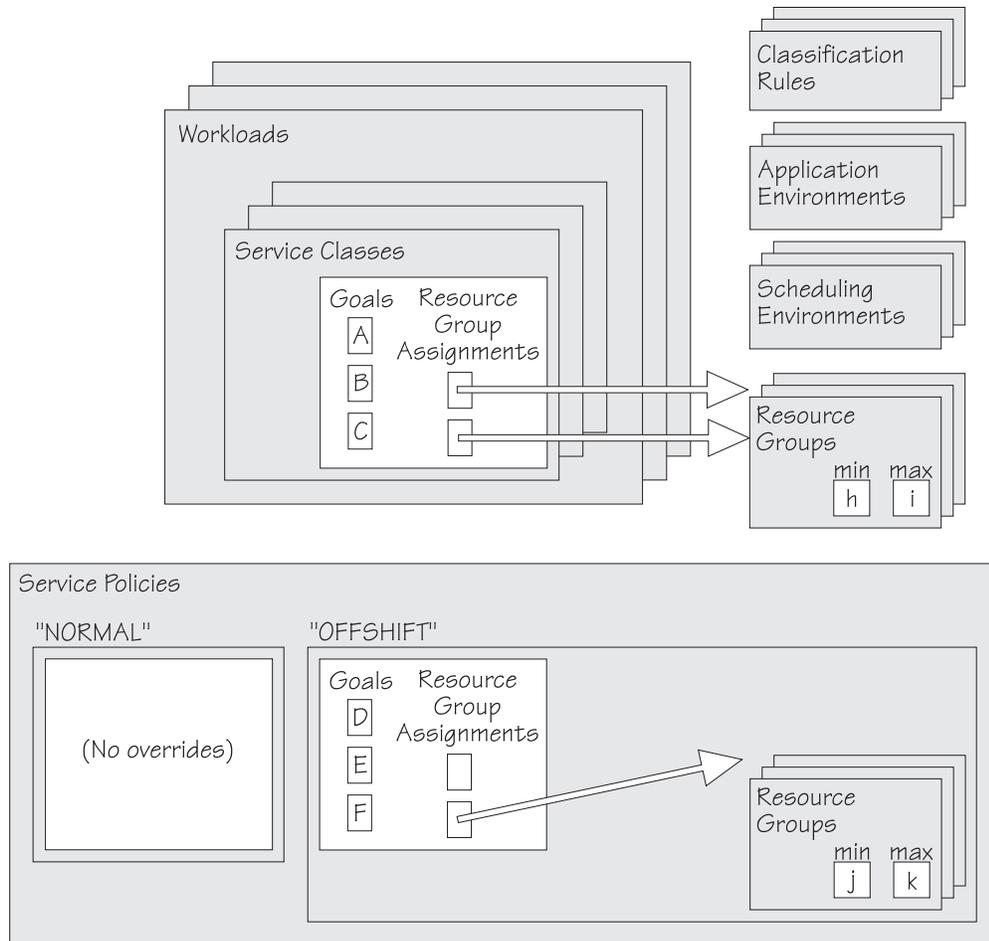


Figure 8. Service Definition, Including Two Service Policies. The NORMAL service policy is empty — therefore the originally defined goals, resource group assignments, and resource group attributes remain unchanged when the NORMAL service policy is in effect. When the OFFSHIFT service is in effect, certain goals, resource group assignments, and resource group attributes are overridden.



---

## Chapter 3. Workload Management Participants

You use the WLM ISPF administrative application to define your service definition. The administrative application requires the following products:

- TSO/E Version 2.5 plus SPEs or later
- ISPF 4.3 or later

This chapter describes the work and reporting environments that support workload management.

---

### Workload Management Work Environments

Cooperation between MVS and the subsystem work managers enable the sysplex-wide specification and reporting of goals defined for work in the service policy. You can define goals for the following kinds of work:

- IMS, if you have IMS/ESA Release 5 or later
- CICS, if you have CICS/ESA 4.1 or later
- OS/390 UNIX System Services
- JES2
- JES3
- APPC
- TSO/E
- SOMobjects
- Component Broker (CB) Objects
- LSFM
- DDF, if you have DB2 V4.1 or later
- DB2, if you have DB2 V 5 or later
- IWEB, if you have Internet Connection Server V 2.2 or later, Domino Go Webserver, or IBM http Server for OS/390
- MQSeries Workflow
- SYSH, if you need to manage non z/OS partitions

Arriving work is associated with a service class, and therefore a goal and possibly a resource group. You get feedback from RMF as to how well the system is doing in meeting the goals. RMF 1.3.0 for OS/390 R3 and later provides sysplex-wide workload management reporting, providing sysplex-wide as well as single system feedback on the goals through the Post Processor and Monitor III realtime reports. The reports show kinds of delays seen by subsystem work managers such as CICS and IMS. SDSF displays workload management related information.

## The Workload Management Environment

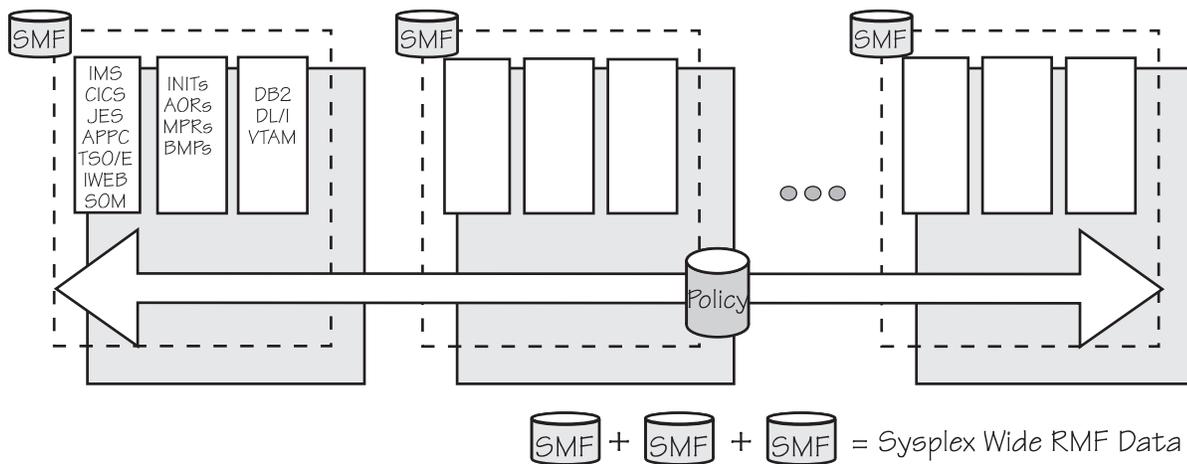


Figure 9. Sysplex View of the Management Environment. The shaded box represents the processor(s) and the dashed box represents the MVS image(s) running on the processor.

Workload management understands which address spaces (INITs, AORs, MPPs, BMPs, TORs, FORs) are involved in processing work within a service class, and matches the resources to meet the goal. Information as to how well each system is doing in processing towards a service class goal is recorded in SMF records on each system. For more information about workload management information in RMF reports, see *z/OS RMF Report Analysis*.

## Subsystem Support for Goal Types and Multiple Periods

The types of goals a subsystem supports depends on the workload management services it uses. For the following subsystems that have address space-oriented transactions or use enclaves, you can specify any goal type and multiple periods:

- **Subsystems that have address space-oriented transactions:**
  - APPC
  - JES2
  - JES3
  - OS/390 UNIX System Services
  - TSO/E
- **Subsystems that use enclaves:**
  - Component Broker
  - DB2
  - DDF
  - IWEB
  - MQSeries Workflow
  - LSFM
  - SOMobjects

**Note:** Enclaves are transactions that can span multiple dispatchable units in one or more address spaces, and in the case of multisystem enclaves, one or more address spaces on multiple systems in a parallel sysplex. See *z/OS MVS Programming: Workload Management Services* for more information on enclaves.

The CICS and IMS subsystems do not use enclaves, but use a different set of WLM services to support their transactions to WLM. Therefore, they support only response time goals, either percentile or average, and single period service classes.

## Subsystem-Specific Performance Hints

Based on installation experiences, here are some subsystem-specific performance hints:

- **Watch out for increased CPU usage by the WLM address space due to a high CICS MAXTASK setting.** For CICS 4.1 and later releases, WLM collects delay samples for each performance block. Because the number of performance blocks created is based on the MAXTASK value (a value of 100 means 100 performance blocks created per region), a MAXTASK value that is too high can cause a large sampling overhead when a CICS workload is switched to goal mode. If MAXTASK has been set to an arbitrarily high value, it should be reduced to a true “high water mark” value.
- **Watch out for work defaulting to SYSOTHER.** Work in subsystems that use enclaves (see “Subsystem Support for Goal Types and Multiple Periods” on page 18 for a list of these subsystems) can suffer performance degradation if left unclassified in the service definition. If you do not add classification rules for this work in your service definition, it will be assigned to the SYSOTHER service class, which has a discretionary goal. Using the WLM application, you need to add classification rules to assign the work to service classes with appropriate response time or velocity goals.

As a general rule, it’s a good idea to keep an eye on the SYSOTHER service class through RMF or another monitor. Any service accumulating in the SYSOTHER service class is a signal that you have unclassified work in your system.

For the latest information on these topics and others, see the WLM/SRM web site at:

<http://www.ibm.com/servers/eserver/zseries/zos/wlm/>

---

## Workload Balancing

Workload management allocates resources to meet goals of the work that arrives. System programmers must use the existing methods of routing and scheduling work for subsystems except for those listed below. For subsystems not exploiting workload balancing or routing services, if you want to balance your work across all MVS images in a sysplex, the system programmer must set the routing controls to either balance the arrival of work, or to ensure that all MVS images are equal candidates for processing work.

Examples of subsystems that can automatically balance work in a sysplex include:

- CPSM provides goal-oriented routing based on the goal defined for CICS work in the workload management service policy.
- DB2 V 4.1 provides automatic and dynamic work balancing in a sysplex for distributed data facility (DDF) work.
- DB2 V 5 provides additional automatic work balancing through its sysplex query parallelism function.
- SOMobjects in OS/390 R3 and later uses application environments to help balance object class binding requests in a sysplex.
- CICS V 4.1, DB2 V 4.1, TSO/E V 2.5, and APPC cooperate with VTAM 4.4 and OS/390 R3 workload management and later in a sysplex to manage session placement. New sessions for these subsystems are directed to the appropriate systems in the sysplex to balance work and meet performance goals.
- OS/390 R4 JES2 and OS/390 R8 JES3 provides automatic and dynamic placement of initiators for WLM-managed job classes.

- Component Broker cooperates with WLM in a sysplex to balance work among application control regions and to meet performance goals.

## Workload Management in a CPSM Environment

Figure 10 shows workload management in a CICS with CPSM environment.

WLM recognizes that the terminal owning region (TOR) and the application owning region (AOR) on one or more systems are involved in processing CICS transactions. Using RMF Monitor I, you can get reporting information on the CICS response times, and on any execution delays experienced by a service class period for a single system or for the sysplex.

Workload Management In a CICS Environment

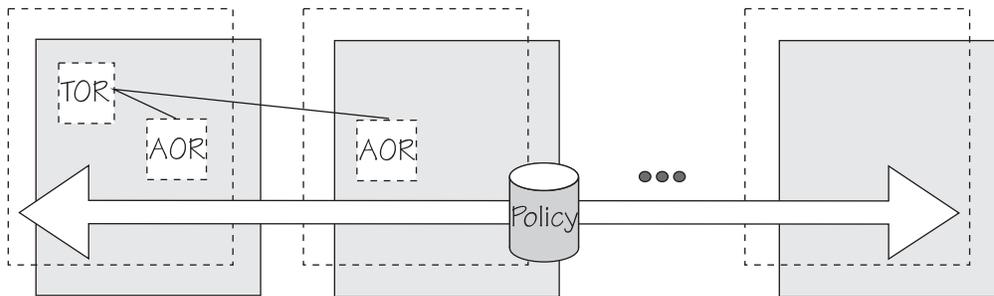


Figure 10. Workload Management in a CICS Environment

In a CPSM environment, WLM provides the CICS service class goal to CPSM. If the goal is an average response time goal and you specified the dynamic goal-oriented algorithm for CPSM, then CPSM uses the transaction's goal to help decide where to route the transaction for processing to meet the goal. If the goal is a percentile response time goal, CPSM reverts to its "shortest queue" algorithm since it only has average response time data available to it. Percentile goals are still preferred for any workload that can have a few unusually long transactions distorting the average response time. Following is a summary of the two CPSM algorithms:

### Shortest Queue

Send the CICS transaction to the AOR that has the shortest queue of pending work, but prefer AORs on a local MVS image, and be sure to avoid "unhealthy" AORs.

**Goal** Send the transaction to an AOR that has been successfully meeting the goal, but prefer AORs on a local image and be sure to avoid "unhealthy" AORs. If the goal is a percentile goal, use the shortest queue.

For more information, see *CICSplex SM Managing Workloads*.

## Workload Management in a DB2 Distributed Data Facility (DDF) Environment

The definition of a response time for the enclaves used to manage DDF transactions depends upon several parameters, including DB2 install parameters and the attributes used when binding the package or plan. For more information, see *DB2 for OS/390 V5 Administration Guide*.

---

## Batch Workload Management

Workload management can dynamically manage the number of batch initiator address spaces in a JES2 or JES3 environment. You can selectively turn over control of the batch initiator management to WLM for one or more job classes. WLM will start new initiators, as needed, to meet the performance goals of this work.

By specifying or defaulting **MODE=JES** on the JES2 **JOBCLASS** statement or the JES3 **GROUP** statement, you indicate that the initiators for the job class should be JES-managed, as in the past. By specifying **MODE=JES**, you keep the job class in JES-managed mode. (JES will manage the batch initiators for that job class, in the same way it has in prior releases.) By specifying **MODE=WLM**, you put that class into WLM-managed mode.

You can switch as many job classes to WLM-managed mode as you wish. You can easily switch any job class back to JES-managed mode by using the JES2 **\$TJOBCLASS** command or the JES3 **MODIFY** command.

### Notes:

1. If you have velocity performance goals set for the work running on WLM-managed batch initiators, be aware that the initiation delays will be figured into the velocity formula. This will affect your velocity values and probably require some adjustment to your goals. See Chapter 8, “Defining Service Classes and Performance Goals” on page 39 for information on defining velocity goals.
2. All jobs with the same service class should be managed by the same type of initiation. For example, if jobs in job classes A and B are classified to the HOTBATCH service class, and **JOBCLASS(A)** is **MODE=WLM**, while **JOBCLASS(B)** is **MODE=JES**, workload management will have a very difficult time managing the goals of the HOTBATCH service class without managing class B jobs.
3. You can use the JES2 **JOBCLASS** parameter to specify a default scheduling environment, thereby saving the effort of changing JCL jobcards or writing a specific JES2 exit to assign the scheduling environment.
4. If a WLM initiator is experiencing ABEND822s, there are two ways to recycle the initiator:
  - If you can determine the ASID of the initiator that is abending, you can stop it by issuing the command **P INIT,A=asid**. The initiator does not need to be idle at the time that you enter the command. If the initiator is busy processing a job, it will stop itself after the job finishes. WLM will automatically replace the initiator with a new one.
  - If you cannot determine the ASID, or if you want to recycle all initiators as part of a regular cleanup process, you can enter the commands **\$P XEQ** and **\$S XEQ**. The **\$P XEQ** command causes all WLM initiators on that system to be “flagged” to terminate. The **\$S XEQ** command enables WLM to start new initiators (without needing to wait for the old initiators to end). Beware that the **\$P XEQ** command purges WLM’s history which tells it how many initiators are needed for each service class. It may take some time for WLM to build up the same number of initiators that existed before.

See the following JES2 and JES3 manuals for more information on WLM-managed JES2 job classes:

- *z/OS JES2 Initialization and Tuning Guide*
- *z/OS JES2 Initialization and Tuning Reference*
- *z/OS JES2 Commands*

- *z/OS JES3 Initialization and Tuning Guide*
- *z/OS JES3 Initialization and Tuning Reference*
- *z/OS JES3 Commands*

The following other functions exist to help you manage batch work in a JES environment:

- A new work qualifier, **PRI**, which allows you to use the job priority when defining work classification rules. See Chapter 9, “Defining Classification Rules” on page 49.
- Scheduling environments, which allow you to define resource requirements for incoming work, ensuring that the work will be scheduled on a system within the sysplex only if the resource settings on that system satisfy those requirements. See Chapter 14, “Defining Scheduling Environments” on page 113.

---

## Multisystem Enclave Support

With multisystem enclave support, enclaves can run in multiple address spaces spanning multiple systems within a parallel sysplex. As in a single system enclave, the work will be reported on and managed as a single unit.

UNIX System Services Parallel Environment uses multisystem enclaves to run parallel jobs. With all tasks of the job running in the same enclave, WLM can manage all of the work to a single performance goal.

See Chapter 17, “Defining a Coupling Facility Structure for Multisystem Enclave Support” on page 155 for more information on setting up the SYSZWLM\_WORKUNIT coupling facility structure, a prerequisite to multisystem enclave support.

See the “Creating and Using Enclaves” chapter in *z/OS MVS Programming: Workload Management Services* for more information on multisystem enclaves.

See *z/OS UNIX System Services Parallel Environment: Operation and Use* for more information on UNIX System Services Parallel Environment.

---

## Intelligent Resource Director

The Intelligent Resource Director (IRD) extends the concept of goal-oriented resource management by allowing you to group logical partitions that are resident on the same physical server, and in the same sysplex, into an “LPAR cluster.” This gives Workload Management the ability to manage resources, both processor and DASD I/O, not just in one single image but across the entire cluster of logical partitions.

See Chapter 18, “The Intelligent Resource Director” on page 159 for more information.

---

## Non z/OS Partition CPU Management

Note that this function is shipped disabled with z/OS V1R2 . A PTF for APAR OW50221 is needed to enable this function. This PTF will be provided by 12/01. If you have logical partitions running LINUX or other non z/OS systems, you can manage CPU resources across these logical partitions in accordance with workload goals. Non z/OS partition CPU management does not support the management of partitions running OS/390 systems.

Non z/OS partition CPU management allows WLM to shift weight from either z/OS or non z/OS partitions to another z/OS or non z/OS partition in the same LPAR cluster.

A new subsystem type SYSH is introduced for LINUX CPU management. SY and PX are the qualifiers which are valid for SYSH. No resource groups can be associated with service classes for SYSH. Also CPU protection cannot be assigned. SYSH supports velocity goals with single periods, but no discretionary goals. In order to activate non z/OS CPU Management, see “Enabling Non z/OS CPU Management” on page 167.

---

## Workload Management, IBM License Manager, and Workload License Charges

As part of the z/OS support of IBM License Manager and Workload License Charges, you can set a defined capacity limit, also called a “soft cap,” for the work running in a logical partition. This defined capacity is measured in millions of service units per hour (MSUs). It allows for short-term spikes in the CPU usage, while managing to an overall, long-term rolling 4 hour average. It applies to all work running in the partition, regardless of the number of individual workloads the partition may contain.

See the *PR/SM Planning Guide* for more information on actually setting a defined capacity limit.

WLM enforces the defined capacity limit by tracking the partition’s CPU usage and continually averaging it over the past 4 hours. Spikes above the defined capacity limit are possible, as long as they are offset by low points that keep the 4 hour average at or below the limit. When this 4 hour average goes over the defined capacity limit, then WLM caps the partition. At that point, it can use no more than the defined capacity limit, until the average drops below the limit.

Consider the following example:

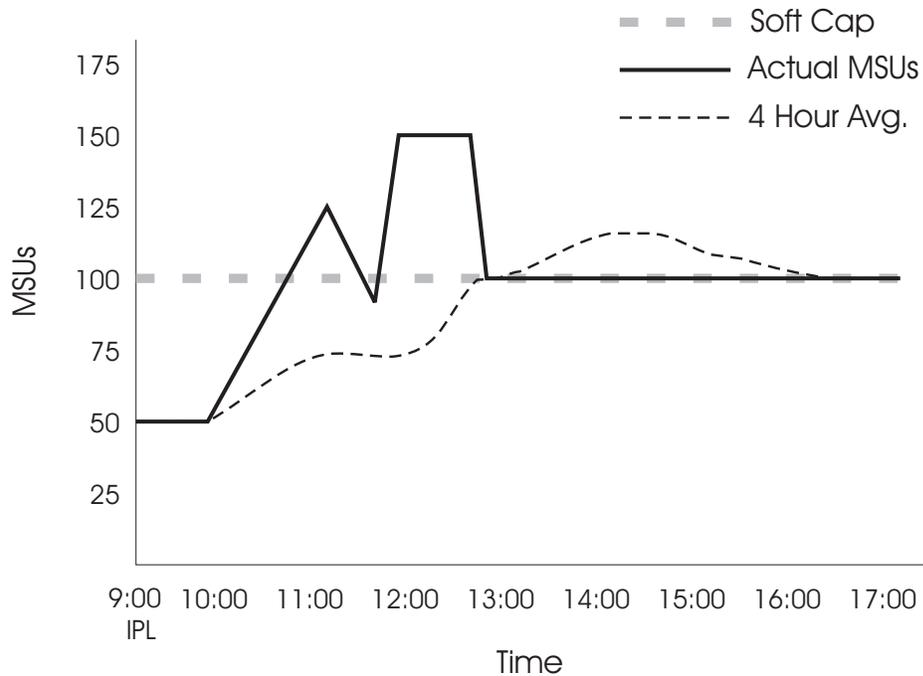


Figure 11. Workload Management Enforcing a Defined Capacity Limit

This partition has a 100 MSU defined capacity limit, as shown by the gray dotted line. The solid black line is the actual MSU consumption of the partition, and the dotted black line is the 4 hour MSU average. The partition's CPU usage is 50 MSUs at 9:00, when the system is IPLed. By 12:00, the partition's CPU usage has gone up to 150 MSUs. The 4 hour average, however, is still below the defined capacity limit at that point, so no action is taken. (Note that the 4 hour average is based on the 3 hours the system has been up since 9:00 — when the system has been up for 4 hours, the average will start rolling.) Just before 13:00, the 4 hour average hits the defined capacity limit. Now WLM caps the partition's CPU usage at 100 MSUs.

WLM allows the CPU usage to remain at the defined capacity limit. Therefore, the 4 hour average can continue to go up. In this example, this occurs because the low usage numbers (50 MSUs between 9:00 and 10:00) are falling off the back end of the 4 hour "horizon," and being replaced by the new usage numbers (150 MSUs between 13:00 and 14:00). This is a natural consequence of managing to the 4 hour average.

In summary, WLM caps a partition's CPU usage only when the 4 hour average reaches the defined capacity limit. Before that point, you may see CPU usage spikes over the limit. You may see the 4 hour average go over the limit after capping, until the passage of time gradually brings the average back down to the limit.

## Workload Management with Other Products

For subsystem work managers, if you do not have the product or product release supporting workload management, you can define goals only for the subsystem regions. Since subsystem regions are treated as batch jobs or started tasks, a velocity goal is most appropriate. All started tasks are managed on the system where they are started.

For example, if you have a CICS release that is not at least at the CICS/ESA 4.1 level, you could define velocity goals for the TORS and AORS. You could not manage CICS transactions to response time goals.

For reporting or monitoring products, you should check whether they support workload management. Vendor documentation should explain whether or not they report on the workload management activity.



---

## Chapter 4. Setting up a Service Definition

This chapter describes how to set up a service definition based on your performance objectives. A service definition contains all of the information necessary for workload management processing. Based on the information in this chapter, you should be able to set up a preliminary service definition using the worksheets provided. Then, you can enter your service definition into the ISPF administrative application with ease.

The service definition is the way to express your installation's business goals to your sysplex. In order to do this, you must understand your installation's business environment from the following areas:

- What are your installation's revenue-earning workloads?
- Is there a business priority to the workloads?
- Do you understand what kind of service users expect and the time-frame they expect it in?
- Do you understand the service you can deliver?  
Even if your installation does not currently have an SLA, or other written performance objectives, users most often have service expectations.
- How can you use monitoring and performance products to gather information?

---

### Specifying a Service Definition

You define a minimum of one policy in a service definition. A service definition contains workloads, service classes, and classification rules. Optionally, it contains resource groups, report classes, and application environments. A service definition also includes one or more policies. A policy is a set of overrides to the goals and resource group limits in the service definition. In a service definition, you also specify whether you want workload management to manage I/O priorities based on your service class goals. A service definition has an identifying name and description.

There is one notepad available for the service definition. You can update the notepad whenever you are creating or modifying any of the other parts of a service definition, such as the policy, workload, service class, resource group, and classification rules. All of the information in a service definition is available to reporting products, except for the notepad information and the classification rules. The notepad information is provided for a history log, and change management.

**Name** Service definition name

**Description**  
Description of the service definition

**Name (required)**  
Eight character identifier of the service definition.

**Description (optional)**  
An area of 32 characters to describe the service definition. For example, you could include the time period this service definition is intended to cover.

---

## Storing Service Definitions

You can work in the ISPF administrative application with one service definition at a time. In order to make the service definition accessible to all systems in the sysplex, you store the service definition on a WLM couple data set. Only one service definition can be installed on the WLM couple data set at a time.

If you want to work on more than one service definition at a time, you can keep each in a distinct MVS partitioned data set (PDS). As an MVS PDS, the service definition is subject to all the same functions as a PDS. You can restrict access to the PDS, send it, and copy it, as you can any MVS PDS.

The service definition must be installed on a WLM couple data set, and a service policy activated. Only service policies in the service definition installed on the WLM couple data set can be activated. A WLM couple data set can have automatic back-up. For more information about allocating and using a WLM couple data set, see Chapter 16, "Workload Management Migration" on page 131.

---

## Defining the Parts of a Service Definition

The following chapters explain how to set up a service definition by defining each of its parts: policies, workloads, resource groups, application environments, service classes, classification rules, and report classes. When you set up your service definition, you should define its parts in the following order:

**1. Service policy**

Chapter 5, “Defining Service Policies” on page 31.

**2. Workloads**

Chapter 6, “Defining Workloads” on page 35.

**3. Resource groups**

Chapter 7, “Defining Resource Groups” on page 37.

**4. Service Classes**

Chapter 8, “Defining Service Classes and Performance Goals” on page 39.

**5. Service policy overrides**

“Using Policy Overrides” on page 32.

**6. Classification rules**

Chapter 9, “Defining Classification Rules” on page 49.

**7. Report classes**

Chapter 10, “Defining Report Classes” on page 81.

**8. Service coefficients and options**

Chapter 11, “Defining Service Coefficients and Options” on page 83. Service definition options include I/O priority management.

**9. Application environments**

Chapter 13, “Defining Application Environments” on page 101.

**10. Scheduling environments**

Chapter 14, “Defining Scheduling Environments” on page 113.



---

## Chapter 5. Defining Service Policies

A *service policy* is a named collection of service class and resource group specification overrides. When a policy is put into effect, the overrides are merged with the service class and resource group specifications in the service definition. A policy override is a way to change a goal or resource group capacity without having to redefine all of your service classes and resource groups.

See “Summary of Service Definition and Service Policy Concepts” on page 14 for an overview of the relationship between a service definition and a service policy, and Figure 8 on page 15 for a visual overview of how service policy overrides work. Note that in an ideal scenario, you would only have to define your service definition once. As part of that service definition, you would predefine multiple policies to meet varying performance goals or business needs. Once the service definition is installed, you would then activate one policy at a time, and then, when appropriate, switch to another. Note that you *must* define at least one service policy, and you can define up to 99.

When you are creating your service definition, you may choose to define one empty “default” policy with no overrides at all. Next, create your workloads and service classes. Then determine how and when your service classes may have different goals at different times. Define additional policies with the appropriate overrides for these time periods.

### Defining Service Policies

**Name** Service policy name

**Description**

Description of the service policy

**Policy Override**

Changing a service class goal or resource group

#### **Name (required)**

Eight characters identifying the service policy. Every service policy name must be unique in a service definition. The service policy is activated by name in one of the following ways:

- An operator command from the operator console.
- A utility function from the workload management ISPF application.

You can display the name of the active service policy with an operator command, or by viewing a performance monitor, such as RMF.

#### **Description (optional)**

An area of 32 characters describing the service policy. The descriptive text is available to performance monitors for reporting.

#### **Policy Override (optional)**

A way to change a performance goal, a service class-resource group assignment, or a resource group capacity for a service policy. For more information about defining policy overrides, see “Using Policy Overrides” on page 32.

### Examples of Service Policies

- Daytime policy
  - Name = DAYTIME
  - Description = Policy from 9:00 am to 5:00 pm
- Policy for national holidays
  - Name = HOLIDAY
  - Description = Policy for Arbor day
- Weekend policy
  - Name = WEEKEND
  - Description = Policy for Sat and Sun

## Using Policy Overrides

Once you have defined your service classes, you can determine whether any of your service class goals change at different times. If they do, you can define a policy override. With an override, you can change one or more of the following for a service policy:

- A goal for a service class period
- Number and duration of periods
- A service class - resource group assignment
- Resource group attributes
- CPU protection.

### Example 1: Policy Overrides

In this example, the service class BATPIG is in the BATCH workload. It is associated with the resource group LIMIT. Suppose the LIMIT resource group is assigned some maximum capacity. BATPIG is assigned a discretionary goal. Since it is a discretionary goal, it does not have an assigned importance. In the Weekend policy, however, both the goal and the resource group association is overridden. The resource group association is overridden, so that in the Weekend policy, BATPIG is not assigned to a resource group. It is instead assigned a response time goal of 1 hour, with an importance of 5.

```

----- Base -----
Policy: Standard                                Policy: Weekend
                                                (overrides)

Service Class..... BATPIG
Decription..... All batch CPU hogs
Workload..... BATCH
Resource Group. LIMIT

Period   Goal.... discretionary           1 hour AVG
 1       Import.. n/a                     5
  
```

### Example 2: Policy Overrides

In this example, the CICS HIGH service class is in the CICS workload. In the Standard policy, it is assigned a response time goal of 1 second average, with an importance of 1. It is not assigned a resource group, because there is no business need to limit or guarantee capacity. For the holiday policy, however, it has an overridden goal. Because of a contract with CICS users, the agreed response time for weekends is 2 seconds average, with an importance of 2.

```

----- Base -----
Policy: Standard                                Policy: Weekend
                                                (overrides)

Service Class..... CICSHIGH
Description..... All short CICS transactions
Workload..... CICS
Resource Group. _____

Period   Goal.... 1 sec avg                    2 sec avg
1        Import.. 1                               2
  
```

### Example 3: Policy Overrides

In this example, the resource group DEPT58 is associated with 3 service classes: 58CICS, 58TSO, and 58BATCH. Since the department is willing to pay for more capacity on the weekends, the minimum is overridden for the weekend policy. It is 1500 CPU service units per second. So in the weekend policy, the service classes 58CICS, 58TSO, and 58 BATCH have a minimum of 1500 CPU service units guaranteed.

```

----- Base -----
Policy: Standard                                Policy: Weekend
                                                (overrides)

Resource Group..... DEPT58
Description..... Contracted capacity for dept. 58

Minimum   1000                                1500
Maximum   _____

Service Class 58CICS _
              58TSO  _
              58BATCH_
  
```



---

## Chapter 6. Defining Workloads

A *workload* is a named collection of work to be reported as a unit. You can arrange workloads by subsystem (CICS, IMS), by major application (Production, Batch, Office) or by line of business (ATM, Inventory, department). Logically, a workload is a collection of service classes.

### Defining Workloads

**Name** Workload name

**Description**  
Description of workload

#### **Name (required)**

Eight characters identifying the workload. Every workload name must be unique for all defined workloads in a service definition.

#### **Description (optional)**

An area of 32 characters describing the workload. The descriptive text is available to performance monitors for reporting.

---

## Defining a Departmental Workload

In order to set up a departmental workload that crosses subsystem boundaries, you must keep in mind how you can assign the work to service classes. You should review Chapter 9, "Defining Classification Rules" on page 49 and find out the best way you can assign your work to service classes. The following table shows examples of what would be required for some of the subsystems:

- TSO** You must set up TSO userids or account numbers according to department structure, so that the userids correspond to a specific department.
- JES** You must have unique batch classes or account numbers by department.
- CICS** You must have unique CICS regions for each department.
- IMS** You must have a separate IMS/VS resource lock manager (IRLM), IMS control region, and IMS message processing region (MPR) for each workload.

For more information, see Chapter 9, "Defining Classification Rules" on page 49.

### Examples of Workloads

- By subsystem

**Name** IMS

**Description**

All work in classes IMS1 IMSA IMSS TIMS IMSV

- By department/location

**Name** DEVELOP

**Description**

All work in classes STSO IMSB and LINKA

---

## Chapter 7. Defining Resource Groups

A *resource group* is an amount of processor capacity across the sysplex. It is optional. Unless you have some special need to limit or protect processor capacity for a group of work, you should skip defining resource groups and let workload management manage all of the processor resource to meet performance goals. You use a resource group to:

- Limit the amount of processing capacity available to one or more service classes.
- Set a minimum processing capacity for one or more service classes in the event that the work is not achieving its goals.

You can specify a minimum and maximum amount of capacity in unweighted CPU service units to a resource group. The minimum and maximum capacity applies to all systems in the sysplex. You can assign only one resource group to a service class. You can assign multiple service classes to the same resource group. You can define up to 32 resource groups per service definition.

Keep in mind your service class goals when you assign a service class to a resource group. Given the combination of the goals, the importance level, and the resource capacity, some goals may not be achievable when capacity is restricted.

If work in a resource group is consuming resources above the specified maximum capacity, the system throttles the associated work to slow down the rate of resource consumption. The system may use several mechanisms to slow down the rate of resource consumption, including swapping the address spaces, changing its dispatching priority, and capping the amount of service that can be consumed. Reporting information reflects that the service class may not be achieving its goals because of the resource group capping.

By setting a minimum processing capacity, you create an overriding mechanism to circumvent the normal rules of importance. If the work in a resource group is not meeting its goals, then workload management will attempt to provide the defined minimum amount of CPU resource to that resource group.

### Defining Resource Groups

**Name** Resource Group name

**Description**

Description of resource group

**Capacity Maximum**

Maximum unweighted CPU service units per second

**Capacity Minimum**

Minimum unweighted CPU service units per second

**Note:** For more information on unweighted CPU service units, see the “System Resources Manager” chapter of *z/OS MVS Initialization and Tuning Guide*.

**Name**

Eight characters that identify the name of the resource group. Each resource group must be unique within a service definition.

**Description**

Up to 32 characters that describe the resource group.

**Capacity**

Identifies the amount of available capacity you want workload management to allocate to the resource group. Capacity is in unweighted CPU service units per second, and it includes cycles in both TCB and SRB mode. The table in Appendix B, "CPU Capacity Table" on page 225 shows the service units per second by CPU model. Resource Group minimum can equal resource group maximum. For a given resource group, you can vary the capacity minimum and maximum by service policy.

**Maximum**

CPU service units per second this resource group may use. *Maximum* specified for this resource group applies to all service classes in that resource group combined. *Maximum* is enforced. There is no default maximum value.

**Minimum**

CPU service units per second that should be available for this resource group when work in the group is missing its goals. The default is 0. If a resource group is not meeting its minimum capacity and work in that resource group is missing its goal, workload management will attempt to give CPU resource to that work, even if the action causes more important work (outside the resource group) to miss its goal. If there is discretionary work in a resource group that is not meeting its minimum capacity, WLM will attempt to give the discretionary work more CPU resource if that action does not cause other work to miss its goal.

The minimum capacity setting has no effect when work in a resource group is meeting its goals.

**Note:** You cannot assign a resource group to service classes representing transaction-oriented work, such as CICS or IMS transactions. The ISPF application notifies you with an error message if you attempt to define one. If you want to assign a minimum or a maximum capacity to CICS or IMS work, you can do so by assigning a resource group to their regions. For example, suppose you have three service classes representing your CICS works: CICSTRN, CICSAORS, and CICSTORS. CICSTRN represents all of your online CICS transactions, and has one period with a short response time goal. CICSAORS and CICSTORS represent all of your CICS AORs and TORs, respectively, that process the online transactions. To assign a maximum capacity to your CICSTRN work, define a resource group, and assign it to the regions. So you assign the resource group to the CICSAORS and CICSTORS service classes.

---

## Chapter 8. Defining Service Classes and Performance Goals

A *service class* is a named group of work within a workload with similar performance characteristics:

- Performance goals
- Resource requirements
- Business importance to the installation.

Workload management manages a service class period as a single entity when allocating resources to meet performance goals. A service class can be associated with only one workload. You can define up to 100 service classes.

You can assign the following kinds of *performance goals* to service classes: average response time, response time with percentile, velocity, and discretionary. You assign an importance level to the performance goal. Importance indicates how vital it is to the installation that the performance goal be met relative to other goals.

Because some work has variable resource requirements, workload management provides *performance periods* where you specify a series of varying goals and importances. You can define up to eight performance periods for each service class. You can also assign a service class to a resource group if its CPU service must be either protected or limited.

This chapter explains the parts of a service class, how to define performance goals, and how to use performance periods.

## Defining Service Classes and Performance Goals

**Name** Service class name

**Description**  
Service class description

**Workload**  
The name of the workload associated with this service class.

**Resource Group**  
The name of the resource group associated with the work in this service class.

**Performance Period**  
One goal per period.

**Duration**  
Number of service units for this performance period.

**Average Response Time**  
Average response time for transactions completing within the period in terms of hours, minutes, and seconds. Decimal points are accepted. Response time varies from 15 milliseconds to 24 hours.

**Response Time and Percentile**  
A percentile of work to be completed in the specified amount of time. Percentile boundaries vary from 1 to 99. Amount of time is in hours, minutes, or seconds. Decimal points are accepted. Response time ranges from 15 milliseconds to 24 hours.

**Velocity**  
Measure of how fast work should run when ready, without being delayed for WLM-managed resources. Velocity ranges from 1 to 99.

**Discretionary**  
Workload management defined goal. Work is run as system resources are available.

**Importance**  
How important it is to the installation that the goal be achieved.

**CPU Protection**  
Whether long-term CPU protection should be assigned to this service class.

### **Name (required)**

Eight characters describing the service class. Service class names must be unique within a service definition.

### **Description (optional)**

An area of 32 characters describing the service class. The descriptive text is available to performance monitors for reporting.

### **Workload (required)**

The name of the workload associated with the service class. You can associate only one workload per service class in a service definition. The workload must have been previously defined.

### **Resource Group (optional)**

The resource group name associated with this service class. You can assign

only one resource group per service class in a service policy. You can override the resource group assigned to a service class in each service policy. For more information about resource groups, see Chapter 7, “Defining Resource Groups” on page 37.

### **Performance Period**

A performance goal, importance, and *duration* for a service class. You set up multiple performance periods for work that has changing performance requirements as work consumes more and more resources. You can specify up to eight performance periods.

### **Duration**

Specifies the length of the period in service units. For a definition of service units, see Chapter 11, “Defining Service Coefficients and Options” on page 83. If the work included in this service class period does not complete when the number of service units have been used, the work moves into the next performance period. You do not specify a duration on the last defined period.

### **Response Time**

The expected amount of time required to complete the work submitted under the service class, in milliseconds, seconds, minutes and hours. Specify either an average response time, or response time with a percentile. Percentile is the percentage of work in that period that should complete within the response time. Percentile must be a whole number. You must specify a system response time goal, not “end-to-end”. That is, workload management does not control all aspects of system performance, so response time scope is confined to the time workload management has control of the work. This time includes the time the work is using or waiting for CPU, storage, or I/O service.

**Note:** Workload management does not delay work, or limit it, to achieve the response time goal when extra processing capacity exists.

### **Velocity**

A measure of how fast work should run when ready, without being delayed for WLM-managed resources. Velocity is a percentage from 1 to 99. See “Velocity Formula” on page 42 for a description of the calculations needed to determine velocity.

### **Discretionary**

Workload management defined goal. Associate this goal with work for which you do not have a specific performance goal. Work with a discretionary goal is run when excess resources are available.

### **Importance**

The relative importance of the service class goal. Importance is a reflection of how important it is that the service class goal be achieved, Workload management uses importance only when work is not meeting its goal. Importance indicates the order in which work should receive resources when work is not achieving its goal. Importance is required for all goal types except discretionary. Importance applies on a performance period level and you can change importance from period to period. Importance is in five levels: 1 to 5, 1 being the highest importance.

### **CPU Protection**

By specifying YES in the “CPU Critical” field when defining a service class, you ensure that work of lower importance will always have a lower dispatch priority. See Chapter 12, “Defining Special Protection Options for Critical Work” on page 89 for more information.

---

## Velocity Formula

The formula for velocity is:

$$\frac{\text{using samples}}{\text{using samples} + \text{delay samples}} \times 100$$

where:

### using samples

include:

- The number of samples of work using the processor
- The number of samples of work using non-paging DASD I/O resources (in a state of device connect or device disconnect).<sup>2</sup>

The I/O samples are derived from actual time measurements.

### delay samples

include:

- The number of samples of work delayed for the processor
- The number of samples of work delayed for storage
- The number of samples of work delayed for non-paging DASD I/O.<sup>2</sup>

Work delayed for storage includes:

- Paging
- Swapping
- Multiprogramming level (MPL)
- Server address space creation delays
- Initiation delays for batch jobs in WLM-managed job classes.

MPL is the SRM-controlled function that adjusts the number of address spaces allowed to be in central storage and ready to be dispatched.

I/O delays include:

- IOS queue
- Subchannel pending
- Control unit queue delays

The samples for subchannel pending and control unit queue delay are derived from actual time measurements.

---

## Defining Performance Goals

This section explains how to define performance goals in your service definition. If you have an SLA today, you should consider a few things:

- Does it contain end-to-end response time?

If it does, then you need to keep in mind that workload management processes towards system response times, and make the adjustment when you define the

---

<sup>2</sup> If I/O priority management is off, these samples are *not* included.

performance goal. Section “Determining System Response Time Goals” explains how you can determine the system response times of work.

- For what type of workloads do you need a different goal?

You may have some throughput type goals, which you need to convert into either response time goals, or velocity goals.

## Determining System Response Time Goals

Goal mode introduces several changes in the definition of a work request’s response time. The changes more accurately reflect end-user expectations of response time.

The number of batch transactions equals number of jobs. Defining a response time goal may not be appropriate for some types of batch work, such as jobs with very long execution times. Work that is appropriate for a response time goal should have at least three transaction completions per 20 minutes of elapsed time. If there are too few completions, use a velocity or discretionary goal.

TYPRUN=HOLD and TYPRUN=JCLHOLD times are *not* included in batch response times.

## Examples

### Examples of Service Classes with Response Time Goals

- A service class representing TSO/E work with multiple periods.

<b>Service Class</b>	TSO	
<b>Period</b>	1	
	<b>Response Time</b>	85% 0.5 second
	<b>Importance</b>	1
	<b>Duration</b>	400 Service Units
<b>Period</b>	2	
	<b>Response Time</b>	80% 1 second
	<b>Importance</b>	3
	<b>Duration</b>	1000 Service Units
<b>Period</b>	3	
	<b>Response Time</b>	60% 15 second
	<b>Importance</b>	4

Note that the percentile in period 1 and 2 refer to the transactions ending in each period, not the total TSO/E transactions.

- A service class representing CICS transactions.

<b>Service Class</b>	CICSHOT	
<b>Period</b>	1	
	<b>Response Time</b>	0.5 second AVG
	<b>Importance</b>	1

- A service class representing IMS transactions.

<b>Service Class</b>	IMSCAT1	
<b>Response Time</b>	95% .3 Second	
<b>Importance</b>	1	

- A service class representing IMS transactions.

<b>Service Class</b>	OIMSCAT3	
<b>Response Time</b>	5 sec AVG	
<b>Importance</b>	3	

## Defining Velocity Goals

This section describes where to find information to set a velocity goal, and what kind of work is appropriate for velocity goals. Velocity goals define the acceptable amount of delay for work when work is ready to run. Velocity goals are intended for subsystems which use address spaces or enclaves to represent individual work requests. Velocity goals are not supported for work in the IMS and CICS subsystem work environments because velocity data is accounted to the region, not to the individual transaction. Velocity is a goal to consider for long-running jobs.

For a service class with multiple periods, you cannot switch from a velocity goal in one period to a response time goal in a later period. See “Subsystem Support for Goal Types and Multiple Periods” on page 18 for a list of subsystems for which you can specify multiple periods.

## Adjusting Velocity Goals Based on Samples Included in Velocity Calculation

You can adjust your velocity goals based on whether or not the following samples are to be included in the velocity calculation:

- I/O samples (included in the velocity calculation if I/O priority management is turned on)
- Initiation delay samples (included in the velocity calculation if you have WLM-managed batch initiators).

In the RMF Monitor I workload activity report, there are two fields, **I/O PRTY** and **INIT MGMT**, which indicate the following:

### **I/O PRTY**

If you have I/O priority management turned off, then the **I/O PRTY** value shows you what your velocity would be if you were to turn I/O priority management on.

### **INIT MGMT**

If you are not currently using WLM-managed batch initiators, then the **INIT MGMT** value shows you what your velocity would be if you turned over control of all batch initiators in this service class to WLM.

**Note:** For each of the above two fields, it is assumed that the *other* setting is unchanged. For example, the **INIT MGMT** field assumes that your current I/O priority management setting remains the same.

These fields may help you to adjust a current velocity goal in anticipation of including these samples.

## Using Velocity Goals for Started Tasks

Velocity goals are the most appropriate goal for started tasks and long running work. Instead of figuring out a specific velocity goal for your started tasks, you should divide your started tasks into a high, a medium, and a low importance service class, and define a velocity that suffices for each category.

You can also take advantage of the system supplied service classes for started tasks: **SYSTEM** and **SYSSTC**. Workload manager recognizes special system address spaces (like **GRS**, **SMF**, **CATALOG**, **MASTER**, **RASP**, **XCFCAS**, **SMXC**, **CONSOLE**, **IOSAS**, **WLM**), puts them into the **SYSTEM** service class, and treats them accordingly. Address spaces in the **SYSSTC** service class are kept at a very high dispatching priority.

**Note:** You can also assign address spaces to the **SYSTEM** and **SYSSTC** service classes as part of your work classification rules. See “System-Provided Service Classes” on page 71.

For information about how to define service classes and associated classification rules for started tasks, see “Using the System Supplied Service Classes” on page 76.

Velocity is also appropriate for the “server” started tasks, that is, the address spaces that do work on behalf of a transaction manager or resource manager, such as CICS AOR, or an IMS control region. Since the server address spaces are processing work that also has an assigned performance goal, the velocity goal that you assign to servers applies only during address space start-up. Then workload management manages resources to meet the goals defined for the work the servers are processing, and not towards the goals defined for the servers.

If you have a version of a work manager such as CICS and IMS that does not support workload management, you cannot define a goal to the work manager’s transactions, but you can define a velocity goal for its server address spaces.

## Using Discretionary Goals

With discretionary goals, workload management decides how best to run this work. Since workload management’s prime responsibility is matching resources to work, discretionary goals are used best for the work for which you do not have a specific performance goal. For a service class with multiple performance periods, you can specify discretionary only as the goal in the last performance period.

Discretionary work is run using any system resources not required to meet the goals of other work. If certain types of other work are overachieving their goals, that work may be “capped” so that the resources may be diverted to run discretionary work. See “Migration Considerations for Discretionary Goal Management” on page 152 for more information on the types of work that are eligible for resource donation, and how you may want to adjust those goals.

### Examples of Service Classes with Discretionary Goals

- Discretionary goal as last period goal

<b>Service Class</b>	DEVBATCH	
<b>Period</b>	1	
	<b>Response Time</b>	80% 1 minute
	<b>Importance</b>	2
	<b>Duration</b>	2000 Service Units
<b>Period</b>	2	
	<b>Response Time</b>	80% 5 minutes
	<b>Importance</b>	3
	<b>Duration</b>	10000 Service Units
<b>Period</b>	3	
	<b>Discretionary</b>	

- Discretionary goal for leftover work

<b>Service Class</b>	ASDBATCH
<b>Discretionary</b>	

---

## Using Performance Periods

Performance periods are available for work that has variable resource requirements and for which your goals change as the work uses more resources. You specify a goal, an importance, and a duration for a performance period. Duration is the amount of service that period should consume before going on the next goal. Duration is specified in service units. For more information about defining durations, see Chapter 11, “Defining Service Coefficients and Options” on page 83.

You can define multiple performance periods for work in subsystems which use address spaces or enclaves to represent individual work requests. For a list of subsystems for which you can specify multiple periods, see “Subsystem Support for Goal Types and Multiple Periods” on page 18.

Multiple periods are not supported for work in the IMS and CICS subsystem work environments because service units are accumulated to the address space, not the individual transactions. So, the system cannot track a duration for those transactions.

Multiple periods are also not supported for work in the SYSH subsystem.

## Defining Goals Appropriate for Performance Periods

As you go from one performance period to the next, you can change the type of goal. Goals should become less stringent going from one period to the next. A prime example would be changing to a velocity or discretionary type goal in the last period.

## Using Importance Levels in Performance Periods

Importance levels should stay the same or decrease as the transactions move from one performance period to the next. Remember that importance applies only if a goal is not being met during the duration of the period.

### Examples of Multiple Performance Period Goals

Decreasing stringency of goal and decreasing importance from one period to the next

Service Class = BATCHX

Period 1

Velocity = 50  
Importance = 3  
Duration = 2500 SU

Period 2

Velocity = 15  
Importance = 5



---

## Chapter 9. Defining Classification Rules

*Classification rules* are the rules you define to categorize work into service classes, and optionally report classes, based on *work qualifiers*. A work qualifier is what identifies a work request to the system. The first qualifier is the subsystem type that receives the work request.

There is one set of classification rules in the service definition for a sysplex. They are the same regardless of what service policy is in effect; a policy cannot override classification rules. You should define classification rules after you have defined service classes, and ensure that every service class has a corresponding rule.

The full list of work qualifiers and their abbreviations is:

<b>AI</b>	Accounting information
<b>CI</b>	Correlation information
<b>CN</b>	Collection name
<b>CT</b>	Connection type
<b>LU</b>	Logical Unit name
<b>NET</b>	Netid
<b>PC</b>	Process name
<b>PF</b>	Perform
<b>PK</b>	Package name
<b>PN</b>	Plan name
<b>PR</b>	Procedure name
<b>PRI</b>	Priority
<b>PX</b>	Sysplex name This is the same as cluster name for the SYSH subsystem.
<b>SE</b>	Scheduling environment name
<b>SI</b>	Subsystem instance
<b>SPM</b>	Subsystem parameter
<b>SSC</b>	Subsystem collection name
<b>SY</b>	System name
<b>TC</b>	Transaction class/job class
<b>TN</b>	Transaction name/job name
<b>UI</b>	Userid

### Notes:

1. Not all work qualifiers are valid for every subsystem type; they are subsystem dependent. For details on which qualifiers are valid for which subsystems, see Table 4 on page 55.
2. For many of these qualifiers, you can specify **classification groups** by adding a G to the type abbreviation. (For example, a transaction name group would be TNG.) See "Using Groups" on page 75 for more information.

A singular classification rule is a work qualifier and an associated service class or report class. The following example is a classification rule:

Example of a Classification Rule						
Subsystem Type . . :		IMS		Fold qualifier names?		Y (Y or N)
Description . . .		IMS medium interactive				
		-----Qualifier-----			-----Class-----	
Action	Type	Name	Start		Service	Report
				DEFAULTS:	IMSMED	
___	1	___	___		___	___

**Note:** The “Fold qualifier names” option, set to the default Y, means that the qualifier names will be folded to upper case as soon as you type them in and then press Enter. If you set this option to N, then the qualifier names will remain in the case they are typed in. Leave this option set to Y unless you know that you need mixed case qualifier names in your classification rules.

This example shows that all work coming into any IMS subsystem is associated with service class IMSMED. Service class IMSMED is the default service class for the IMS subsystem type. You can also assign a default report class to a subsystem type.

Since you might not want all work coming into a subsystem assigned to the same service class, or the same report class, you can specify multiple classification rules.

Figure 12 shows two classification rules. In the example, the incoming work request has work qualifiers of subsystem type, job name, job class, accounting information, and userid.

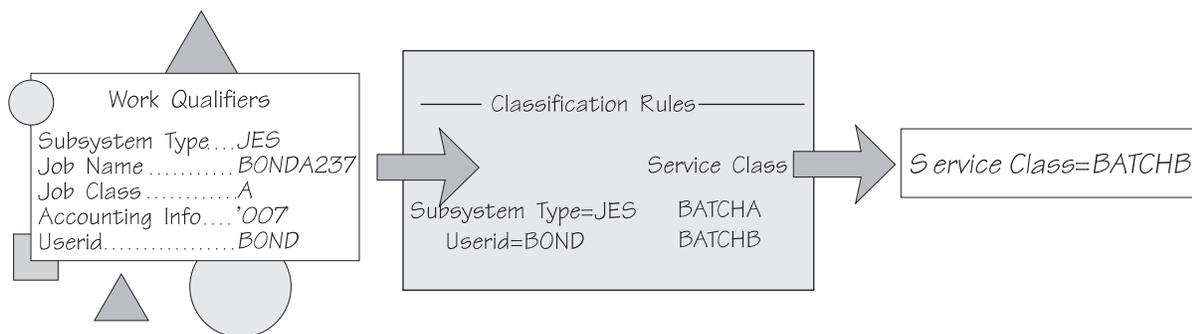


Figure 12. Using Classification Rules to Assign Work to Service Classes

In the example, the service administrator set up classification rules to assign all work coming into JES into service class BATCHA, unless the work had a userid of BOND, in which case, it should be assigned to service class BATCHB. For JES classification, you do not need to specify JES2 or JES3.

For example, if you want all CICS work to go into service class CICSB except the following:

- You want work originating from LU name LONDON to run in service class CICS D
- You want work originating from LU name PARIS to run in service class CICS A, unless:

- The work is from the PAYROLL application, in which case you want it to go into service class CICSC.

You could specify the following classification rules:

Example of Multiple Classification Rules				
Subsystem Type . . . . . CICS				
-----Qualifier-----			-----Class-----	
Type	Name		Service	Report
			DEFAULT: CICSB	_____
1	LU	LONDON	CICSD	_____
1	LU	PARIS	CICSA	_____
2	TN	PAYROLL	CICSC	_____

The previous example has two classification rules with level 1 qualifiers: LU name LONDON and LU name PARIS. Under PARIS there is a level 2 qualifier with transaction name PAYROLL. The PAYROLL qualifier applies only to transactions associated with the level 1 qualifier of PARIS.

In this example, if a work request comes in from an LU other than London or Paris, then it is assigned CICSB. If another work request comes in from Paris, and is from the payroll application, it is assigned to the service class CICSC. If a work request is from the payroll application, but came in from a system in London, then it is assigned to the service class CICSD.

The order of the nesting and the order of the level 1 qualifiers, determine the hierarchy of the classification rules. The application supports eight characters for each rule. For more information about defining the hierarchy of the classification rules, see “Defining the Order of Classification Rules” on page 68.

---

## Defining Classification Rules for Each Subsystem

Work qualifiers depend on the subsystem that first receives the work request. When you are defining the rules, start with the service classes you have defined, and look at the type of work they represent. Determine which subsystem type or types process the work in each service class. Then understand which work qualifiers your installation could use for setting up the rules. It may be that your installation follows certain naming conventions for the qualifiers for accounting purposes. These naming conventions could help you to filter work into service classes. Also, understand which work qualifiers are available for each subsystem type. You can then decide which qualifiers you can use for each service class.

The following table shows the IBM-supplied subsystem types that workload management supports, the kind of work they run, whether they use address space-oriented transactions or enclaves (see special note for CICS and IMS), and where to go for more information. (Unless otherwise noted below, look for “Workload Manager” in each book.) A comparison of the various transaction types is shown in Table 3 on page 55.

**Note:** Enclaves are transactions that can span multiple dispatchable units in one or more address spaces. See *z/OS MVS Programming: Workload Management Services* for more information on enclaves.

Table 1. IBM-Defined Subsystem Types

Subsystem Type	Work Description	Enclave, Address Space, or LPAR	For More Information, See:
<b>ASCH</b>	The work requests include all APPC transaction programs scheduled by the IBM-supplied APPC/MVS transaction scheduler.	Address Space	<ul style="list-style-type: none"> <li>• <i>z/OS MVS Planning: APPC/MVS Management</i></li> </ul>
<b>CB</b>	The work requests include all Component Broker client object method requests.	Enclave	<ul style="list-style-type: none"> <li>• The online information included with the Component Broker system management user interface</li> </ul>
<b>CICS</b>	The work requests include all transactions processed by CICS Version 4 and above.	(See note below.)	<ul style="list-style-type: none"> <li>• <i>CICS/ESA Performance Guide</i></li> <li>• <i>CICS/ESA Dynamic Transaction Routing in a CICSplex</i></li> </ul>
<b>DB2</b>	The work requests include only the queries that DB2 has created by splitting a single, larger query and distributed to remote systems in a sysplex. The local piece of a split query, and any other DB2 work, is classified according to the subsystem type of the originator of the request (for example, DDF, TSO, or JES).	Enclave	<ul style="list-style-type: none"> <li>• <i>DB2 for OS/390 and z/OS Data Sharing: Planning and Administration</i></li> </ul>
<b>DDF</b>	The work requests include all DB2 distributed data facility (DB2 Version 4 and above) work requests.	Enclave	<ul style="list-style-type: none"> <li>• <i>DB2 for OS/390 and z/OS Data Sharing: Planning and Administration</i></li> </ul>
<b>IMS</b>	The work requests include all messages processed by IMS Version 5 and above.	(See note below.)	<ul style="list-style-type: none"> <li>• <i>IMS/ESA Administration Guide: System</i></li> </ul>
<b>IWEB</b>	The work requests include all requests from the world-wide-web being serviced by the Internet Connection Server (ICS), Domino Go Webservice, or IBM http Server for OS/390. As of OS/390 Release 6, these requests also include those handled by the Secure Sockets Layer (SSL). As of OS/390 Release 7, this also includes transactions handled by the Fast Response Cache Accelerator.	Enclave	<ul style="list-style-type: none"> <li>• <i>z/OS HTTP Server Planning, Installing, and Using</i></li> </ul>
<b>JES</b>	The work requests include all jobs that JES2 or JES3 initiates.	Address Space	<ul style="list-style-type: none"> <li>• <i>z/OS JES2 Initialization and Tuning Guide</i></li> <li>• <i>z/OS JES3 Initialization and Tuning Guide</i></li> </ul>
<b>LSFM</b>	The work requests include all work from LAN Server for MVS.	Enclave	
<b>MQ</b>	The work requests include MQSeries Workflow work such as new client server requests, activity executions, activity responses, and sub-process requests.	Enclave	<ul style="list-style-type: none"> <li>• <i>MQSeries Workflow for OS/390 System Management Guide</i></li> </ul>
<b>OMVS</b>	The work requests include work processed in OS/390 UNIX System Services forked children address spaces. (Work that comes from an enclave is managed to the goals of the originating subsystem.)	Address Space	<ul style="list-style-type: none"> <li>• <i>z/OS UNIX System Services Planning</i></li> </ul>
<b>SOM</b>	The work requests include all SOM client object class binding requests.	Enclave	<ul style="list-style-type: none"> <li>• <i>OS/390 SOMobjects Configuration and Administration Guide</i></li> </ul>

Table 1. IBM-Defined Subsystem Types (continued)

Subsystem Type	Work Description	Enclave, Address Space, or LPAR	For More Information, See:
STC	The work requests include all work initiated by the START and MOUNT commands. STC also includes system component address spaces such as the TRACE and PC/AUTH address spaces.	Address Space	<ul style="list-style-type: none"> <li>“Using the System Supplied Service Classes” on page 76</li> </ul>
TSO	The work requests include all commands issued from foreground TSO sessions.	Address Space	<ul style="list-style-type: none"> <li><i>z/OS MVS Initialization and Tuning Guide</i> (Look for the discussion of TSO/E transactions in the “System Resources Manager” chapter.)</li> </ul>
SYSH	Identifies non z/OS partitions (e.g. LINUX partition) in the LPAR cluster which needs to be managed by WLM according to business goals set for the partition.	LPAR	<ul style="list-style-type: none"> <li>“Non z/OS Partition CPU Management” on page 22 and Chapter 18, “The Intelligent Resource Director” on page 159</li> </ul>

#### Important Note on CICS and IMS Transactions

CICS and IMS do not use enclaves, but use a different set of WLM services to provide transaction management.

CICS and IMS transactions can be assigned only response time goals (either percentile or average) within single period service classes. If you do not define any goals at all for CICS or IMS work, then the work will be managed to the velocity goals of the address spaces. Once you have defined a transaction goal for CICS or IMS work, then *all* subsequent work will be managed to those transaction goals, *not* to the velocity goals of the address spaces.

For example, you may initially be managing all CICS work to the velocity goals of the CICS address space. If you define a response time goal for a CICS transaction, you will be required to declare a default goal as part of that definition. Now all CICS transactions will be managed to those response time goals, even if they must accept the default.

Table 2 on page 54 shows some selected non-IBM subsystems, with information on how customers can define classification rules for these WLM exploiters. The information here may be incomplete and may not reflect recent changes to the exploiting subsystems. Always check the related subsystem documentation for the latest information.

Table 2. Non-IBM Subsystems

Subsystem	Work Description (with Allowable Work Qualifiers, Goals, and Periods)	For More Information, See:
<p><b>SAP R/3</b></p>	<p>Each SAP R/3 work process has an associated type, corresponding to the allowable transaction name qualifiers listed below. The ICLI server maps each work process type to one of eight pre-existing enclaves, each with its own performance goal. These enclaves are started at ICLI server startup time, and are deleted when the ICLI server is shut down. Therefore, only velocity goals are appropriate, with single periods.</p> <p><b>Allowable Work Qualifiers:</b></p> <ul style="list-style-type: none"> <li>• Transaction Name — must be one of:               <ul style="list-style-type: none"> <li>– DIALOG</li> <li>– BATCH</li> <li>– UPDATE</li> <li>– UPDATE2</li> <li>– GENERIC</li> <li>– SPOOL</li> <li>– ENQUEUE</li> <li>– UNKNOWN</li> </ul> </li> <li>• Userid — the userid of the user who started the ICLI server</li> </ul> <p>When defining the classification rules for SAP, use the Userid (UI) qualifier with the user who started the ICLI server, and then the Transaction Name (TN) qualifiers (with the exact names shown above) as the subrules. Example:</p> <pre>Subsystem Type . . . . SAP  ----Qualifier-----          -----Class----- Type      Name                Service  Report                 DEFAULT: DISCRET  _____ 1  UI      ICLIRUN              _____ 2  TN      GENERIC              LOWVEL   _____ 2  TN      DIALOG                DIALOG   _____ 2  TN      UPDATE                 UPDATE   _____ 2  TN      UPDATE2               HIGHVEL  _____ 2  TN      SPOOL                  LOWVEL   _____ 2  TN      BATCH                  BATCH    _____</pre> <p>Note: In this example, ENQUEUE and UNKNOWN work defaults to the DISCRET service class.</p>	<ul style="list-style-type: none"> <li>• <i>SAP R/3 on DB2 for OS/390: Planning Guide</i></li> </ul>

If you have a subsystem *not* included in either of the above tables, you should check its documentation for the kind of work requests supported by workload management and the applicable work qualifiers.

The following table summarizes the key differences among the service classes for enclave transactions, address space-oriented transactions, and IMS/CICS transactions.

Table 3. Enclave Transactions, Address Space-Oriented Transactions, and CICS/IMS Transactions

Transaction Type	Allowable Goal Types	Allowable Number of Periods	RMF (or Other Monitor) Reporting
<b>Address Space-Oriented</b>	<ul style="list-style-type: none"> <li>Response Time</li> <li>Execution Velocity</li> <li>Discretionary</li> </ul>	Multiple	<ul style="list-style-type: none"> <li>IOC, CPU, MSO, and SRB service consumption reported</li> <li>Execution delays reported</li> </ul>
<b>Enclave</b>	<ul style="list-style-type: none"> <li>Response Time</li> <li>Execution Velocity</li> <li>Discretionary</li> </ul>	Multiple	<ul style="list-style-type: none"> <li>CPU service consumption reported</li> <li>Execution delays reported</li> <li>“Served by” reported for enclaves using TCBS</li> </ul>
<b>CICS/IMS</b>	<ul style="list-style-type: none"> <li>Response Time</li> </ul>	1	<ul style="list-style-type: none"> <li>No service consumption reported (reported under regions)</li> <li>No execution delays reported (reported under regions)</li> <li>“Service Classes Being Served” reported (for service classes assigned to the server address spaces)</li> <li>“Response Time Breakdown in Percentage” reported</li> </ul>

The ISPF application provides these subsystem types as a selection list on the classification rules panel. You can add any additional subsystem type if it supports workload management on the same panel.

Table 4 shows which work qualifiers are supported by each IBM-defined subsystem.

Table 4. Qualifiers Supported by Each IBM-Defined Subsystem Type

	A S C H	C B	C I C S	D B 2	D D F	I M S	I W E B	J E S	L S F M	M Q	O M V S	S O M	S T C	T S O	S Y S H
Accounting Information	●			●	●			●			●		●	●	
Collection Name		●		●	●							●			
Connection Type				●	●										
Correlation Information				●	●										
LU Name			●	●	●	●									
Netid				●	●	●									
Package Name				●	●										
Perform								●					●	●	
Plan Name				●	●										
Priority				●				●		●					
Procedure Name				●	●										
Process Name										●					

Table 4. Qualifiers Supported by Each IBM-Defined Subsystem Type (continued)

	A S C H	C B	C I C S	D B 2	D D F	I M S	I W E B	J E S	L S F M	M Q	O M V S	S O M	S T C	T S O	S Y S H
Scheduling Environment Name				●				●							
Subsystem Collection Name				●				●							
Subsystem Instance		●	●	●	●	●	●	●	●	●					
Subsystem Parameter				●			●			●		●	●		
Sysplex Name	●	●	●	●	●	●	●	●	●	●	●	●	●	●	●
System Name	●										●		●	●	●
Transaction Class/Job Class	●	●		●		●	●	●		●					
Transaction Name/Job Name	●	●	●	●		●	●	●	●	●	●		●		
Userid	●	●	●	●	●	●	●	●		●	●	●	●	●	

For information about which qualifiers are used by other subsystems supporting workload management, see their supporting documentation.

## Defining Work Qualifiers

The name field for work qualifiers is 8 characters long. You can use nesting for the work qualifiers that run longer than 8 characters. These are accounting information, correlation information, collection name, procedure name, and subsystem parameter. See “Organizing Work for Classification” on page 73 for how to nest using the start position.

You can use masking and wild card notation to group qualifiers that share a common substring. For accounting information and the subsystem parameter, you can use a start position to indicate how far to index into the character string.

The following list explains each of the qualifiers.

• **Accounting Information**

- ASCH** The information passed on the JOB statement.
- DB2** Accounting information is associated with the originator of the query; for example, the accounting information from TSO, JES, or DDF.
- DDF** Accounting information is the value of the DB2 accounting string associated with the DDF server thread.
- JES** The information passed on the JOB statement, not the EXEC statement.

**Note:** WLM/SRM ignores this if the JES2 init deck has a JOBDEF statement with parm ACCTFLD=OPTIONAL specified.

**OMVS**

Accounting data is normally inherited from the parent process of an OS/390 UNIX System Services address space. In addition, when a daemon creates a process for another user, accounting data is taken from the WORKATTR of the RACF user profile. A user can also assign accounting data by setting the `_BPX_ACCT_DATA` environment variable or by passing accounting data on the interface to the `_spawn` service. For more information about OS/390 UNIX System Services accounting information, see *z/OS UNIX System Services Planning*.

- STC** The information passed on the JOB statement.
- TSO** The accounting information specified by the TSO/E user when logging on to the system.

Because JCL supports 143 characters in accounting information, and the application allows only eight characters per rule, the application allows “nesting” for accounting information. See “Using Inheritance in Classification Rules” on page 68 for more information.

**Example of Nesting Accounting Information**

In this example, you can “nest” accounting information (AI) in the classification rules.

Subsystem Type . . . . . JES (Required)  
 Description . . . . . All batch rules

-----Qualifier-----			-----Class-----	
Type	Name	Start	Service	Report
			DEFAULTS: BATH00EY	_____
1	AI	43876AAA 2		_____
2	AI	DEPT58* 10	BATHBEST	_____

This example shows the classification rules for the JES subsystem. You can classify with more than the allowed 8 characters by nesting accounting information. In the example, all work with accounting information ‘43876AAADEPT58’ starting in position 2 for 14 characters is associated with service class BATHBEST.

- **Collection Name**

**CB** The logical server group name defined using the Component Broker system management utility. This represents a set of Component Broker objects that are grouped together and run in a logical server. For more information, see the online information included with the Component Broker system management user interface.

**DB2** The collection name associated with the originator of the query; for example, the collection name from DDF.

**DDF** The DB2 collection name of the first SQL package accessed by the distributed relational database architecture (DRDA) requestor in the work request.

**SOM** The logical server name defined using the SOM REGIMPL utility (defined to REGIMPL as the application alias). This represents a set of SOM objects that are grouped together and run in a logical server. For more information, see *OS/390 SOMobjects Configuration and Administration Guide*.

- **Connection Type**

**DB2** Connection type associated with the originator of the query; for example, the connection type from DDF.

**DDF** The DB2 connection type of the DDF server thread. The thread contains the value 'DIST ' indicating it is a server.

- **Correlation Information**

**DB2** Correlation ID associated with the originator of the query; for example, the correlation ID from DDF.

**DDF** The DB2 correlation ID of the DDF server thread.

- **LU name and Netid**

LU name and netid are used mostly for qualifying CICS, DB2, DDF, and IMS work. If you want to filter on the fully qualified name, you can use the LU name value of up to 8 characters, and then define a sub-rule of netid for up to 8 characters.

**CICS** Only the LU name is available for CICS. It is the 8-byte NETNAME of the principal facility of the transaction instance. For details of the value of this parameter for non-VTAM terminals, and for transaction-routed transactions see *CICS/ESA Performance Guide*.

**DB2** The LU name and netid associated with the originator of the query; for example, the LU name and netid from DDF. For more information about the format of the LU name, see *DB2 SQL Reference*.

**DDF** The VTAM LU name and netid of the system that issued the structured query language (SQL) request. For more information about the format of the LU name, see *DB2 SQL Reference*.

**IMS** If a transaction comes in from an LU 6.2 device, you can specify both the LU name and the netid. Otherwise, you can specify only the 8 byte LU name.

- **Package Name**

**DB2** Package name of the originator of the query; for example, the package name from DDF.

**DDF** The name of the first DB2 package accessed by the DRDA requestor in the work request.

- **PERFORM**

**JES** The performance group number specified using the PERFORM keyword on the JCL JOB statement.

**STC** One of the following:

- The performance group number specified using the PERFORM keyword on the START command.

- The performance group number specified using the PERFORM keyword on the JCL JOB statement.

**TSO** The performance group number specified on the logon panel.

- **Plan Name**

**DB2** Plan name associated with the originator of the query; for example, the plan name from DDF.

**DDF** The DB2 plan name associated with the DB2 server thread. For DB2 private protocol requestors and DB2 Version 3 or later DRDA requestors, this is the DB2 plan name of the requesting application. For non-DB2 requestors and other DRDA requestors, this is not applicable.

- **Priority**

**DB2** Priority associated with the originator of the query; for example, the priority from a batch job.

**MQ** A value between 0 and 9, the priority associated with the MQSeries Workflow message.

**Note:** For MQ work, 0 is the highest priority and 9 is the lowest. (Contrast with JES work, where 15 is the highest priority and 0 is the lowest.)

**JES** A value between 0 and 15, the priority associated with the batch job submitted through JES2 or JES3. Use of this qualifier requires OS/390 JES2 V2R4 or above, or JES3 V2R8 or above.

**Note:** For JES work, 15 is the highest priority and 0 is the lowest. (Contrast with MQ work, where 0 is the highest priority and 9 is the lowest.)

When you use priority as a work qualifier, you can use operators such as greater-than ('>') and less-than ('<') to group a range of priorities into one service or report class.

**Priority Example**

To put priority 8 and higher work into service class BATCH020, and put all other work into service class BATCH005, you would code the following:

```
Subsystem Type . . . . . JES (Required)
Description . . . . . Job Priority
```

-----Qualifier-----			-----Class-----	
Type	Name	Start	Service	Report
			DEFAULTS: BATCH005	_____
1	PRI	>=8	BATCH020	_____

- **Procedure Name**

**DB2** Procedure name associated with the originator of the query; for example, the DB2 stored procedure name from DDF.

**DDF** If the first SQL statement issued by the DDF client is a CALL statement, this field contains the unqualified name of the DB2 stored procedure. In all other cases, this field contains blanks.

- **Process Name**

**MQ** The MQWIH\_ServiceName from the message's work information header.

- **Scheduling Environment Name**

**Note:** This qualifier is available on OS/390 Release 10 with APAR OW43813 installed.

**JES** The scheduling environment name assigned to the job.

**DB2** Scheduling environment name associated with the originator of the query.

- **Subsystem Collection Name**

**Note:** This qualifier is available on OS/390 Release 10 with APAR OW43813 installed.

**JES** The XCF group name.

**DB2** Subsystem collection name associated with the originator of the query.

- **Subsystem Instance**

You can use subsystem instance to isolate multiple instances of a subsystem. For example, use subsystem instance if you have a CICS production system as well as a CICS test system.

**CB** The Component Broker-specific server name.

**CICS** The VTAM applid for the subsystem instance. For more information, see *CICS/ESA Dynamic Transaction Routing in a CICSplex*.

**DB2** The subsystem type associated with the originator of the query; for example:

- TSO for requests from TSO/E
- JES for requests from batch jobs
- DDF for requests from DDF

**DDF** The DB2 server's MVS subsystem name. For more information about the name, see *DB2 Administration Guide*.

**IMS** The IMS subsystem name, as defined on the IMSID positional parameter in the IMS DFSMPR procedure. It is a 1- to 4-character value that uniquely identifies the control region. The generation default is IMSA. For more information, see *IMS/ESA System Definition Reference*

**IWEB** The subsystem name from the application environment definition. (Note that this is identical to bits 0-7 of the Subsystem Parameter qualifier for IWEB).

**JES** The JES2 or JES3 subsystem name from the IEFSSNxx parmlib member.

**LSFM** The proname of address space in which LAN Server for MVS is running.

**MQ** The MQSeries Workflow subsystem name.

- **Subsystem Parameter**

If you have a vendor or home-grown subsystem type that has a qualifier other than the IBM supported ones, it could use the subsystem parameter. You should check your subsystem documentation to determine whether your subsystem supports the subsystem parameter, and in what parameter format.

Because the subsystem parameter is up to 256 characters long, you can nest to use more than the limit of eight characters. Nesting is allowed only for qualifiers longer than 8 characters. These are: accounting information, correlation information, connection type, and subsystem parameter.

**DB2** The subsystem parameter, if any, associated with the originator of the query.

**IWEB** A 47-byte string formatted as follows:

<b>1-8</b>	Subsystem name
<b>9</b>	Blank
<b>10-24</b>	Source IP address
<b>25</b>	Blank
<b>26-40</b>	Target IP address
<b>41</b>	Blank
<b>42-47</b>	Target port

For more information, see *Internet Connection Server User's Guide*.

**MQ** The 32-byte *application\_environment\_name* from the APPLICID attribute of the process definition associated with the WLM-managed queue.

**SOM** A 246-byte string consisting of two 123-byte fields:

- Field 1 — class name
- Field 2 — method name

For more information, see *OS/390 SOMobjects Configuration and Administration Guide*.

**STC** Indicates the system-provided service class name that will be assigned if a started task created with the high dispatching priority, privileged, or system task attribute is not assigned to a service class.

Values:

- SYSTEM — Started task was created with high dispatching priority attribute.
- SYSSTC — Started task is privileged or is a system task.
- (blank) — Started task was not created with the high dispatching attribute, is not privileged, and is not a system task.

- **Sysplex Name**

For all subsystem types, use the sysplex name qualifier if you have a common service definition in multiple sysplexes, and need to assign different service classes or report classes based on the specific sysplex in which the work is running.

**Note:** This qualifier is available on OS/390 Release 10 with APAR OW43813 installed.

- **System Name**

The system name qualifier is supported for address spaces whose execution system is known at classification time. Note that JES is not eligible for this qualifier, as the system on which classification occurs may not be the system on which the job is run. Subsystem-defined transactions (CICS/IMS) and enclave-based transactions are not bound to an execution system at classification time, and are therefore not eligible either.

**Note:** This qualifier is available on OS/390 Release 10 with APAR OW43813 installed.

**ASCH** The name of the execution system.

**OMVS**

The name of the execution system.

**STC** The name of the execution system.

**TSO** The name of the execution system.

**SYSH** The name of the execution system.

- **Transaction Class/Job Class**

**ASCH** The job class used for work selection.

**CB** Name resulting from mapping the URI to a name. For more information, see *WebSphere Application Server for z/OS and OS/390: Installation and Customization* manual.

**DB2** The job or transaction class associated with the originator of the query; for example, the job class from JES.

**IMS** The CLASS keyword on the PGMTYPE=parameter in the APPLCNT macro. For more information, see *IMS/ESA System Definition Reference*.

**IWEB** The arbitrary class name specified in the APPLENV directive in the Webserver's administrative file. Using the filtering function in the webserver, you can assign transactions to transaction classes based on the requested URL. The transaction classes can then be assigned unique service classes using this Transaction Class qualifier. Note that this is probably the most useful qualifier for IWEB work, because of its flexibility. For more information, see *z/OS HTTP Server Planning, Installing, and Using*.

**Special Note for the Fast Response Cache Accelerator**

In OS/390 Release 7, there is a new function called the Fast Response Cache Accelerator for high performance handling of cached static web pages. **You must classify this work as described below. Otherwise, it will be assigned the default service class for IWEB work.**

The transactions handled by the Cache Accelerator are all joined to a single, long-lived enclave. This enclave should be assigned a unique transaction class (as specified on the Webserver FRCAWLMParms directive). This transaction class should then be assigned to a service class with a single period and a velocity goal in the service policy under the IWEB subsystem type. Neither response time goals nor multiple periods are appropriate for this work, as WLM is not aware of the individual Cache Accelerator requests. (Because each individual transaction is so trivial, it would cost more resource to manage them than to just process them.) In RMF reports, you will see zero ended transactions for the Cache Accelerator service class (assuming you have no other work running in this service class), but you will see some amount of accumulated service for this single enclave.

**JES** The job class used for work selection.

**MQ** A value of either ONLINE, meaning that the server is immediately available, or BACKLOG, meaning that the message was queued pending availability of a server.

- **Transaction Name/Job Name**

**ASCH** The jobname in the JCL JOB statement in the APPC/MVS transaction program (TP) profile.

**CB** The method name, for example, GET, POST, or DELETE. For more information, see *WebSphere Application Server for z/OS and OS/390: Installation and Customization* manual.

**CICS** A parameter on many CICS commands. It is often referred to as the CICS transaction identifier, or *transid*. For more information, see *CICS/ESA Resource Definition Guide*.

**DB2** The transaction or job name associated with the originator of the query; for example, the job name from JES.

**IMS** The CODE= parameter on the IMS TRANSACT macro. For more information, see *IMS/ESA System Definition Reference*.

**IWEB** The method name, for example, GET, HEAD, POST, DELETE, or PUT.

**JES** The jobname of the JES managed job. For example, you may run a CICS region as a batch job in your installation. You would define it in classification rules as a transaction name in the JES subsystem type.

**LSFM** One of the following:

- LSFMMMTX - multi-media transactions
- LSFMFITX - file transactions
- LSFMAMTX - administration transactions
- LSFMCMTX - communication transactions

**MQ** The MQWIH\_ServiceStep value from the message's work information header.

**OMVS**

The jobname for the OS/390 UNIX System Services address space. By default, fork and spawn set jobname values to the userid with a number (1-9) appended. However, daemons or users with appropriate privileges can set the \_BPX\_JOBNAME environment variable to change the jobname for forked or spawned children. In this way, servers and daemons in OS/390 UNIX System Services address spaces can easily be assigned performance attributes different than other OS/390 UNIX System Services address spaces.

**STC** One of the following:

- The name as specified on the JOBNAME= parameter of the START command
- The name specified on the MOUNT command
- The system address space name
- The name on the JOB statement.

For example, you may run your IMS regions as started tasks in your installation. You would define these as transaction names in your STC subsystem type in classification rules. However, if IMS V5.1 is present, the STC rule for IMS is ignored and the rules for the IMS subsystem type are used when IMS becomes a server.

- **Userid**

**ASCH** The userid of the user requesting the APPC/MVS service.

**CB** The userid of the user requesting the Component Broker service.

**CICS** The userid specified at LOGON time, which is the RACF (or other access control facility) defined resource. For more information about CICS userids, see *CICS/ESA CICS-RACF Security Guide*.

**DB2** The userid associated with the originator of the query; for example, the userid from TSO, JES, or DDF.

**DDF** The DDF server thread's primary AUTHID, after inbound name translation.

**IMS** The userid specified at LOGON time, which is the RACF (or other access control facility) defined resource.

**IWEB** The userid of the web server address space (not the original requestor's userid). Note that because this userid will generally be the same for all transactions, using this qualifier for IWEB work will have limited usefulness.

**JES** The userid specified on the JOB statement on the RACF USER keyword.

**MQ** The first 8 bytes of the 12-byte message header field MQMD\_USERIDENTIFIER.

**OMVS**

The RACF userid associated with the address space. This userid is either inherited from the parent process or assigned by a daemon process (for example, the rlogin or telnet daemon). For more information about OS/390 UNIX System Services userids, see *z/OS UNIX System Services Planning*.

**SOM** The userid of the user requesting the SOM service.

**STC** The userid assigned to the started task by RACF (or other access control facility).

**TSO** The userid specified at LOGON time, which is the RACF (or other access control facility) defined user profile.

---

## Defining the Order of Classification Rules

When the subsystem receives a work request, the system searches the classification rules for a matching qualifier and its service class or report class. Because a piece of work can have more than one work qualifier associated with it, it may match more than one classification rule. Therefore, the order in which you specify the classification rules determines which service classes are assigned.

Previously, prior to z/OS 1.3, with the IEAICSxx, the system used a set order for searching the work qualifiers. With workload management, there is no set order. You define it. Only the subsystem type must be the first level qualifier for classification rules. You determine the rest of search order by the order in which you specify the classification rules. You can use a different hierarchy for each subsystem.

### Example of Defining an Order of Rules

Suppose you are defining your JES classification rules. For all work requests coming into the JES2 subsystem instance, you want to assign work with userid MNGRBIG and jobname PERFEVAL to service class BATSLOW. All other work with userid beginning with MNGR should be assigned to service class BATHOT. So in this case, the hierarchy for the JES subsystem is:

1. Subsystem type, because it is always the first
2. Subsystem instance
3. Userid
4. Jobname

## Defining a Subsystem Service Class Default

A service class default is the service class that is assigned if no other classification rule matches for that subsystem type. If you want to assign any work in a subsystem type to a service class, then you must assign a default service class for that subsystem — except for STC. You are not required to assign a default service class for STC, even if you assign started tasks to different service classes.

Optionally, you can assign a default report class for the subsystem type. If you want to assign work running in a subsystem to report classes, then you do not have to assign a default service class for that subsystem.

### Using Inheritance in Classification Rules

Keep in mind that if you leave out the service class or report class on a classification rule, the work inherits the service class or report class from the higher level rule.

You can use this to your advantage by using the qualifier that applies to the most service classes in the subsystem as your first level qualifier. You can then list the exceptions in the subsequent levels.

### Example of Using Inheritance

For example, for CICS, the available qualifiers include subsystem instance, userid, transaction name, and LU name.

```

Subsystem Type . . . . . : CICS
Description . . . . . : CICS subsystem
-----Class-----
Service   Report
DEFAULTS: CICSB
          CICSA   ATMA
          _____
          CASHA   CASH
          _____
          DEPOSITA DEPOSIT
          _____
          BIGDEP
          _____
          CICSC   ATMC
          _____
          CASHC   CASH
          _____
          DEPOSITC DEPOSIT
          _____
          BIGDEP
  
```

In this example, the installation set up their userids for their CICS work according to the ATM set-up that they have. Since all of their interactive work is related to the ATMs, they chose userid as their first level qualifier. Then, they wanted to separate out their cash transactions from their deposit transactions for reporting purposes, so they set up a report class for each.

The transactions do not have a service class explicitly assigned to them, so they inherit the service class from the rule one level before.

In addition, for the deposit transactions, they wanted to separate out those deposits coming from ATMs on Wall St location, because that area had been having some service troubles. So they defined a report class at a level 3 under the DEPOSIT transactions for each ATM userid.

You cannot nest most qualifier types within themselves. For example, if you choose userid as a first level, you cannot use userid as a second level, or sub-rule qualifier. Nesting is allowed only for qualifiers longer than 8 characters. These are: accounting information, correlation information, collection name, procedure name, and subsystem parameter. See "Organizing Work for Classification" on page 73 for how to nest using the start position.

Keep in mind that the system sequentially checks for a match against each level one rule. When it finds a match, it continues just in that family through the level two rules for the first match. Similarly, if a match is found at any given level, then its sub-rules are searched for further qualifier matches. The last matching rule is the one that determines the service class and report class to be assigned.

**Example 1: Using the Order of Rules**

Suppose you have defined the following classification rules for your IMS work:

			-----Class-----
			Service      Report
		DEFAULTS:	IMSB      _____
1	Tran Name	5128	IMSA      _____
2	LU Name	BERMU	IMSC      ____      _____

Suppose the following kind of work requests enter the system:

	Transaction 1	Transaction 2	Transaction 3
Tran Name	5128	6666	5128
LU Name	BERMU	BERMU	CANCUN

- Transaction 1 is assigned service class IMSC, because the transaction 5128 is from the BERMU LU name.
- Transaction 2 is assigned service class IMSB, the subsystem default, because it is not transaction 5128, and therefore the system never checks any sub-rules.
- Transaction 3 is assigned service class IMSA, because it is associated with LU name CANCUN.
- If you specified the classification rules as:

1	LU Name	BERMU		IMSC	_____
2	Tran Name	5128		IMSA	____      _____

Then all work from LU BERMU is assigned in service class IMSC, except work with Tran name 5128.

### Example 2: Using the Order of Rules

Suppose you have defined the following rules for your IMS work:

-----Qualifier----		-----Class-----	
Type	Name	Service	Report
		DEFAULTS: PRDIMSR	_____
1	SI	IMST	TRNIMSR
2	TC	15	TRNIMSNR
1	TC	15	PRDIMSNR
1	SI	IMSM	MDLIMSR
2	TC	15	MDLIMSNR

- If a work request in transaction class (TC) 15 enters the system from subsystem instance (SI) IMST, it is assigned service class TRNIMSNR.
- If a work request in transaction class (TC) 15 enters the system from subsystem instance (SI) IMSM, it is assigned service class PRDIMSNR and not MDLIMSNR as you might expect.

This is because the level 1 classification rule

1	TC	15	PRDIMSNR	_____
---	----	----	----------	-------

is above the more explicit classification rule:

1	SI	IMSM	MDLIMSR	_____
2	TC	15	MDLIMSNR	_____

The system stopped on the first level one match that it encountered. You can re-order the rules so that this does not occur. Put the most explicit rule above the more general one as shown below.

-----Qualifier----		-----Class-----	
Type	Name	Service	Report
		DEFAULTS: PRDIMSR	_____
1	SI	IMST	TRNIMSR
2	TC	15	TRNIMSNR
1	SI	IMSM	MDLIMSR
2	TC	15	MDLIMSNR
1	TC	15	PRDIMSNR

### System-Provided Service Classes

If some work comes into a system for which there is no associated service class defined in the classification rules, workload management assigns it to a default service class. There are three such default service classes:

#### SYSTEM

For all system address spaces designated 'high dispatching priority' (X'FF') address spaces. The high dispatching priority address spaces include: MASTER, GRS, DUMPSRV, SMF, CATALOG, RASP, XCFAS, SMXC, CONSOLE, IOSAS, and others. For a list of the high dispatching priority address spaces in your installation, see the RMF Monitor II report and look for the x'FF' dispatching priority.

You do not need to set up service classes for these system address spaces. Workload management recognizes these as special system address spaces and treats them accordingly.

If for some reason you do want to control these address spaces, you can do the following:

- Define a service class for them

- Set up a classification rule in the STC subsystem type which assigns the address space to a service class other than the default STC service class.

**Note:** You cannot control \*MASTER\*, INIT, or WLM in a service class.

When you assign a service class other than SYSTEM to a started task eligible for the SYSTEM service class, it loses the high dispatching priority attribute and runs at the dispatching priority of the assigned service class period. The high dispatching priority attribute can be restored by one of the following methods:

- You can use the RESET command to change the started task's service class to SYSTEM.
- You can change the classification rules to explicitly classify the started task to SYSTEM and activate a policy.

You can also assign work to the SYSTEM service class as part of your work classification rules. You can *only* do this, however, for classification rules in the STC subsystem type, and only for address spaces that are designated as "high dispatching priority" address spaces.

**Note:** If you assign a started task to either the SYSTEM or the SYSSTC service class in a sysplex running an MVS level earlier than OS/390 R4, a toleration PTF for APAR OW25831 is needed on the downlevel systems.

For more information about using SYSTEM in classification rules for started tasks, see "Using the System Supplied Service Classes" on page 76.

## SYSSTC

For all started tasks not otherwise associated with a service class. Workload management treats work in SYSSTC just below special system address spaces in terms of dispatching.

You can also assign work to the SYSSTC service class as part of your work classification rules. You can do this for classification rules in the following subsystem types:

- ASCH
- JES
- OMVS (OS/390 UNIX System Services)
- STC
- TSO

Some address spaces normally created when running MVS are neither high dispatching priority, privileged, nor a system task, such as NETVIEW. These address spaces must be explicitly assigned to a service class such as SYSSTC.

**Note:** If you assign a started task to either the SYSTEM or the SYSSTC service class in a sysplex running an MVS level earlier than OS/390 R4, a toleration PTF for APAR OW25831 is needed on the downlevel systems.

For more information about using SYSSTC in classification rules for started tasks, see "Using the System Supplied Service Classes" on page 76.

## SYSOTHER

For all other work not associated with a service class. This is intended as a 'catcher' for all work whose subsystem type has no classification. It is assigned a discretionary goal.

---

## Organizing Work for Classification

There are some ways you can organize your qualifiers for easier classification. You can use masking or wildcard notation as a way of grouping work to the same service class or report class. Or, you could set up a qualifier group when the qualifier is a standard 8 character value.

If you have more than five rules at a given level for the same classification qualifier within a subsystem type, there may be performance implications. Qualifier groups are quicker to check than a single rule, so it may make sense to use them for performance sensitive subsystems like CICS and IMS.

For qualifiers longer than 8 characters, you can use the start position. Accounting information, correlation information, connection type, procedure name, and the subsystem parameter are greater than 8 characters. The following sections explain each kind of notation.

### Using the Start Position

For work qualifiers longer than 8 characters, you can use a start position to indicate how far to index into the character string for a match. For example, you can assign all TSO/E users in a department to the same service class, assuming you follow a naming convention for accounting information for the department.

The ISPF administrative application provides the **Start** field where you can specify the starting position. The name field for a work qualifier is 8 characters long. If you are matching on a string less than 8 characters, you must use wildcard notation (asterisk) at the end of the string. Otherwise, the qualifier is padded with blanks to be 8 characters, and the blanks are used when making a match.

### Examples of Using the Start Position

- Assume you want to associate all JES2 work from department IRS with the service class JESFAST. You assigned the default for JES2 work as service class JESMED. If all JES2 accounting information from department IRS have the characters 'DIRS' starting in the seventeenth position, you enter a rule with qualifier DIRS\* to match on just the four characters. If you want to filter out those jobs with the eight characters 'DIRS ' starting in the seventeenth position, you need another rule with qualifier DIRS to assign those jobs to JESMED. The example below shows the rules:

```
Subsystem Type . . . . . : JES
Description . . . . . All JES2 service classes
```

-----Qualifier-----			-----Class-----		
Type	Name	Start	Service	Report	
			DEFAULT: JESMED		
1	AI	DIRS	JESMED		_____
1	AI	DIRS*	JESFAST		_____

In this case, all jobs that have accounting information with the eight characters 'DIRS ' starting in the seventeenth position are assigned to JESMED. All other jobs that have the four characters 'DIRS' starting in the seventeenth position are assigned to JESFAST. All other work coming into JES is assigned to service class JESMED.

## Using Masking Notation

You can use masking notation to replace a single character within a qualifier. This allows any character to match the position in the rule. Use a % in the position where the character would be. You can use multiple masks successively for multiple character replacement. If you specify a mask at the end of a character string, it could match on a null value OR a single character.

### Examples of Masking Notation

For example, suppose you have a naming convention for all users in your IS services department that all userids start with DEPT58, followed by a letter A-F (depending on which division), ended by an I. Suppose also, you would like to bill your IS services department separately. You could use masking notation in setting up the classification rules as shown below.

-----Qualifier-----			-----Class-----		
Type	Name	Start	Service	Report	
			DEFAULTS: BATREG		
1	UI	DEPT58%I			DEPT58

In the example, all work in this subsystem is associated with service class BATREG, and all work from the IS services department is associated with the report class DEPT58.

## Using Wildcarding Notation

You can also use *wildcard* notation for multiple character replacement in a character string. The wildcard character is an asterisk (\*). You can use the wildcard character as the last position of a character string, or by itself. If a character string contains

an asterisk in a position other than the last, it is treated as a character; for example, if you specify a character string of *CI\*S*, the third character in a matching character string **must** have an asterisk (\*) as the third character. An asterisk by itself indicates a match for all characters.

#### Examples of Wildcard Notation

For example, suppose your installation has a naming convention for your CICS AORS and TORS. You can use the following wildcard notation in your CICS classification rules. Note that the subsystem instance of *CI\*S* is not wildcard notation, a matching subsystem instance must be *CI\*S*.

-----Qualifier-----			-----Class-----	
Type	Name	Start	Service	Report
			DEFAULTS: CICSSTC2	_____
1 TN	TOR*	___	CICSSTC1_	
1 TN	AOR*	___	CICSSTC3_	
1 SI	CI*S	___	CICSTEST_	

#### Important Note

Be careful when putting specific definitions below wildcards which might cause an unwanted early match. In the example below, the rule for TOR11 is useless, because a TOR11 transaction will match the TOR\* rule before it.

-----Qualifier-----			-----Class-----	
Type	Name	Start	Service	Report
			DEFAULTS: CICSSTC2	_____
1 TN	TOR*	___	CICSSTC1_	
1 TN	TOR11	___	CICSSTC4_	

## Using Groups

Groups are available for grouping together work qualifiers to make classification simpler. You can create groups to collect together work when you don't have a standard naming convention that allows masking or wildcarding. A *group* is a collection of the same work qualifiers. For example, you may want to create a group of started tasks because you want to assign them all to the same service class.

Groups are allowed for the following work qualifiers:

- Connection type
- LU name
- Netid
- Package name
- Perform
- Plan name
- Subsystem instance
- System Name
- Transaction class
- Transaction name
- Userid

Group types are specified by adding G to the type abbreviation. For example, a transaction name group is indicated as TNG.

Qualifier groups of more than 5 members are quicker to check than single instances in the classification rules. So if you have, for example, a long list of CICS or IMS transaction names that you want to group in a service class or report class, consider setting up a group.

### Examples of Groups

If you want to assign a large number of CICS transactions to the same service class, you can create a transaction name group (TNG). You name the group, for example *CICSCONV*, and list all the transaction names you want included in the group.

```
Qualifier type . . . . . : Transaction Name
Group name . . . . . : CICSCONV (required)
Description . . . . . : CICS Conversational Group
```

--Qualifier--

```
Name
CDBC
CDBI
CDBM
CEBR
CECI
CECS
CEDA
CEDB
```

```
Qualifier type . . . . . : Transaction Name
Group name . . . . . : CICSLONG (required)
Description . . . . . : CICS Long-Running Transactions
```

--Qualifier--

```
Name
CDB0
CSGX
CSNC
CSNE
CSSX
```

Then you use those group names in the classification rules, as shown below.

```
Subsystem Type . . . . . : CICS
Description . . . . . : CICS transactions
```

-----Qualifier-----			-----Class-----	
Type	Name	Start	Service	Report
			DEFAULTS:	
1	TNG	CICSCONV	CICSMED	_____
1	TNG	CICSLONG	CICSCONV	_____
			CICSLONG	_____

## Using the System Supplied Service Classes

You can also take advantage of the system-supplied service classes to simplify the process of defining service classes and classification rules for started tasks.

Use the system-provided service classes *SYSTEM* and *SYSSTC* for your STC service classes. WLM recognizes certain system address spaces when they are created (like *GRS*, *SMF*, *CATALOG*, *MASTER*, *RASP*, *XCFAS*, *SMXC*, *CONSOLE*, *IOSAS*, *WLM*), puts them into the *SYSTEM* service class, and treats them accordingly. If a started task is not assigned to a service class, WLM manages the started task in the *SYSSTC* service class. Started tasks in *SYSSTC* are assigned a high dispatching priority. This is appropriate for started tasks such as *JES* and *VTAM*. Not all

started tasks are appropriate for SYSSTC, as a CPU-intensive started task could use a large amount of processor cycles. However, if your processor is lightly loaded, or in a 6-way, 8-way, or 10-way MP, SYSSTC might be appropriate, because that one task may not affect the ability of the remaining processors to manage the important work with goals.

**Notes:**

1. \*MASTER\*, INIT and WLM always run in the SYSTEM service class and cannot be reassigned via the service definition.
2. If you assign a started task to either the SYSTEM or the SYSSTC service class in a sysplex running an MVS level earlier than OS/390 R4, a toleration PTF for APAR OW25831 is needed on the downlevel systems.

**Using a Sysplex with Systems at OS/390 R4 or Later**

The following example implements started task classification with these assumptions:

- Any started tasks not explicitly classified are given low priority.
- Started tasks are defined in three transaction name groups: HI\_STC, MED\_STC, and LOW\_STC.
- System defaults are used for MVS-owned address spaces.
- Separate reporting is used for Master, DUMPSRV, and GRS.

```

Subsystem Type . . . . . : STC
Description . . . . . : All started tasks
-----Qualifier-----
Type      Name      Start      Service      Report
-----
1 TN      %MASTER%  _____  SYSTEM      MASTER
1 TN      GRS        _____  SYSTEM      GRS
1 TN      DUMPSRV   _____  SYSTEM      DUMPSRV
1 SPM     SYSTEM    _____  SYSTEM
1 SPM     SYSSTC    _____  SYSSTC
1 TNG     HI_STC    _____  SYSSTC
1 TNG     MED_STC   _____  VEL35I3
1 TNG     LOW_STC   _____  VEL15I5
1 TN      %%%DBM1  _____  SYSSTC      DBM1S
1 TNG     DB2       _____  VEL60       DB2S

```

The first rule assigns \*MASTER\* to report class MASTER. (Note that "SYSTEM" is shown in the service class column. In this case, it is not a true assignment as \*MASTER\* must *always* run in the SYSTEM service class anyway.) \*MASTER\* could also be classified using the SPM SYSTEM rule, but since separate reporting is desired for \*MASTER\*, it is classified separately. GRS and DUMPSRV are handled similarly.

The SPM rules assign started tasks created with the high dispatching priority attribute to SYSTEM, and other privileged or system tasks to SYSSTC. This allows you to let MVS manage started tasks that it recognizes as special.

**Notes:**

1. Note that explicitly defining the SPM rules, as in the above example, is optional. If they were removed from the example, then the high dispatching priority work would still be assigned to SYSTEM, and the other privileged or system tasks would still be assigned to SYSSTC, as those are the defaults. The reason for explicitly defining them here with the SPM rules is to protect yourself from inadvertently assigning them elsewhere in the rules that follow the SPM rules.

- Note, also, that the placement of these SPM rules is crucial. If they had been declared first, then the three TN rules intended to assign \*MASTER\*, DUMPSRV, and GRS to report classes would have never been reached.

System tasks are those given the privileged and/or system task attribute in the IBM-supplied program properties table or in the SCHEDxx parmlib member. See the SCHEDxx chapter in *z/OS MVS Initialization and Tuning Reference* for a list of system tasks.

Started tasks in the HI\_STC transaction name group are run in service class SYSSTC, MED\_STC in VEL35I3, and LOW\_STC in VEL15I5. The goals on the service classes use importance to guarantee that started tasks of low value will be sacrificed if necessary to keep medium-value started tasks meeting goals. Note that this does not guarantee any relationship between the dispatching priorities that will be observed for service classes VEL35I3 and VEL15I5.

All instances of IRLM, across all subsystems, are run in SYSSTC. The matching regions are classified to customer-defined service classes to guarantee that IRLM's dispatching priority is always above the dispatching priority of the regions it is serving.

A transaction name group is defined to match all DB2 subsystem regions with 4-character subsystem names not already classified. Such a group for DB2 V4 could be defined as follows:

```
Qualifier type . . . . . : Transaction Name
Group name . . . . . : DB2
Description . . . . . All non-IRLM DB2 regions
Fold qualifier names? . . . . Y (Y or N)

Qualifier Name
%%%DBM1
%%%MSTR
%%%DIST
%%%SPAS
```

Started tasks which do not match any of the classification rules are assigned the default service class DISC. If the default service class were left blank, these started tasks would be assigned to SYSSTC.

### Using a Sysplex with Some Pre-OS/390 R4 Systems

Use the following technique for classifying started tasks if you have any pre-OS/390 R4 systems in your sysplex.

To assign started tasks to the SYSSTC service class, you need to leave the default service class for the STC subsystem type blank. Suppose you have grouped your started tasks into two categories based on how you would like to control them. Suppose that you have defined two transaction name groups for started tasks HI\_STC and MED\_STC. You also want to use the SYSSTC service class for the started tasks listed in HI\_STC. HI\_STC includes started tasks such as VTAM, JES, LLA, and RMF. MED\_STC contains started tasks that require a medium velocity goal, and you want to assign them to the VEL35 service class. You want to assign all other started tasks that are ineligible for SYSTEM or SYSSTC to the DISC service class.

To assign the started tasks to the appropriate service class, specify the following in the classification rules for STC:

Subsystem Type . . . . . : STC  
 Description . . . . . All started tasks

-----Qualifier-----			-----Class-----	
Type	Name	Start	Service	Report
			DEFAULTS: _____	
1	TNG	HI_STC	_____	_____
1	TNG	MED_STC	VEL35	_____
1	TN	*	DISC	_____

Note that you must leave the default service class blank. Also, in the first classification rule, you must leave the transaction name group HI\_STC service class assignment blank. This allows those started tasks in the group HI\_STC to run in the SYSSTC service class.

The second rule assigns the tasks you identified as medium into a service class called VEL35. And the last rule says everything else receives a service class called DISC. Note that workload management recognizes the standard system address spaces (like GRS, DUMPSRV, MASTER, WLM, etc), and runs them in either the appropriate SYSTEM or SYSSTC service class.



---

## Chapter 10. Defining Report Classes

Optionally, classification rules can assign incoming work to a report class. Workload management provides data for reporting on all of the service definition terms on a service class period and workload basis. Report classes can be used to report on a subset of transactions running in a single service class but also to combine transactions running in different service classes within one report class. In the first case, a report class is called homogeneous, in the second case it is called heterogeneous. Heterogeneous report classes can cause incorrect performance data, since the data collected is based on different goals, importance, or duration. WLM allows a caller of the workload reporting service IWMRCOLL to detect whether a report class is homogeneous or heterogeneous for a given report interval.. See *z/OS MVS Programming: Workload Management Services* for more information on how a reporting product can receive an indication whether a report class is homogeneous or heterogeneous.

The data available for report classes include:

- Number of transactions completed
- Average response times
- Resource usage data
- State samples.
- Response time distribution buckets.
- work manager delay data.

Each item in the list is related to a report class period. The reporting products are advised to report response time distributions only for homogenous report class periods, where only 1 service class contributed data to the report class.

Note that for CICS and IMS workloads, resource usage data and state samples are reported with the service classes for the regions, not the service classes assigned to the transactions.

Note that a report class can has many periods as the largest service class within the same service definition has periods. Not all of them may be used, though. Report classes inherit periods automatically based on the transaction's service class period that is contributing to the report class. A reporting product can determine whether a report class is homogeneous within a given reporting interval and is advised to report the data as one period if the report class is heterogeneous. Based on its capabilities and options, a reporting product can report the data of a homogeneous report class as one period or as multiple periods.

You can assign up to a maximum of 999 report classes, with a maximum of one report class per work request or transaction.

### Defining Report Classes

**Name** Report class name

**Description**

Description of the report class

**Name (required)**

Eight character identifier of the report class.

**Description (optional)**

An area of 32 characters to describe the report class.

**Example 1: Defining Report Classes**

Suppose you have defined the following rules for your CICS work:

```
Subsystem Type . . . . . : CICS
Description . . . . . : CICS subsystem
```

			-----Class-----	
			Service	Report
			DEFAULTS: CICSB	
1	UI	ATMA	CICSA	ATMA
2	TN	CASH	_____	CASHA
2	TN	DEPOSIT	_____	DEPOSITA
3	LU	WALLST	_____	BIGDEP
1	UI	ATMC	CICSC	ATMC
2	TN	CASH	_____	CASHC
2	TN	DEPOSIT	_____	DEPOSITC
3	LU	WALLST	_____	BIGDEP

In this example, the cash transactions are separated from their deposit transactions for reporting purposes. Report class ATMA therefore, does not include all transactions with userid ATMA, because it does not include the cash or deposit transactions.

---

## Chapter 11. Defining Service Coefficients and Options

The following options must be specified for workload management:

- How the system calculates the amount of resources work consumes with service coefficients
- Whether workload management is to dynamically set I/O priorities based on performance goals
- Whether workload management is to dynamically reassign parallel access volume alias addresses based on performance goals.

All of these options can be set using the Service Coefficient/Service Definition Options panel in the WLM ISPF application, as shown in “Working with Service Coefficients and Options” on page 195.

---

### Service Definition Coefficients

The amount of system resources an address space or enclave consumes is measured in *service units*. Service units are calculated based on the CPU, SRB, I/O, and storage (MSO) service an address space consumes.

Service units are the basis for period switching within a service class that has multiple periods. The *duration* of a service class period is specified in terms of service units. When an address space or enclave running in the service class period has consumed the amount of service specified by the duration, workload management moves it to the next period. The work is managed to the goal and importance of the new period.

Because not all kinds of service are equal in every installation, you can assign additional weight to one kind of service over another. This weight is called a service coefficient.

### Changing Your Coefficient Values

Prior to z/OS 1.3, you could use the same coefficients as used in the IEAIPsxx parmlib member. Then you could directly compare RMF data, and determine your durations properly.

However, if you plan to use workload management, it is probably a good time for you to re-think your coefficients. The current defaults are inflated, given the size and processing capability of processors. Processors can consume much higher amounts of service, and as a result, service unit consumption numbers are very high. These high numbers can cause problems if they reach the point where they wrap in the SMF fields. If they wrap, you may see abnormally large transaction counts, and last period work may be restarted in the first period.

It is possible for you to make them smaller, yet still maintain the same relationship between the coefficient values. Consider changing your definitions to the following:

```
CPU    1
IOC    0.5
MSO    0
```

SRB 1

If you do decide to change the coefficients, you must re-calculate your durations and accounting procedures.

HINT: If you want to gather storage service information by service class, but don't want it affecting your durations or accounting procedures, use an MSO coefficient of 0.0001. This results in very low MSO service unit numbers, but still allows you to obtain storage service information through RMF.

Since changing the coefficients affects durations and accounting values, the defaults are meant to be consistent with settings seen in the field today. If you do not define the service coefficients, the defaults are:

CPU 10.0  
IOC 5.0  
MSO 0.0  
SRB 10.0

---

## Specifying I/O Priority Management

I/O priority queueing is used to control non-paging DASD I/O requests that are queued because the device is busy. You can optionally have the system manage I/O priorities in the sysplex based on service class goals.

The default for I/O priority management is **no**, which sets I/O priorities equal to dispatching priorities. This is identical to how I/O priorities were handled prior to OS/390 R3. If you specify **yes**, workload management sets I/O priorities in the sysplex based on goals.

WLM dynamically adjusts the I/O priority based on how well each service class is meeting its goals and whether the device can contribute to meeting the goal. The system does not micro-manage the I/O priorities, and changes a service class period's I/O priority infrequently.

When I/O priority management is on, I/O samples are used in the velocity formula. See "Velocity Formula" on page 42 for more information.

## Considerations for I/O Priority Management

If you specify I/O priority management, workload management dynamically sets I/O priorities based on goals and I/O activity, and includes the I/O information when calculating execution velocity. So you might see some changes in your velocity values.

The new DASD I/O using and DASD I/O delay samples are reported even when I/O priority management is turned off. This allows you to calculate the new velocity values to plan for velocity changes.

See Table 5 on page 85 for the recommended I/O priority management values, according to the system levels in the sysplex and the state of the velocity goal.

Table 5. Recommended Values for I/O Priority Management

Condition	Recommended Value
Pre-OS/390 R3 systems exist in the sysplex, or new velocity data is being collected.	NO
All systems in the sysplex are at the OS/390 R3 level or later, and velocity goals have been adjusted for new values.	YES

## Specifying Dynamic Alias Management

### Workload Management Considerations

As part of the Enterprise Storage Subsystem's implementation of parallel access volumes, the concept of base addresses versus alias addresses is introduced. While the base address is the actual unit address of a given volume, there can be many alias addresses assigned to a base address, and any or all of those alias addresses can be reassigned to a different base address. With dynamic alias management, WLM can automatically perform those alias address reassignments to help work meet its goals and to minimize IOS queueing.

When you specify **yes** for this value on the Service Coefficient/Service Definition Options panel, you enable dynamic alias management globally throughout the sysplex. WLM will keep track of the devices used by different workloads and broadcast this information to other systems in the sysplex. If WLM determines that a workload is not meeting its goal due to IOS queue time, then WLM attempts to find alias devices that can be moved to help that workload achieve its goal. Even if all work is meeting its goals, WLM will attempt to move aliases to the busiest devices to minimize overall queueing.

As of z/OS 1.3, in addition to automatically managing aliases, WLM ensures during system initialization that a minimum number of aliases is assigned to parallel access volumes with page datasets. The minimum number of aliases for a volume is:

$$2 * \text{number\_of\_page\_datasets} - 1$$

This allows each page dataset to have two I/Os active at a time, and ensures that paging intensive activities such as system dumping are not delayed by IOS queueing. This automatic enforcement of the minimum aliases only happens if dynamic parallel access volumes management is active for the device.

**IMPORTANT:** If you enable dynamic alias management, you must also enable I/O priority management. So you need to specify **yes** for both of these options on the panel. If I/O priority management is set to **no**, you will get only the efficiency part of dynamic alias management and not the goal-oriented management. This means that WLM will make alias moves that minimize overall IOS queueing, but these moves will not take service class goals into consideration.

### HCD Considerations

While you can globally enable or disable dynamic alias management on the WLM ISPF panel, you can also individually enable or disable dynamic alias management on a given device via HCD. You can do this by specifying **WLMPAV=YES** or **NO** in that device's HCD definition.

Note, however, that there is no consistency checking for dynamic alias management between different systems in a sysplex. If at least one system in the sysplex specifies WLMPAV=YES for a device, then dynamic alias tuning will be enabled on that device for *all* systems in the sysplex, even if other systems have specified WLMPAV=NO. It is recommended *not* to use dynamic alias management for a device unless all systems sharing that device have dynamic alias management enabled. Otherwise, WLM will be attempting to manage alias assignments without taking into account the activity from the non-participating systems.

Note, also, that you can specify WLMPAV=YES or NO on both base and alias devices. The WLMPAV settings on an alias device, however, is only meaningful when the alias device is bound to a base device that is offline, as follows:

- If the base device is *offline*, then only alias devices with WLMPAV set to YES will be reassigned to other base devices.

The WLMPAV setting on the base device itself is irrelevant when the base device is offline, for either “giving” or “receiving” aliases. (Even if WLMPAV was set to YES on the base device, it cannot have new aliases assigned to it, as it is offline.)

- For any base device that is offline to one or more systems in the sysplex and online to others, the WLMPAV keyword in HCD needs to be set to NO for the base device and its aliases. You need to statically assign the desired number of aliases for the base device via the ESS Specialist. If you try to use dynamic alias management for such a device, WLM will make unpredictable alias moves.
- If the base device is *online*, then the WLMPAV settings on the aliases are ignored, as follows:
  - If WLMPAV is set to YES on the base device, then the aliases can be reassigned regardless of their WLMPAV settings.
  - If WLMPAV is set to NO on the base device, then no aliases can be reassigned, regardless of their WLMPAV settings.

For a WLMPAV=YES base device, the aliases initially assigned to it should be allowed to default to YES. The only situation where you might want to change an alias to WLMPAV=NO is if the alias is initially assigned to a WLMPAV=NO base device. Because the base is set to NO, the aliases initially assigned to it will not be moved to other bases by WLM. Then, because the aliases are set to NO, if the base is ever varied offline, the aliases remain assigned to that base and cannot be reassigned by WLM to other bases. Certain combinations of WLMPAV settings are not recommended, as described in the table below:

Table 6. Effects of WLMPAV Settings on Base and Alias Devices

Base Device WLMPAV Setting	Alias Device WLMPAV Setting	Effects/Recommendations
YES	YES	<ul style="list-style-type: none"> <li>• <b>If base is online:</b> Base is WLM-managed. Aliases can be freely moved to and from the base device by WLM.</li> <li>• <b>If base is offline:</b> Aliases become unbound and are available to WLM to assign to other WLM-managed bases.</li> </ul>
YES	NO	<b>Not recommended.</b> If base is WLM-managed, then it is not predictable which aliases will remain bound to that base when the base goes offline. If the base device is set to YES, then you should set the aliases to YES as well. (See previous option.)
NO	YES	<b>Not recommended.</b> If the base is not WLM-managed, then you risk losing all of its aliases when the device goes offline. (See next option.)

Table 6. Effects of WLMPAV Settings on Base and Alias Devices (continued)

Base Device WLMPAV Setting	Alias Device WLMPAV Setting	Effects/Recommendations
NO	NO	<ul style="list-style-type: none"> <li>• <b>If base is online:</b> Base is not WLM-managed. The initial aliases assigned to this base remain there.</li> <li>• <b>If base is offline:</b> Aliases remain bound to the offline base device and are not available to WLM for reassignment. When the base comes back online, it retains its initial alias assignments.</li> </ul>

In order for dynamic alias management to be most effective, you should try to spread out your aliases in the initial definition. If one base device has several alias devices while other base devices have none, it will take more time for WLM to reassign the aliases appropriately. Ideally, you should have at least two aliases assigned to each base at the outset.

See *z/OS HCD Planning* for more information on HCD definitions.



## Chapter 12. Defining Special Protection Options for Critical Work

In OS/390 Release 10 and later, new options are available to help performance administrators protect critical work. Although applicable to several other subsystem types, CICS and IMS work will particularly benefit from the enhancements described in this section:

- Long-term storage protection (available on Release 10 with APAR OW43810 installed)
- Long-term CPU protection (available on Release 10 with APAR OW43855 installed)
- Exemption from transaction response time management (available on Release 10 with APAR OW43812 installed).

These three options are described below, and then illustrated by example in “Sample Scenarios” on page 91.

### Important

The use of these options limits WLM’s ability to manage the system. This may affect system performance and/or reduce the system’s overall throughput.

## Long-Term Storage Protection

When you assign long-term storage protection to critical work, WLM restricts storage donations to other work. This option can be useful for work that needs to retain storage during long periods of inactivity because it cannot afford paging delays when it becomes active again. With long-term storage protection assigned, this work will lose storage only to other work of equal or greater importance that needs the storage to meet performance goals.

You assign long-term storage protection with the “Storage Critical” option, found by scrolling right on the “Modify Rules for the Subsystem Type” panel:

```

====> _____ Modify Rules for the Subsystem Type _____ Row 1 to 2 of Command
                SCROLL <====> PAG

Subsystem Type . : CICS          Fold qualifier names?  Y (Y or N)
Description . . . CICS Transactions

Action codes:  A=After      C=Copy      M=Move      I=Insert rule
                B=Before    D=Delete row  R=Repeat    IS=Insert Sub-rule
                <==== More

Action   -----Qualifier-----   Storage   Manage Region
        Type      Name      Start   Critical   Using Goals Of
----- 1  TN      COMBL*   _____ NO         N/A
----- 2  UI      COMBLD   _____ NO         N/A
----- 2  UI      COMFTP   _____ YES        N/A
  
```

Figure 13. Specifying the “Storage Critical” Option

You can assign storage protection to all types of address spaces using classification rules for subsystem types ASCH, JES, OMVS, STC, and TSO. By specifying YES in

the “Storage Critical” field for a classification rule, you assign storage protection to *all address spaces that match that classification rule*. An address space must be in a service class that meets two requirements, however, before it can be storage-protected:

- The service class must have a single period.
- The service class must have either a velocity goal, or a response time goal of over 20 seconds.

**Note:** These 2 requirements only apply to the address spaces classified under subsystem types ASCH, JES, OMVS, STC, and TSO.

For CICS and IMS work, you can assign long-term storage protection by specifying YES in the “Storage Critical” field in the rules for specific transactions. Once you specify YES for one transaction in a CICS/IMS service class, *all CICS/IMS transactions in that service class* will be storage-protected. If a CICS or IMS region is managed as a server by WLM (managed to the response time goals of the transactions it serves) and *any* of the transaction service classes it serves is assigned storage protection, then the CICS/IMS region itself is automatically storage-protected by WLM.

As an alternative to assigning storage protection based on specific transaction service classes, you can instead choose to assign storage protection to the region in which the transactions run. You do this by adding or modifying the STC or JES classification rule that assigns the service class to the region.

---

## Long-Term CPU Protection

When you assign long-term CPU protection to critical work, you ensure that less important work will generally have a lower dispatch priority. (There are some rare exceptions, such as when other work is promoted because it is holding an enqueue for which there is contention.) This protection can be valuable for work which is extremely CPU-sensitive, such as certain CICS and IMS transactions.

Use the “Cpu Critical” option on the “Modify a Service Class” panel to assign long-term CPU protection to a specific service class:

Create a Service Class ROW 1 TO 2 OF 2

Service Class Name . . . . . APPC9 (Required)

Description . . . . . \_\_\_\_\_

Workload Name . . . . . APPC (name or ?)

Base Resource Group . . . . . \_\_\_\_\_ (name or ?)

**Cpu Critical** . . . . . **YES** (YES or NO)

Specify BASE GOAL information. Action Codes: I=Insert new period, E=Edit period, D=Delete period.

Figure 14. Specifying the “Storage Critical” Option

You can assign CPU protection to service classes handling address space-oriented work, enclave work, or CICS/IMS transactions, but the service class must have only one period, and it cannot have a discretionary goal. If a CICS or IMS region is managed as a server by WLM (managed to the response time goals of the transactions it serves) and *any* of the transaction service classes it serves is assigned CPU protection, then the CICS/IMS region itself is automatically CPU-protected by WLM.

## Exemption from Transaction Response Time Management

Use the “Manage Region Using Goals Of:” field in the “Modify Rules for the Subsystem Type” panel to declare that a specific CICS/IMS region will not be managed to the response times of the CICS/IMS transactions that it processes (note that other regions are not affected by what is in this column):

```

====> _____ Modify Rules for the Subsystem Type Row 1 to 2 of Command
                SCROLL <====> PAG

Subsystem Type . : STC          Fold qualifier names?  Y (Y or N)
Description . . . IBM-defined subsystem type
Action codes:  A=After    C=Copy    M=Move    I=Insert rule
                B=Before  D=Delete row R=Repeat  IS=Insert Sub-rule
                <==== More
Action  -----Qualifier-----  Storage  Manage Region
        Type      Name      Start  Critical  Using Goals Of
----- 1  SY      SYST1      _____  NO        TRANSACTION
----- 2  TN      CICS*      _____  YES       TRANSACTION
----- 2  TN      CICS*      _____  YES       TRANSACTION
  
```

Figure 15. Specifying the “Manage Region Using Goals Of:” Option

If you specify TRANSACTION in this field (the default), the region will be managed as a CICS/IMS transaction server by WLM. If you specify REGION in this field, the region will be managed to the performance goal of the service class assigned to that region (address space). In other words, it will *not* be managed as a CICS/IMS transaction server by WLM. Note that this option can only be used in STC and JES classification rules.

If you specify TRANSACTION, RMF will report performance information as follows:

- Response time data is reported in the WLMGL SCPER report for the service class in which those transactions are running
- Response time data is also reported in the WLMGL RCLASS report for the report class in which those transactions are running
- The service classes served by the region are reported in the WLMGL SCLASS report for the service class in which the region is running.

If you specify REGION, only the information in the SCPER report will change. Transaction response times reported by these regions will not be reported in any service class. This response time data will still be reported in the RCLASS report for the transaction report class, and the service classes served by the region will still be reported in the SCLASS report for the service class in which the region is running. In both of these cases, this information will be useful if you are migrating the CICS/IMS region to transaction response time management, and need both a transaction response time benchmark and a list of the service classes the region is serving.

## Sample Scenarios

The following scenarios illustrate different configurations and how they would benefit from these options:

## Scenario 1

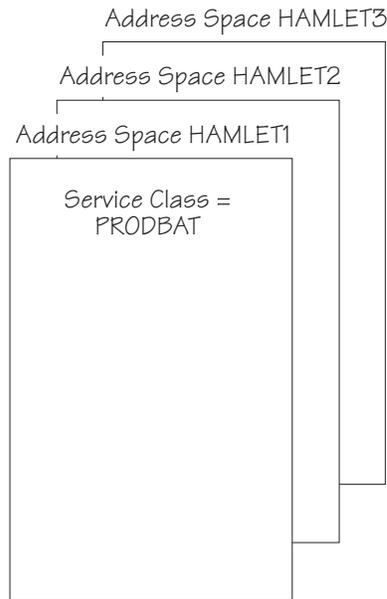


Figure 16. Scenario 1: Address Spaces

In this scenario, you wish to assign storage protection and/or CPU protection to address spaces. Suppose you have the following classification rules:

Subsystem Type . . : JES		Fold qualifier names? Y (Y or N)		
Description . . . IBM-defined subsystem type				
		-----Qualifier-----	-----Class-----	Storage
Action	Type	Name	Start	Service Report Critical
___ 1	TN	HAMLET1	___	PRODBAT ___ NO
___ 1	TN	HAMLET*	___	PRODBAT ___ NO

**Note:** Many of the “panels” shown in this section are actually composites of the information displayed after scrolling right from the “Modify Classification Rules” panels.

To assign storage protection to the address space HAMLET1, change the “Storage Critical” value to field to YES in the classification rule for HAMLET1:

Subsystem Type . . : JES		Fold qualifier names? Y (Y or N)		
Description . . . IBM-defined subsystem type				
		-----Qualifier-----	-----Class-----	Storage
Action	Type	Name	Start	Service Report Critical
___ 1	TN	HAMLET1	___	PRODBAT ___ <b>YES</b>
___ 1	TN	HAMLET*	___	PRODBAT ___ NO

To assign storage protection to the address spaces HAMLET2 and HAMLET3 (but *not* HAMLET1), change the “Storage Critical” value to field to YES in the wild card rule for HAMLET\*:

```

Subsystem Type . . : JES          Fold qualifier names?  Y  (Y or N)
Description . . . . IBM-defined subsystem type

Action   -----Qualifier-----   -----Class-----   Storage
Type    Name   Start   Service   Report   Critical
-----  ---   ----   -----   -----   -----
  1     TN     HAMLET1  PRODBAT  _____ NO
  1     TN     HAMLET*  PRODBAT  _____ YES

```

Note that in this instance, HAMLET1 will *not* be protected, as it will match the HAMLET1 rule first. To protect *all* of the address spaces, you would have to specify YES in both rules.

The default values (with no CPU protection assigned) in the service class definitions would be:

```

Service Class Name . . . . . PRODBAT (Required)
Description . . . . . Production Batch _____
Workload Name . . . . . JES (name or ?)
Base Resource Group . . . . . (name or ?)
Cpu Critical . . . . . NO (YES or NO)

```

To assign CPU protection to all of these address spaces, change the value in the "Cpu Critical" field in the service class definition for PRODBAT:

```

Service Class Name . . . . . PRODBAT (Required)
Description . . . . . Production Batch _____
Workload Name . . . . . JES (name or ?)
Base Resource Group . . . . . (name or ?)
Cpu Critical . . . . . YES (YES or NO)

```

## Scenario 2

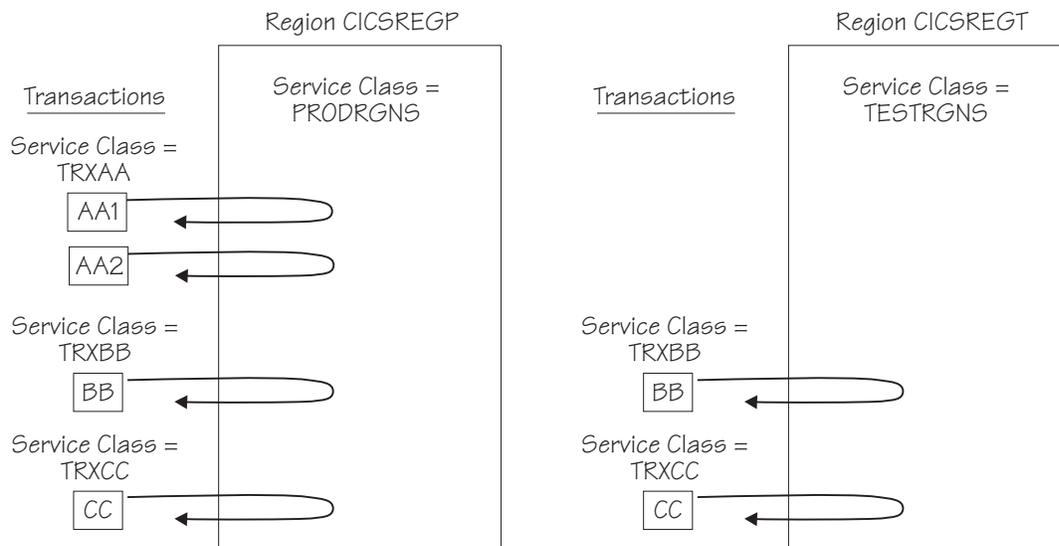


Figure 17. Scenarios 2,3,4,5: CICS/IMS Regions

Suppose you have the following CICS classification rules:

```

Subsystem Type . . : CICS          Fold qualifier names?  Y (Y or N)
Description . . . IBM-defined subsystem type

```

Action	-----Qualifier-----			-----Class-----		Storage
	Type	Name	Start	Service	Report	Critical
___ 1	TN	AA1	___	TRXAA	AA1RPT	NO
___ 1	TN	AA2	___	TRXAA	AA2RPT	NO
___ 1	TN	BB	___	TRXBB	___	NO
___ 1	TN	CC	___	TRXCC	___	NO

Suppose you have the following STC classification rules to classify the CICS regions:

```

Subsystem Type . . : STC          Fold qualifier names?  Y (Y or N)
Description . . . IBM-defined subsystem type

```

Action	-----Qualifier-----			-----Class-----		Storage	Manage Region
	Type	Name	Start	Service	Report	Critical	Using Goals Of:
___ 1	TN	CICSREGP	___	PRODRGNS	___	NO	TRANSACTION
___ 1	TN	CICSREGT	___	TESTRGNS	___	NO	TRANSACTION

Suppose the transaction service classes are defined as follows. (Only TRXAA is shown here. The definitions for TRXBB and TRXCC would also have NO specified in the "Cpu Critical" field.)

```

Service Class Name . . . . . TRXAA (Required)
Description . . . . . Transactions AA1, AA2_____
Workload Name . . . . . JES (name or ?)
Base Resource Group . . . . . _____ (name or ?)
Cpu Critical . . . . . NO (YES or NO)

```

Suppose the regions' service classes are defined as follows (PRODRGNS shown here, TESTRGNS would look the same):

```

Service Class Name . . . . . PRODRGNS (Required)
Description . . . . . Production Regions_____
Workload Name . . . . . STC (name or ?)
Base Resource Group . . . . . _____ (name or ?)
Cpu Critical . . . . . NO (YES or NO)

```

In this scenario, assume that the regions are running normal, non-conversational transactions. Response time goals are appropriate, and there is enough activity so that WLM can manage the regions as servers virtually all of the time. Transaction AA1 is very important to the business, and you wish to give it both storage and CPU protection.

In this case, protection on a transaction service class level is sufficient. This approach allows you to focus on protecting specific transactions rather than the regions that process them. The protection will be inherited by any regions in which the transactions run, as long as WLM is allowed to manage the region to the transactions' goals. ("Scenario 5" on page 97 shows what happens when WLM is *not* allowed to manage the region to the transactions' goals.)

Assign storage protection to transaction AA1 using the CICS classification rule:

```

Subsystem Type . . : CICS          Fold qualifier names?  Y  (Y or N)
Description . . . : IBM-defined subsystem type

```

Action	-----Qualifier-----			-----Class-----		Storage
	Type	Name	Start	Service	Report	Critical
___ 1	TN	AA1	___	TRXAA	AA1RPT	YES
___ 1	TN	AA2	___	TRXAA	AA2RPT	NO
___ 1	TN	BB	___	TRXBB	___	NO
___ 1	TN	CC	___	TRXCC	___	NO

Transaction service class TRXAA runs only in the CICSREGP region, and not in CICSREGT; therefore, CICSREGP will inherit the storage protection, and CICSREGT will *not* inherit the storage protection.

Assign CPU protection to the transaction service class TRXAA in the service class definition:

```

Service Class Name . . . . . TRXAA    (Required)
Description . . . . . CICS Transactions
Workload Name . . . . . CICS        (name or ?)
Base Resource Group . . . . . _____ (name or ?)
Cpu Critical . . . . . YES          (YES or NO)

```

Any region serving any TRXAA transactions, even one serving AA2 only, inherits CPU protection. As was true for storage protection, CICSREGT will *not* inherit CPU protection because it does not serve transaction service class TRXAA.

Reporting products which display data about the regions themselves will not show that storage and CPU protection was specified, but will show that they were protected while serving the transactions. (See “Reporting” on page 98.) Service class reports will show the storage and CPU protection assigned to the TRXAA transaction service class.

### Scenario 3

In this scenario, again using the CICS regions shown in Figure 17 on page 93, assume that the regions are running non-conversational transactions, but with periods of low activity during which WLM may stop managing them as servers. During this time, it is more likely that the regions’ pages will be stolen by competing workloads. In this scenario, assume that transaction BB needs storage and CPU protection.

Protection on a transaction service class level is once again useful, ensuring that the transactions will be protected wherever they run. The regions themselves should also be protected, as WLM may not manage them as servers during the low activity periods.

Assign storage protection to transaction BB using the CICS classification rules:

```

Subsystem Type . : CICS          Fold qualifier names?  Y (Y or N)
Description . . . IBM-defined subsystem type

```

Action	-----Qualifier-----			-----Class-----		Storage
	Type	Name	Start	Service	Report	Critical
___ 1	TN	AA1	___	TRXAA	AA1RPT	NO
___ 1	TN	AA2	___	TRXAA	AA2RPT	NO
___ 1	TN	BB	___	TRXBB	___	<b>YES</b>
___ 1	TN	CC	___	TRXCC	___	NO

Also, assign storage protection to the regions themselves using the STC classification rules:

```

Subsystem Type . : STC          Fold qualifier names?  Y (Y or N)
Description . . . IBM-defined subsystem type

```

Action	-----Qualifier-----			-----Class-----		Storage	Manage Region
	Type	Name	Start	Service	Report	Critical	Using Goals Of:
___ 1	TN	CICSREGP	___	PRODRGNS	___	<b>YES</b>	TRANSACTION
___ 1	TN	CICSREGT	___	TESTRGNS	___	<b>YES</b>	TRANSACTION

Assign CPU protection to the transaction service classes TRXBB (the "Cpu Critical" field in the TRXAA and TRXCC service class definitions would remain NO):

```

Service Class Name . . . . . TRXBB (Required)
Description . . . . . CICS Transactions
Workload Name . . . . . CICS (name or ?)
Base Resource Group . . . . . (name or ?)
Cpu Critical . . . . . YES (YES or NO)

```

And also to the regions themselves (PRODRGNS shown here, TESTRGNS would also specify YES in the "Cpu Critical" field):

```

Service Class Name . . . . . PRODRGNS (Required)
Description . . . . . CICS Regions
Workload Name . . . . . STC (name or ?)
Base Resource Group . . . . . (name or ?)
Cpu Critical . . . . . YES (YES or NO)

```

Note that since both CICSREGP and CICSREGT run transaction BB, *both* regions must be protected.

Reporting products which display data about the regions will show that both CPU and storage protection was specified. (See "Reporting" on page 98.) While the regions are serving transactions, protection will occur if either the regions themselves or any of their served transaction service classes are protected.

## Scenario 4

In this scenario, again using the CICS regions shown in Figure 17 on page 93, assume that the regions are running conversational transactions, and response time goals are not appropriate. By exempting the regions from management to the transaction response time goals, the regions will instead be managed according to the goal of the service class assigned to those regions. (If either storage or CPU protection is needed, that goal must be a velocity goal, since discretionary goals are

not eligible for storage or CPU protection.) In this scenario, assume that only the production region CICSREGP needs protection.

Assign storage protection to the CICSREGP region. Also, in the same panel, exempt both regions from management to the transaction response time goals:

```

Subsystem Type . . : STC          Fold qualifier names?  Y (Y or N)
Description . . . IBM-defined subsystem type

-----Qualifier-----
Action  Type   Name      Start   Service  Report  Storage  Manage Region
-----  ---   ---      ---     ---      ---     ---      Using Goals Of:
-----  ---   ---      ---     ---      ---     ---      REGION
_____  1   TN      CICSREGP  _____  PRODRGNS  _____  YES
_____  1   TN      CICSREGT  _____  TESTRGNS  _____  NO      REGION
  
```

Assign CPU protection to the PRODRGNS service class (the “Cpu Critical” field in the TESTRGNS service class definitions would remain NO):

```

Service Class Name . . . . . PRODRGNS (Required)
Description . . . . . CICS Regions
Workload Name . . . . . STC (name or ?)
Base Resource Group . . . . . (name or ?)
Cpu Critical . . . . . YES (YES or NO)
  
```

Reporting products which display data about the regions will show that CPU or storage protection was specified based on the regions’ storage protection value and the CPU protection value of the regions’ service classes. (See “Reporting” on page 98.)

## Scenario 5

This scenario is similar to “Scenario 2” on page 93, but here you’ll see what happens when WLM is not allowed to manage one of the regions to the transactions’ goals, and how this will prevent protection of a transaction. In this case, assume that it is transaction BB that you wish to give both storage and CPU protection, and assume that at the same time you have exempted region CICSREGT from management to the transaction response time goals.

You’ve assigned the storage protection to transaction BB using the CICS classification rule:

```

Subsystem Type . . : CICS          Fold qualifier names?  Y (Y or N)
Description . . . IBM-defined subsystem type

-----Qualifier-----
Action  Type   Name      Start   Service  Report  Storage
-----  ---   ---      ---     ---      ---     ---
_____  1   TN      AA1      _____  TRXAA   AA1RPT  NO
_____  1   TN      AA2      _____  TRXAA   AA2RPT  NO
_____  1   TN      BB       _____  TRXBB   _____  YES
_____  1   TN      CC       _____  TRXCC   _____  NO
  
```

You’ve also assigned CPU protection to the transaction service class TRXBB in the service class definition:

```

Service Class Name . . . . . TRXBB (Required)
Description . . . . . CICS Transactions_____
Workload Name . . . . . CICS (name or ?)
Base Resource Group . . . . . _____ (name or ?)
Cpu Critical . . . . . YES (YES or NO)

```

In the classification rule for CICSREGT, you have exempted the region from management to the transaction response time goals:

```

Subsystem Type . : STC          Fold qualifier names?  Y (Y or N)
Description . . . IBM-defined subsystem type

Action  Type  Name  Start  Service  Report  Storage  Manage Region
        1  TN   CICSREGP  ___  PRODRGNS  ___  YES  TRANSACTION
        1  TN   CICSREGT  ___  TESTRGNS  ___  YES  REGION

```

As illustrated in Figure 17 on page 93, transaction BB runs in both regions, CICSREGP and CICSREGT. *Since WLM will not manage region CICSREGT to transaction response times, it will not inherit storage or CPU protection from the TRXBB transaction service class.* Transaction BB will therefore *not* run with storage or CPU protection in region CICSREGT.

## Reporting

Because storage protection can be implicitly applied to an entire transaction service class, and because WLM may or may not be honoring a customer's storage or CPU protection assignment at any given time (for example, due to a RESET), there are six different "states" that can be reported:

- Storage protection has been explicitly assigned on a classification rule
- Storage protection has been implicitly assigned to a CICS/IMS transaction service class (because it was assigned to at least one transaction in that service class)
- Storage protection is currently being honored
- CPU protection has been explicitly assigned in a service class definition
- CPU protection is currently being honored
- Exemption from transaction response time management has been explicitly assigned on a classification rule.

These states are reported in SMF type 30 and type 79.1 records. States that apply to an entire service class are also reported in SMF 72.3 records.

## Option Summary

The following table summarizes the effects of the storage protection, CPU protection, and exemption from transaction response time management options:

*Table 7. Summary of Storage Protection, CPU Protection, and Exemption from Transaction Response Time Management Options*

When you...	WLM...
Assign CPU protection to a service class used to manage address spaces and/or enclaves.	Protects any address space or enclave managed according to the goals of that service class. Address spaces being managed as servers are managed according to the goals of the served transactions.
Assign storage protection to an ASCH, JES, OMVS, STC, or TSO address space.	Protects any address space which matches the classification rule, regardless of its server status. Address spaces currently running in multiperiod service classes or in service classes with a short response time goal (20 seconds or less) are excluded from protection.
Assign CPU or storage protection to a CICS or IMS transaction.	Protects any regions recognized as serving that CICS/IMS transaction, unless you prevent WLM from managing the regions as servers. Note that once storage protection is assigned to any transaction in a service class, then <i>all</i> transactions in the same service class become storage protected.
Prevent WLM from managing a CICS or IMS region according to the response time goals of the transactions it is running.	Does not recognize the region as a server. The region is managed using the goal of the service class assigned to the region. Transaction response time data is not reported in the service classes to which the transactions are classified, but is still reported in their report classes, if assigned.
Issue the RESET QUIESCE command.	Will no longer enforce CPU protection. All other options remain unchanged.
Issue the RESET SRVCLASS= or RESET RESUME command.	Will assign CPU protection if the target service class has the CPU protection attribute. All other options remain unchanged.



---

## Chapter 13. Defining Application Environments

An application environment is a group of application functions requested by a client that execute in server address spaces. Workload management can dynamically manage the number of server address spaces to meet the performance goals of the work making the requests. Alternatively, the server address spaces can be started and stopped manually or by automation.

Each application environment should represent a named group of server functions that require access to the same application libraries. Grouping server functions helps simplify library security, application program change control, performance management, and system operation.

For example, an application environment could be one or more DB2 stored procedures. DB2 could have an associated application environment named PAYROLL that handles specific types of stored procedure requests.

---

### Getting Started with Application Environments

The following conditions are required before an application environment can be used:

- The work manager subsystem must have implemented the workload management services that make use of application environments. Examples of IBM-supplied work managers that use application environments are:
  - DB2 (subsystem type DB2)
  - SOMobjects (subsystem type SOM)
  - Component Broker (subsystem CB)
  - Internet Connection Server, Domino Go Webserver, or IBM http Server for OS/390 (subsystem type IWEB)
  - MQSeries Workflow (subsystem type MQ)

Refer to subsystem documentation to determine if the subsystems used in your installation make use of application environments.

- One or more application environments must be defined in the workload management service definition. The subsystem reference information should provide guidance for logically grouping applications into application environments.
- The subsystem's work requests must be associated with the appropriate application environment. This step is unique for each subsystem and should be described in the subsystem reference information.

If you request through the service definition that server address spaces be automatically managed, workload management starts and stops server spaces as needed. For example, when a DB2 stored procedure request comes into the system, workload management determines whether there is a server address space to process the work, and if there is, makes the work available to the server. If there is no server address space available, workload management creates one.

Table 8 on page 102 shows the IBM-supplied subsystems that use application environments, the types of requests made by each subsystem, and where the subsystem stores the information that maps the work to application environments.

Table 8. IBM-supplied Subsystems Using Application Environments

Application Environment Subsystem Type	Request Type	Application Environment Mapping
CB	Component Broker object method requests	Server group name
DB2	Stored procedure requests	DB2 SYSIBM.SYSROUTINES catalog table
IWEB	Hyper Text Transfer Protocol (HTTP) requests	Web configuration file
MQ	MQSeries Workflow requests	MQ process definition for the WLM-managed queue (APPLICID) field
SOM	SOM client object class binding requests	Implementation repository

## Specifying Application Environments to Workload Management

To define an application environment, specify:

- The subsystem type under which the applications are running.
- The JCL procedure to start server address spaces if you wish workload management to automatically manage the number of servers.
- If a JCL procedure is specified, any required start parameters.
- Whether requests can execute in multiple server address spaces and on multiple systems.

### Defining Application Environments

#### Application Environment

Application environment name.

#### Description (optional)

Description of the application environment.

#### Subsystem Type

The subsystem type associated with the application environment, such as SOM, DB2, or IWEB.

#### Procedure Name (optional)

The name of the JCL procedure for starting the address space. The procedure name may be omitted.

#### Start Parameters (optional)

Parameters required by the JCL procedure to start the address space.

#### Limits on address spaces

Specify whether workload management can start multiple or single server address spaces for a subsystem instance.

### Application Environment (required)

One to 32 character name of the application environment. You must use this name when specifying to the subsystem how to map work to the application environments. You also use this name in operator commands when performing actions on the application environment. The name cannot begin with the letters SYS.

For guidance in mapping subsystem work to application environments, see the subsystem reference information.

If you check the authorization of application environment server address spaces through a SAF product, such as RACF, or if you plan to do so, choose an application environment name that does not exceed 27 characters. There is a restriction with respect to the maximum length of a server profile name passed to the SAF product (see also “Authorizing Application Environment Servers” on page 109).

**Description (optional)**

Up to 32 characters describing the application environment.

**Subsystem Type (required)**

Subsystem type is the one to four character name of the subsystem using application environments. This subsystem type is provided to workload management when the subsystem initializes. The types currently in use are listed in Table 8 on page 102. For a subsystem not listed, refer to the subsystem’s documentation for the required information.

**Note:** If you are using DB2 stored procedures, note that the subsystem type DB2 specified here for an application environment is used only for identifying the DB2 subsystem when it begins to use the application environment. There is no connection between this value and classification. For more information on classification, see Chapter 9, “Defining Classification Rules” on page 49.

**Procedure Name (optional)**

Procedure name is the one to eight character name of the JCL procedure that workload management uses to start a server for the application environment work requests. Refer to the appropriate subsystem documentation for sample JCL procedures to use.

To ensure that an application environment uses the same JCL procedure across the sysplex, either (1) identical procedure proclibs must be maintained across the sysplex or (2) all the procedures must be stored in a single, shared proclib.

If you specify a procedure name, “automatic” control is in effect, and workload management manages the number of servers. If you do not specify a procedure name, “manual” control is in effect, and servers must be started manually or by automation. In either case, workload management processes work requests from the application environment according to the goals defined for the work once a server address space is available.

**Start Parameters (optional)**

Start parameters are the parameters required for the JCL procedure defined in **Procedure Name**. These parameters define how workload management should start the server address spaces. Specify parameters here that you would use for starting a server address space with an MVS START command. If you specify the symbol &IWMSSNM, workload management substitutes the subsystem instance name provided to workload management when the subsystem connected to it. Refer to the subsystem reference information to determine the instance name and the appropriate parameters to specify.

**Limits on creating server address spaces for a subsystem instance (required)**

You can limit the number of servers for a subsystem instance. Reasons for limiting the number of servers might be a need to serialize or limit application environment server activity while testing, or a restriction in the subsystem itself. There are three options: ere are three options:

1. No limit

2. Single server address space per system
3. Single server address space per sysplex

The options that are valid for a subsystem depend on its scope as described in the next section. For guidance on deciding which options to use, and to find out which options are valid for subsystems not explicitly covered in the next section, see the subsystem reference information.

#### Example of an Application Environment

To enable workload management to dynamically start server address spaces to process work from a DB2 payroll application environment, do the following:

- Define the application environment

**Application Environment**

PAYROLL

**Description**

DB2 Payroll APPLENV

**Subsystem Type**

DB2

**Procedure Name**

PAYPROC

**Start Parameters**

DB2SSN=&IWMSSNM

**Limit** No limit

When payroll work arrives into the system, workload management manages system resources to meet the goals defined for the work, and dynamically starts and stops server address spaces to process the work.

## Selecting Server Limits for Application Environments

The previous section describes an option that allows you to limit the number of servers used by workload management for an application environment. The limit is applied independently for each instance of a subsystem. For IBM-supplied subsystems, you can usually use the default value supplied by the WLM ISPF application. This section defines subsystem instance as it is used by application environments and tells how to select the server limit option for IBM-supplied subsystems.

A **subsystem instance** for an application environment is defined as a unique combination of:

- Subsystem type, as specified in the service definition for an application environment, and
- Subsystem name, as defined by the work manager subsystem when it connects to workload management.

A subsystem instance using application environments has one of two different scopes, single-system or sysplex, depending on workload management services used by its subsystem type.

If the scope is single-system, all the server address spaces for the subsystem instance are created on the system where the instance connected to workload management.

If the scope is sysplex, the server spaces can be spread across the sysplex, with workload management starting at most one server on each system. The installation may choose to start additional servers on a system through the START command or automation, and these servers are equally eligible to accept application environment work as the one started automatically.

Note that this scope applies only to application environment server management. A subsystem with single-system scope for application environments, can still perform sysplex-wide functions for other purposes.

You can limit the number of servers for a subsystem instance when defining the application environment in the service definition. Reasons for limiting the number of servers might be a need to serialize or limit application environment server activity while testing, or a restriction in the subsystem itself. There are three options:

1. No limit
2. Single server address space per system
3. Single server address space per sysplex

Options 1 and 2 apply when the subsystem type supports single-system scope. Options 1 and 3 apply when the subsystem type supports sysplex scope. “No limit” for single-system scope means any number of servers may be created for the subsystem instance on the system where it connected to workload management. “No limit” for sysplex scope means servers may be created for the subsystem instance on any number of systems in the sysplex.

In OS/390 R3 and later, the IBM-supplied subsystems using application environments, their scopes, and valid server limit options are as follows:

*Table 9. Application Environment Server Characteristics*

Subsystem Type	Scope	Valid server limit options
CB	system	<p><b>1. No limit</b></p> <p>Use this if you want no restrictions on the number of servers. This is the default for the Component Broker.</p> <p><b>2. Single server address space per system</b></p> <p>Use this if you are testing and want to temporarily limit the number of servers. The Component Broker itself does not have a requirement for limiting servers.</p>
DB2	system	<p><b>1. No limit</b></p> <p>Use this if you want no restrictions on the number of servers. This is the default for DB2.</p> <p><b>2. Single server address space per system</b></p> <p>Use this if the DB2 stored procedure cannot execute concurrently in multiple address spaces. This option should be used if a stored procedure is to be run in “debug” mode and writes to a trace.</p>

Table 9. Application Environment Server Characteristics (continued)

Subsystem Type	Scope	Valid server limit options
IWEB	system	<p><b>1. No limit</b></p> <p>Use this if you want no restrictions on the number of servers. This is the default for Internet Connection Server, Domino Go Webserver, or IBM http Server for OS/390.</p> <p><b>2. Single server address space per system</b></p> <p>Use this if you are testing and want to temporarily limit the number of servers. The Internet Connection Server, Domino Go Webserver, or IBM http Server for OS/390 does not have a requirement for limiting servers.</p>
MQ	system	<p><b>1. No limit</b></p> <p>Use this if you want no restrictions on the number of servers. This is the default for MQSeries Workflow.</p> <p><b>2. Single server address space per system</b></p> <p>Use this if you are testing and want to temporarily limit the number of servers. MQSeries Workflow itself does not have a requirement for limiting servers.</p>
SOM	sysplex	<p><b>3. Single server address space per sysplex</b></p> <p>As of OS/390 R3 and later, SOMobjects requires this limit on the number of servers for an application environment. This option is enforced by the WLM application.</p>

For guidance on deciding which option to use, or to find out what options are valid for other subsystems that use application environments, refer to the subsystem reference information.

## Using Application Environments

Application environments can be manually controlled by the installation or automatically controlled by workload management. All of the applications in an application environment are supported by a single JCL startup procedure. Defining the name of this startup JCL procedure to workload management indicates that workload management should control the server address spaces. This is called “automatic” control. If you omit the name of the JCL procedure in the application environment definition, then “manual” control is in effect.

Under “manual” control, the installation must create and delete, as needed, the server address spaces for each application environment. Note that the VARY WLM,APPLENV command can be used to terminate manually started server address spaces (through the quiesce or refresh options), but it will not restart them. For more information on the VARY WLM,APPLENV command, see “Using Operator Commands for Application Environments” on page 107.

Under “automatic” control, workload management creates server address spaces as started tasks using the JCL procedure specified in the application environment definitions. The startup parameters may be contained in either the JCL procedure defined for each application environment or in the application environment definition. When the server address spaces are no longer needed, workload management deletes them.

Under “automatic” control, the quantity of server address spaces is totally controlled by workload management. If an operator or automation starts or cancels the server address spaces under “automatic” control, workload management will:

- Use servers not started by workload management as if they were started by workload management
- Terminate servers not started by workload management if they are not needed.
- Replace a server address space that was unexpectedly cancelled.

**Note:** You should use the VARY WLM,APPLENV command to manage application environment servers rather than the CANCEL command. If there are more than five server cancellations in 10 minutes, workload management stops creating new servers for the application environment. For more information on “stop” conditions, see “Handling Error Conditions in Application Environments” on page 109.

---

## Managing Application Environments

Once an application environment is defined, and there are server address spaces in use by the subsystem, you can use operator commands to manage the application environment. There are options on the VARY WLM,APPLENV command that allow you to quiesce, resume, or refresh application environments. These functions allow you, for example, to make changes to the JCL procedure, start parameters, or application libraries, and ensure that new work requests run with the modified information. The resume function also allows you to recover from error conditions that have caused workload management to stop an application environment.

An action taken for an application environment is saved in the WLM couple data set and is not discarded across an IPL. For example, if a quiesce action is in effect and the system is IPLed, the quiesce action remains in effect. You can query the current state of an application environment using the DISPLAY WLM,APPLENV command. The scope of both the VARY and DISPLAY commands for application environments is sysplex-wide, that is, they affect the application environment on all systems in the sysplex, regardless of the scope of the subsystem using the application environment. The sysplex scope of the command ensures that an application environment remains consistent across the sysplex, especially where there are shared resources.

This section first introduces the commands that can be used to perform actions on an application environment. Then it describes activities that make use of the commands and describes other conditions that affect the state of an application environment.

### Using Operator Commands for Application Environments

An application environment initially enters the AVAILABLE state when the service policy that contains its definition is activated. AVAILABLE means the application environment is available for use, and servers are allowed to be started for it. There are three options on the VARY command that you can use to change the state of an application environment after it has been made available:

- VARY WLM,APPLENV=xxxx,QUIESCE

The quiesce option causes workload management to request the termination of server address spaces for the application environment upon completion of any active requests. Additional work requests are not handled by the servers, although work requests can continue to be queued, waiting for a server. If you do not want work queued, use subsystem functions to stop the queueing.

You can issue a quiesce action for an application environment that is in the AVAILABLE state. When a quiesce action is issued for an application environment, it first enters the QUIESCING state until all servers have been requested to terminate. It then enters the QUIESCED state.

- VARY WLM,APPLENV=xxxx,RESUME

The resume option restarts an application environment that was previously quiesced and is in the QUIESCED state. It indicates to workload management that server address spaces can once again be started for this application environment. The new servers process any queued requests and all new requests.

When a resume action is issued for an application environment, it first enters the RESUMING state until all systems in the sysplex have accepted the action. It then enters the AVAILABLE state.

- VARY WLM,APPLENV=xxxx,REFRESH

The refresh option requests the termination of existing server address spaces and starts new ones in their place. Existing servers finish their current work requests and end. The new servers process any queued requests and all new requests.

You can issue a refresh action for an application environment that is in the AVAILABLE state. When a refresh action is issued for an application environment, it first enters the REFRESHING state until all servers have been requested to terminate. It then enters the AVAILABLE state.

## Making Changes to the Application Environment Servers

The command options described above are intended to allow changes to application environments without having to shut down the application itself. Use the quiesce function when you want to do one of the following:

- Perform maintenance on application program libraries statically allocated to server address spaces.
- Update the JCL procedure for an application environment.

When you are ready to put the changes into effect, quiesce the application environment, make the changes to the libraries or service definition as needed, then use the resume function to start new servers with the changed information.

You can also use the quiesce function to suspend execution after repeated application failures or errors. After the errors are corrected, you can resume the application environment.

You may have an application environment where the servers keep application program executable modules in a private cache. If you update the application program, you need to ensure that all copies of the changed modules are replaced wherever they are cached. Use the refresh function to do this.

## Changing the Definition of an Application Environment

Workload management initiates a refresh when one of the following changes are made to the application environment definition and activated:

- The JCL procedure name is changed.
- The application environment is switched to “automatic” control, that is, the JCL procedure name was previously left blank, but now one is provided.
- The server start parameters are changed.
- The limit on server address spaces is changed; for example, changed from “no limit” to “single address space per system”.

If an application environment is deleted from the service definition, it enters the DELETING state. After workload management requests the termination of all associated servers, the application environment is no longer displayed at all by the DISPLAY WLM,APPLENV command.

## Handling Error Conditions in Application Environments

Workload management stops the creation of new server address spaces when one of the following conditions exist:

- JCL errors in the procedure associated with the application environment.
- Coding errors in the server code which cause five unexpected terminations of server address spaces within ten minutes.
- Five operator cancellations of server address spaces within ten minutes.
- Failure of the server address space to connect to workload management due to an invalid invocation environment or invalid parameters.

The application environment first enters the STOPPING state, then the STOPPED state after all systems in the sysplex have accepted the action. In STOPPED state, no new servers are created. Any existing server address spaces continue to process work, and workload management is able to accept new work. If there are no existing servers, then workload management rejects any new work requests.

In STOPPED state, you can make changes to libraries, change the procedure, or make any other changes needed to repair the condition that caused workload management to stop the application environment. When the problem is resolved, use the resume function to allow workload management to start new servers. The application environment enters the RESUMING state, then the AVAILABLE state after all systems in the sysplex have accepted the action.

**Note:** If you want to ensure all servers are restarted after a STOPPED state, especially after the JCL procedure or libraries have been modified, you should issue a quiesce function prior to the resume. This ensures there are no servers remaining active that are using back-level information.

---

## Authorizing Application Environment Servers

Because the server address spaces started on behalf of an application environment can run in problem program state, workload management enables you to check the validity of a server through an SAF product such as RACF. When the server is being created, workload management makes an SAF call using a new SERVER class to check whether the server is valid for the application environment. If you do not have the SERVER class defined to your SAF product, workload management allows the server address space to be started.

You can use the SERVER and STARTED classes with a SAF product to restrict access to application environment servers. For example, if you are using DB2, SOM or IWEB servers with application environments, you first associate a userid with the MVS procedure name being used to start the server. This is done using the STARTED resource class or by changing ICHRIN03 (started procedures table). Then you use the SERVER resource class to authorize this userid, and possibly others, to become a server for DB2 stored procedures, SOM method requests, or Internet Connection Server web requests.

## — Example for Restricting Access to Application Environment Servers —

In this example, the installation has the following situation:

- MVS JCL procedures for DB2 stored procedure servers: PAY1, PAY2, PER1, PER2

These are the JCL procedures that workload management uses to start the DB2 servers that handle stored procedure calls.

- DB2 subsystem names: DB2A and DB2B

These are the subsystem names used when the DB2 subsystem connects to workload management.

1. Activate STARTED and SERVER classes (if not already done):

```
SETR CLASSACT(STARTED) RACLIST(STARTED) GENERIC(STARTED)
SETR CLASSACT(SERVER) RACLIST(SERVER) GENERIC(SERVER)
```

2. Establish an arbitrary USERID to use in a subsequent RDEFINE command to tie an MVS procedure name to a server.

```
ADDUSER DB2SERV NOPASSWORD
```

The NOPASSWORD keyword here is important, it makes DB2SERV a protected user.

3. Associate the userid with the started task name.

```
RDEFINE STARTED PAY*.* STDATA(USER(DB2SERV) GROUP(SYS1))
RDEFINE STARTED PER*.* STDATA(USER(DB2SERV) GROUP(SYS1))
```

4. Define server profiles in the form:

```
subsys_type.subsys_name.applenv
```

where,

### **subsys\_type**

is the subsystem type, as specified in the service definition

### **subsys\_name**

is the instance name of the subsystem associated with this server. Refer to subsystem reference information for how to determine the subsystem name. The subsystem uses this name when establishing itself as the work manager for application environment server requests.

### **applenv**

is the application environment name, as specified in the service definition

```
RDEFINE SERVER DB2.DB2A.* UACC(NONE)
RDEFINE SERVER DB2.DB2B.* UACC(NONE)
```

**Note:** The maximum length of a server profile name passed to a SAF product is restricted to 41 characters. WLM cannot start server address spaces for application environments that do not follow this restriction. If you ensure that the applenv is at most 27 characters long, it is guaranteed that the maximum server profile name length does not exceed 41 characters.

5. Permit the userid to the servers. This completes the association between the MVS procedure names and the servers:

```
PERMIT DB2.DB2A.* CLASS(SERVER) ID(DB2SERV) ACCESS(READ)
PERMIT DB2.DB2B.* CLASS(SERVER) ID(DB2SERV) ACCESS(READ)
```

- Refresh the classes to refresh the RACF data base and make these changes go into effect:

```
SETR RACLIST(STARTED) REFRESH  
SETR RACLIST(SERVER) REFRESH
```



---

## Chapter 14. Defining Scheduling Environments

A scheduling environment is a list of resource names along with their required states. It allows you to manage the scheduling of work in an asymmetric sysplex where the systems differ in installed applications, or installed hardware facilities. If an MVS image satisfies all of the requirements in the scheduling environment associated with a given unit of work, then that unit of work can be assigned to that MVS image. If any of the resource requirements are *not* satisfied, then that unit of work *cannot* be assigned to that MVS image.

Scheduling environments and resource names reside in the service definition and apply across the entire sysplex. They are sysplex-oriented. Resource states have a different setting in each system in the sysplex and are, therefore, system-oriented.

Each element in a scheduling environment consists of the name of a resource and a required state of either **ON** or **OFF**, as follows:

- If the required state is **ON**, then the resource state must be set to ON on an MVS image for the requirement to be satisfied.
- If the required state is **OFF**, then the resource state must be set to OFF on an MVS image for the requirement to be satisfied.

In theory, each resource name represents the potential availability of a resource on an MVS system. That resource can be an actual physical entity such as a data base or a peripheral device, or it can be an intangible quality such as a certain time of day or a certain day of the week. The resource names are abstract, and have no inherent meaning.

For instance, you could define a resource name to be XXXX with a required state of ON. If on system SYS1 the corresponding XXXX resource state is set to ON, then the requirement is satisfied. WLM does not care what "XXXX" means, or whether the ON setting really does signify the existence of some real resource. (You could use XXXX as nothing more than an arbitrary toggle switch, setting it ON for whatever reason you wish.) As long as the settings match, the requirement is satisfied.

The rest of this chapter shows you how to define the scheduling environments, the resource names, and their required states. It also shows you how to set the resource states on each individual MVS system, and how to associate a scheduling environment name with incoming work.

---

### Getting Started with Scheduling Environments

**Note:** If you define scheduling environments or resources, the service definition functionality level will change. This will have migration implications if any pre-OS/390 R4 systems use the service definition. See Chapter 16, "Workload Management Migration" on page 131 for further information.

The following steps are required to use scheduling environments:

- You must define one or more scheduling environments, and all of the resource names and required states that are listed in those scheduling environments, in

the workload management service definition. See “Specifying Scheduling Environments to Workload Management”.

- For every system in the sysplex on which you want the resource settings to satisfy either ON or OFF requirements, you must set the individual resource states to either ON or OFF, as appropriate. There is also a third setting, RESET, that satisfies neither an ON nor OFF requirement. See “Managing Resource States” on page 117 for more information on the RESET state.
- For each unit of work with resource state requirements that is submitted for execution, you must specify the name of the scheduling environment that should be used to determine which systems can execute that work. See “Associating Scheduling Environments with Incoming Work” on page 121.

---

## Specifying Scheduling Environments to Workload Management

To define a scheduling environment, you need to specify the following information:

### Scheduling Environment Name (required)

One to 16 character name of the scheduling environment.

- You can have up to 999 unique scheduling environments defined in a service definition.
- Alphanumerics and the special characters @, \$, # and \_ are allowed.
- Underscores (\_) must be imbedded (for example, PLEX\_D01 is valid, but PLEX\_ is not).
- Names beginning with SYS\_ are reserved for system use.

### Description (optional)

Up to 32 characters describing the scheduling environment.

Once you have defined a scheduling environment, you can start selecting its resource names and required states, as follows:

### Resource Name (required)

One to 16 character name of the resource. There can be more than one resource name listed in a scheduling environment.

- You can have up to 999 unique resource names defined in a service definition.
- Alphanumerics and the special characters @, \$, # and \_ are allowed.
- Underscores (\_) must be imbedded (for example, PLEX\_D01 is valid, but PLEX\_ is not).
- Names beginning with SYS\_ are reserved for system use.

### Resource Description (optional)

Up to 32 characters describing each resource.

When you select a resource name to become part of the scheduling environment, you also need to specify a required state:

### Required State (required)

For each resource name in a scheduling environment, you must specify a required state of either ON or OFF:

- **ON** specifies that the resource name must be set to ON on a given system for the work associated with this scheduling environment to be assigned to that system.

- **OFF** specifies that the resource name must be set to OFF on a given system for the work associated with this scheduling environment to be assigned to that system.

### Scheduling Environment Example 1

To define a scheduling environment called DB2LATE that contains the following requirements:

- The “DB2A” resource must be set to ON. (In this example, we’ll say that DB2A has been defined to represent the existence of the DB2 subsystem.)
- The “PRIMETIME” resource must be set to OFF. (In this example, we’ll say that PRIMETIME has been defined to be ON during the normal weekday business hours, and OFF for all other times.)

You would define the following scheduling environment:

#### Scheduling Environment

DB2LATE

#### Description

Offshift DB2 Processing

Resource Name	Required State	Resource Description
DB2A	ON	DB2 Subsystem
PRIMETIME	OFF	Peak Business Hours

## Scheduling Environment Example 2

An installation wishes to run an IEEE floating point application program. For testing purposes, the application's performance using OS/390 Release 6 simulation is adequate, so there is no specific need for IEEE floating point hardware. For production use of the program, however, both OS/390 Release 6 and IEEE floating point hardware are required. Two scheduling environments are created, IEEE\_FP\_FAST and IEEE\_FP\_SIM, with the following requirements:

- **IEEE\_FP\_FAST:**
  - The "OS390R6" resource must be set to ON. (In this example, OS390R6 has been defined to represent the installation of OS/390 Release 6 on the system.)
  - The "IEEE\_FP" resource must be set to ON. (In this example, IEEE has been defined to represent the presence of IEEE floating point hardware in the system.)
- **IEEE\_FP\_SIM:**
  - The "OS390R6" resource must be set to ON.
  - The "IEEE\_FP" resource can be set to either ON or OFF. Because there is no required state for this resource name, it is *not* included in the scheduling environment definition, as shown below.

The installation would define the following two scheduling environments:

### Scheduling Environment

IEEE\_FP\_FAST

#### Description

IEEE Floating Point Production

Resource Name	Required State	Resource Description
OS390R6	ON	OS/390 Release 6
IEEE_FP	ON	IEEE Floating Point Hardware

### Scheduling Environment

IEEE\_FP\_SIM

#### Description

IEEE Floating Point Simulation

Resource Name	Required State	Resource Description
OS390R6	ON	OS/390 Release 6

**Null scheduling environments:** If you no longer need to restrict where work executes in a sysplex, you can remove all the resource state requirements from a scheduling environment. A null or empty scheduling environment always allows work to be scheduled; that is, any system in the sysplex is satisfactory for work associated with a null scheduling environment. This is a migration aid when you initially have resources that exist on only some of the systems in a sysplex, but later make the resources available to every system. It saves the effort of having to remove the scheduling environment specification from all the incoming work.

Refer to “Working with Scheduling Environments” on page 196 to see how to use the WLM ISPF application to create and modify scheduling environments.

---

## Managing Resource States

For every resource name that is referenced by a scheduling environment, a corresponding resource state must be set on each system in the sysplex. The resource state can be:

- **ON**, which will satisfy a resource state requirement of ON.
- **OFF**, which will satisfy a resource state requirement of OFF.
- **RESET**, which will not satisfy any resource state requirement. Resources are put into the RESET state when:
  - A system is IPLed.
  - A policy is activated that defines a resource name that did not exist in the previously active policy.

These resource states can be manipulated in three ways:

- The operator command:

```
F WLM,RESOURCE=resource_name,setting
```

where *setting* can be **ON**, **OFF**, or **RESET**.

For example, to set DB2A to ON on system SYS1, here is the command you would enter *on system SYS1*, along with the response you would receive:

```
F WLM,RESOURCE=DB2A,ON
IWM039I  RESOURCE DB2A IS NOW IN THE ON STATE
```

- The equivalent WLM application programming interface IWMSESET. See *z/OS MVS Programming: Workload Management Services* for more information on using this interface.
- Using SDSF, you can change resource state settings directly on the panel displaying the current states.

**Note:** Do not attempt to issue the F WLM,RESOURCE= command from the COMMNDxx parmlib member, as this member is processed too early during system initialization. If you want resource states to be set on every system IPL, this needs to be done through an automation product such as System Automation for OS/390 as soon as that automation product comes up during system initialization.

It is expected that, in most cases, the mechanics of managing resources states will be handled by installation-provided automation, as opposed to having a human operator issue a modify command every time a resource state is changed. Two examples of how automation could manage resource settings:

- By listening for messages from a subsystem that indicate that that subsystem has completed its initialization and is ready to accept work. The automation script could issue the appropriate F WLM,RESOURCE=subsystem,ON command on that system. When messages are issued indicating that the subsystem is about to terminate, the script could issue the appropriate F WLM,RESOURCE=subsystem,OFF command on that system.

- For time-related resource settings, a simple script can turn settings ON and OFF at certain times of the day (like at the beginning and end of peak business hours).

See *z/OS MVS Programming: Workload Management Services* for more information on automation of resource states, WLM services, and coordination with other job scheduling programs.

When you modify the resource state settings on a given system, you do so *on that system only*. If you modify the DB2A resource state on system SYS1, it has no effect on the DB2A setting on SYS2. If you wish to modify the settings on both systems, you would have to explicitly direct the commands to each individual system.

When *all* of the resource state settings on a particular system match the resource names and required states defined in a particular scheduling environment, only then is that system eligible to receive work associated with the scheduling environment.

### Resource States Example

Using the DB2LATE scheduling environment defined in the first example, here's how the resource states might be set, and how that would affect the eligibility of work scheduled with DB2LATE to run on each system in the sysplex.

1. The resource names DB2A and PRIMETIME have just been defined (and not set on the individual systems yet) or the systems have just IPLed:

Resource State	SYS1 Settings	SYS2 Settings
DB2A	RESET	RESET
PRIMETIME	RESET	RESET

2. Because of the existence of DB2 on SYS1 only, the DB2A resource state is modified to ON on that system and OFF on SYS2. Also, automation has been set up to modify the PRIMETIME setting according to the time of day. At the moment, it is 10:00 a.m. on a Monday morning, so PRIMETIME is set to ON on both systems:

Resource State	SYS1 Settings	SYS2 Settings
DB2A	ON	OFF
PRIMETIME	ON	ON

At the moment, the DB2LATE scheduling environment is not satisfied by the resource state settings of either system. Therefore, any work submitted that is associated with the DB2LATE scheduling environment cannot yet be executed on either system.

3. At 5:00 p.m., automation modifies PRIMETIME to OFF on both systems:

Resource State	SYS1 Settings	SYS2 Settings
DB2A	ON	OFF
PRIMETIME	OFF	OFF

SYS1 now finally has all of the correct resource state settings to satisfy the DB2LATE scheduling environment. The work submitted associated with the DB2LATE scheduling environment can now be executed on SYS1.

In the previous example, work associated with the DB2LATE scheduling environment was assigned to SYS1, because that system satisfied all of the DB2LATE requirements. In the case of a single-system sysplex, this ability to choose the "right" system is no longer applicable. But scheduling environments may still be useful. Consider if, in the previous example, there was only a SYS1 in the sysplex. Only when the DB2A resource setting was set to ON (signalling that DB2 was up and running) and the PRIMETIME resource setting was set to OFF (signalling that the peak business hours were over) would DB2LATE work be processed on SYS1. The DB2LATE scheduling environment would act as a "ready" flag, holding the work until all of the requirements were met.

Figure 18 on page 120 summarizes the relationship between scheduling environments and the resource state settings on several sample systems.

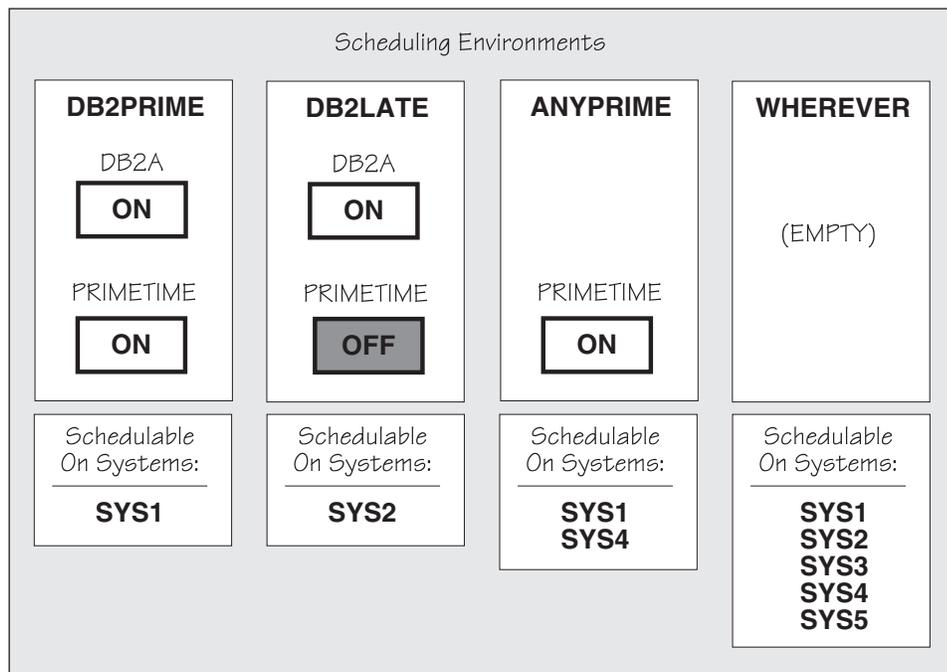
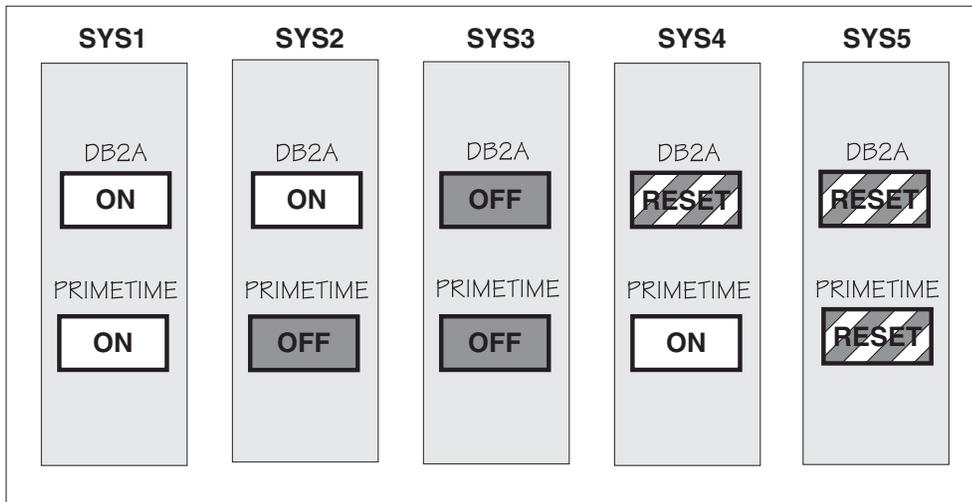


Figure 18. Sample Systems and Scheduling Environments

Note that in this figure:

- Work associated with the **DB2PRIME** scheduling environment is schedulable only on SYS1, as only that system satisfies both of the requirements.
- Similarly, work associated with the **DB2LATE** scheduling environment is schedulable only on SYS2, as only that system satisfies both of the requirements.
- Work associated with the **ANYPRIME** scheduling environment is schedulable on either SYS1 or SYS4, as both of those systems satisfy the sole requirement (PRIMETIME must be ON). This scheduling environment does not care about the DB2A setting. Therefore the RESET state for DB2A on SYS4 is irrelevant.
- Work associated with the **WHEREVER** scheduling environment is schedulable on any system in the sysplex, because it is empty (it has *no* requirements at all).

---

## Associating Scheduling Environments with Incoming Work

Having defined scheduling environments to WLM, and having set the resource states on the individual systems, all you need now is a way to associate the scheduling environment name with actual work.

The **SCHENV=** parameter on the JES2 or JES3 JCL JOB statement associates the job with the scheduling environment as shown in the following JCL example:

```
//SHAMIL JOB (C003,6363),'Steve Hamilton',  
//          MSGLEVEL=(1,1),  
//          REGION=4096K,  
//          CLASS=A,  
//          SCHENV=DB2LATE,  
//          MSGCLASS=0  
//          ...
```

This specification associates this batch job with the DB2LATE scheduling environment. It can be coded by the end user, or automatically supplied by installation-provided exits.

If the scheduling environment name specified is not defined in the active WLM policy, the job will fail with a JCL error during conversion, accompanied by an appropriate error message.

Existing JES2 or JES3 exits can be used to change the scheduling environment name associated with batch jobs. This can be done during JCL conversion. These exits can also be used to dynamically generate scheduling environment associations as work is submitted. This could be useful in migrating from another scheduling mechanism to scheduling environments. See *z/OS JES2 Installation Exits* or *z/OS JES3 Customization* for more information.

---

## Displaying Information about Scheduling Environments and Resource States

Once you have defined scheduling environments, you can issue several different operator commands, both from MVS and from JES2 or JES3 to display information about the scheduling environments and about the resource states.

### MVS Operator Commands

To display sysplex-level information about a scheduling environment, you can issue the following command from an MVS console:

```
D WLM,SCHENV=scheduling_environment
```

For example, to display information about the DB2LATE scheduling environment used above, here is the command you would issue and the response you would receive:

```
D WLM,SCHENV=DB2LATE  
  
IWM036I 12.21.05 WLM DISPLAY 181  
SCHEDULING ENVIRONMENT: DB2LATE  
DESCRIPTION:           Offshift DB2 Processing  
AVAILABLE ON SYSTEMS:  SYS1 SYS3
```

The AVAILABLE ON SYSTEMS field shows that at the time this command was issued, only systems SYS1 and SYS3 satisfied the requirements of the DB2LATE scheduling environment.

To display information about *all* scheduling environments in a sysplex, issue the command with an asterisk (\*) in the *scheduling\_environment* field, as in this example:

```

D WLM,SCHENV=*

IWM036I 12.21.05 WLM DISPLAY 181
SCHEDULING ENVIRONMENT: DB2LATE
DESCRIPTION:           Offshift DB2 Processing
AVAILABLE ON SYSTEMS:  SYS1  SYS3

SCHEDULING ENVIRONMENT: IMSPRIME
DESCRIPTION:           Primetime IMS Processing
NOT AVAILABLE ON ANY SYSTEM

```

In this example, NOT AVAILABLE ON ANY SYSTEM is shown for the IMSPRIME scheduling environment, meaning that *no* systems in the sysplex currently satisfy the IMSPRIME requirements.

To display system-level information about a scheduling environment, use the **SYSTEM=system\_name** parameter. You will see all of the resource names included in that scheduling environment, along with their required and current states. Requirements that are *not* satisfied are marked with an asterisk. So assuming that the DB2LATE scheduling environment is satisfied on SYS1 but not on SYS2, here is what the commands and responses might look like if you wanted the information for both systems:

```

D WLM,SCHENV=DB2LATE,SYSTEM=SYS1

IWM037I 12.21.05 WLM DISPLAY 181
SCHEDULING ENVIRONMENT: DB2LATE
DESCRIPTION:           Offshift DB2 Processing
SYSTEM:                SYS1
STATUS:                AVAILABLE

RESOURCE NAME   REQUIRED   CURRENT
STATE           STATE     STATE
DB2A            ON       ON
PRIMETIME      OFF      OFF

D WLM,SCHENV=DB2LATE,SYSTEM=SYS2

IWM037I 12.21.05 WLM DISPLAY 181
SCHEDULING ENVIRONMENT: DB2LATE
DESCRIPTION:           Offshift DB2 Processing
SYSTEM:                SYS2
STATUS:                NOT AVAILABLE

RESOURCE NAME   REQUIRED   CURRENT
STATE           STATE     STATE
*DB2A          ON       OFF
PRIMETIME      OFF      OFF

```

To display information about a specific resource state on a specific system, use the command:

```

D WLM,RESOURCE=resource_name,SYSTEM=system

```

To display information about the DB2A resource state on system SYS1, for example, here is the command and response:

```
D WLM,RESOURCE=DB2A,SYSTEM=SYS1
```

```
IWM038I 12.21.05 WLM DISPLAY 181
RESOURCE:      DB2A
DESCRIPTION:   DB2 Subsystem
SYSTEM STATE
SYS1          ON
```

To display information about *all* resource settings, use the asterisk in the *resource\_name* field. Also, to display information on *all* systems in the sysplex, use the **SYSTEMS** keyword in place of the **SYSTEM=system\_name** parameter. So to display information about all the resource states on all systems in the sysplex, here is a typical command and response:

```
D WLM,RESOURCE=*,SYSTEMS
```

```
IWM038I 12.21.05 WLM DISPLAY 181
RESOURCE:      DB2A
DESCRIPTION:   DB2 Subsystem
SYSTEM STATE   SYSTEM STATE   SYSTEM STATE
SYS1          ON    SYS2          OFF    SYS3          RESET

RESOURCE:      PRIMETIME
DESCRIPTION:   Peak Business Hours
SYSTEM STATE   SYSTEM STATE   SYSTEM STATE
SYS1          OFF    SYS2          OFF    SYS3          OFF
```

## JES2/JES3 Operator Commands

The JES2 display command **\$D JOBQ** can be used to list those queued jobs associated with a scheduling environment, as in this example:

```
$djobq,schenv=DB2LATE
JOB00007 $HASP608 JOB(SHAMIL)
$HASP608 JOB(SHAMIL) STATUS=(AWAITING EXECUTION),CLASS=S,
$HASP608 PRIORITY=1,SYSAFF=(ANY),HOLD=(NONE),
$HASP608 CMDAUTH=(LOCAL),OFFS=(),SECLABEL=,
$HASP608 USERID=DEALLOC,SPOOL=(VOLUMES=(SPOOL1),
$HASP608 TGS=2,PERCENT=0.3809),ARM_ELEMENT=NO,
$HASP608 SRVCLASS=HOTBATCH,SCHENV=DB2LATE
```

For JES3, use the **\*INQUIRY,Q** command, as in this example:

```
*I,Q,SCHENV=IMSPROD
IAT8674 JOB JOB123 (JOB32787) P=02 CL=Z MAIN(ALLOCATE)
IAT8674 JOB JOBABC (JOB32790) P=02 CL=Z MAIN(ALLOCATE)
IAT8674 JOB JOBDEF (JOB32791) P=02 CL=Z MAIN(ALLOCATE)
IAT8674 JOB JOBGHI (JOB32800) P=02 CL=Z MAIN(ALLOCATE)
IAT8674 JOB JOBJKL (JOB32987) P=02 CL=Z MAIN(ALLOCATE)
IAT8674 JOB JOBMNO (JOB33101) P=02 CL=Z MAIN(ALLOCATE)
```

See *z/OS JES2 Commands* or *z/OS JES3 Commands* for more information on using JES2 or JES3 operator commands.

## SDSF Commands

SDSF can display scheduling environment information, resource information, and allows modification of resource states.



---

## Chapter 15. Worksheets for Defining a Service Definition

The following worksheets help you define a service definition. These worksheets, in order of presentation, are to help you define your:

1. Workloads and service classes
2. Resource groups
3. Classification rules for each of your subsystem types
4. Application environments
5. Scheduling environments.

You should use the worksheets to set up your service definition before you enter any information into the ISPF administrative application.

# Base

Policy \_\_\_\_\_

Policy \_\_\_\_\_  
(Over-rides)

Workload  
Description \_\_\_\_\_  
\_\_\_\_\_

Service Class \_\_\_\_\_  
Description \_\_\_\_\_

Resource Group \_\_\_\_\_

Period 1	Goal	_____	_____
	Importance	_____	_____
	Duration	_____	_____

Period 2	Goal	_____	_____
	Importance	_____	_____
	Duration	_____	_____

Period 3	Goal	_____	_____
	Importance	_____	_____

Service Class \_\_\_\_\_  
Description \_\_\_\_\_

Resource Group \_\_\_\_\_

Period 1	Goal	_____	_____
	Importance	_____	_____

Service Class \_\_\_\_\_  
Description \_\_\_\_\_

Resource Group \_\_\_\_\_

Period 1	Goal	_____	_____
	Importance	_____	_____

# Base

Policy \_\_\_\_\_

Policy \_\_\_\_\_  
(Over-rides)

Resource Group \_\_\_\_\_  
Description \_\_\_\_\_

Min Capacity \_\_\_\_\_  
Max Capacity \_\_\_\_\_

\_\_\_\_\_  
\_\_\_\_\_

Service Classes \_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_

\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_

Resource Group \_\_\_\_\_  
Description \_\_\_\_\_

Min Capacity \_\_\_\_\_  
Max Capacity \_\_\_\_\_

\_\_\_\_\_  
\_\_\_\_\_

Service Classes \_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_

\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_

# Classification Rules for \_\_\_\_\_ Subsystem

Service Class

Report Class

Default: \_\_\_\_\_

\_\_\_\_\_

Qualifier Type

Qualifier Name

\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_

\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_

\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_

\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_  
\_\_\_\_\_

Application Environment \_\_\_\_\_

Description \_\_\_\_\_

Subsystem Type \_\_\_\_\_

Procedure Name \_\_\_\_\_

Start Parameters \_\_\_\_\_

Limit on starting server address spaces for a subsystem instance (select one)

1. No limit
2. One address space per system
3. One address space per sysplex

Scheduling Environment \_\_\_\_\_

Description \_\_\_\_\_

Resource Name	Required State	Description
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____

---

## Chapter 16. Workload Management Migration

This chapter covers two different workload management migration scenarios, depending on whether you already have an existing service definition, or need to create one for the first time:

- **Creating a service definition for the first time.** In this scenario, you do not presently have a service definition, but would like to start exploiting workload management functions.

For the required steps, see “Creating a Service Definition for the First Time”.

- **Migrating to a new z/OS Release with an existing service definition.** In this scenario, you already have a service definition created on a previous release of z/OS or OS/390. You may need to be aware of a new service definition functionality level, possibly reallocate the couple data set (see note below), and accommodate any other changes introduced by the new release level.

For the required steps, see “Migrating to a New z/OS Release With an Existing Service Definition” on page 135.

**Note: The z/OS Release 2 version of the WLM administrative application only works with a WLM couple data set allocated by OS/390 Release 4 or above.** Because of changes to the structure of the WLM couple data set, you must use a CDS format level that is appropriate to your version of the administrative application. At present there are three CDS format levels (as shown in Table 11 on page 143). CDS format level 3 is required for use with the z/OS Release 2 version of the WLM administrative application. If your WLM couple data set was allocated by OS/390 Release 3 or below (CDS format level 1 or 2), you must reallocate it. This is required regardless of whether you want to use the new functions in OS/390 Release 4 and above. Previous versions of the WLM administrative application can be used with a WLM couple data set formatted by OS/390 Release 4 or above, provided that new function is not exploited.

This chapter contains a checklist for each of the above two scenarios. The major migration activities are described in detail following the checklist sections, as follows:

- “Restricting Access to the WLM Service Definition” on page 138
- “Start the Application and Enter/Edit the Service Definition” on page 139
- “Calculate the Size of the WLM Couple Data Set” on page 145
- “Allocate a WLM Couple Data Set” on page 145
- “Make a WLM Couple Data Set Available to the Sysplex for the First Time” on page 149
- “Make a Newly Formatted Couple Data Set Available to the Sysplex” on page 150
- “Migration Considerations for Velocity” on page 151
- “Migration Considerations for Discretionary Goal Management” on page 152
- “Migration Considerations for Protection of Critical Work” on page 153

---

### Creating a Service Definition for the First Time

Use the following checklist to create a service definition for the first time, allocate the WLM couple data set, and then activate a service policy.

1. **Set up performance objectives.** If your installation already has performance objectives, Chapter 4, “Setting up a Service Definition” on page 27 is still helpful in setting up a service policy with realistic performance goals.

2. **Set up a service definition** from the performance objectives.

To accomplish these two steps using the worksheets in Chapter 15, “Worksheets for Defining a Service Definition” on page 125, see Chapter 4, “Setting up a Service Definition” on page 27 through Chapter 14, “Defining Scheduling Environments” on page 113.

3. **Migrate to Goal mode before installing z/OS Release 4.**

In order to use workload management, you must be running in a sysplex. You can be running in a single system sysplex (monoplex) or a multi-system sysplex. For information about how to set up a sysplex, see *z/OS MVS Setting Up a Sysplex*. You should run MVS in compatibility mode (processing with your existing IPS and ICS) on your current release until you complete the steps for migrating to workload management and are comfortable switching into goal mode.

If you are running with mixed MVS releases on a sysplex, you need to install the PTFs for the appropriate compatibility APARs on all downlevel systems to enable different levels of workload management to coexist until you can upgrade the entire sysplex to the new release. Note that this is a cumulative list. An MVS/SP 5.1 system, for example, would need all of the APARs listed below.

Release	Install APARs:
MVS/SP 5.1	OW08866, OW20913, OW25831, OW30930, OW43856
MVS/SP 5.2, OS/390 R1, R2	OW20913, OW25831, OW30930, OW43856
OS/390 R3	OW25831, OW30930, OW43856
OS/390 R4, R5	OW30930, OW43856
OS/390 R6, R7, R8, R9	OW43856
OS/390 R10	None
z/OS R1	None

**Note:** The WLM service definition is stored in ISPF tables. When a new release adds certain specifications to the service definition, structural changes to the ISPF tables are required. When you open a service definition whose table structure is different (older) than the one currently used by the WLM application, the WLM application automatically updates the service definition structure. Once this occurs, the service definition cannot be read by older levels of the WLM application unless the compatibility APARs are installed.

The following releases changed the ISPF table structure:

- MVS 5.1
- MVS 5.2
- OS/390 Release 3
- OS/390 Release 4
- OS/390 Release 6
- OS/390 Release 10

If you ship the service definition PDS between systems (for example, via a TSO transmit), you need to verify that the sending and receiving systems are at compatible levels.

For complete information on installing z/OS, please see *z/OS and z/OS.e Planning for Installation*.

4. **Restrict access to the ISPF administrative application.**

For more information, see “Restricting Access to the WLM Service Definition” on page 138.

5. **Start the application** and enter the service definition.

For more information, see “Start the Application and Enter/Edit the Service Definition” on page 139.

6. **Upgrade the sysplex couple data set**

Make sure that you have installed the z/OS release and allocated your sysplex couple data set with the IXCL1DSU utility. The level of the sysplex couple data set must be at least MVS/ESA SP Version 5 in order to use the WLM couple data set.

For information about how to format the sysplex couple data set, see *z/OS MVS Setting Up a Sysplex*.

7. **Allocate a WLM couple data set.**

For more information, see “Allocate a WLM Couple Data Set” on page 145.

8. **Make the WLM couple data set available** for use in the sysplex for the first time by either:

- Issuing the SETXCF command
- Updating the COUPLExx parmlib member and re-ipling

For more information, see “Make a WLM Couple Data Set Available to the Sysplex for the First Time” on page 149.

9. **Install a service definition** on the WLM couple data set.

Before you can activate a service policy, you need to install the service definition on the WLM couple data set. To do this, use either the WLM ISPF application, or the Install Definition Utility (new to z/OS R3).

To use the ISPF application, go into it specifying the name of the PDS containing your service definition. From the Definition Menu, go to UTILITIES on the action bar. Then select the pull-down option “Install Service Definition”.

To use the Install Definition Utility, configure the sample JCL (member IWMINSTL, shipped in SYS1.SAMPLIB) as directed in the prolog. Once the JCL has been prepared, it can be started from the command console or submitted as a batch job. A simple service definition has been provided (member IWMSSDEF, also shipped in SYS1.SAMPLIB) that is available for those customers without any other definition.

10. **Adjust SMF recording**

Before you run your systems in goal mode with z/OS V1R3, you should be aware of the changes in your SMF recording. There are several changes to SMF records for goal mode.

In particular, you should turn off SMF type 99 records. They trace the actions SRM takes while in goal mode, and are written frequently. SMF type 99 records are for detailed audit information only. Before you switch your systems into goal mode, you should make sure you do not write SMF type 99 records unless you want them.

If you do chargeback based on SMF record type 30 or record type 72 records, you may need to update your accounting package.

For more information about SMF record changes for goal mode, see *z/OS MVS System Management Facilities (SMF)*.

#### 11. Activate a service policy.

Once you have installed a service definition, you can activate a service policy. You can activate a policy either from the administrative application, with the VARY operator command, or the WLM Install Definition Utility.

To activate a service policy from the application, choose the Utilities option from the action bar on the definition menu.

To activate a service policy with the VARY command, specify  
VARY WLM,POLICY=xxxx

where *xxxx* is the name of a policy defined in the installed service definition.

To activate a service policy with the Install Definition Utility, start or submit the sample JCL (member IWMINSTL in SYS1.SAMPLIB).

Once you issue the command, there is an active policy for the sysplex. Systems will start managing system resources to meet the goals defined in the service policy.

For more information about the VARY command, see *z/OS MVS System Commands*.

---

## Migrating to a New z/OS Release With an Existing Service Definition

Use the following checklist to migrate to z/OS Release 4 with an existing service definition. You may or may not need to reallocate the WLM couple data set, depending on which release you are migrating from.

1. **Evaluate your service definition.** At some point, either before or after you migrate to the new release, you may need to make one or more adjustments to your service definition, depending on which release you are migrating from:
  - **If you are migrating from OS/390 Release 2 or earlier:**

If you currently have more than 100 service classes defined, you must reduce the number to 100 or less. OS/390 Release 3 and later releases support at most 100 service classes.

**Important:** You need to make this adjustment **before** installing z/OS Release 2.

Also, see “Specifying I/O Priority Management” on page 84 for information on automatically managing I/O priorities in the sysplex based on service class goals.

Note that the following bullets apply to your system, as well.
  - **If you are migrating from OS/390 Release 3 or earlier:**

The addition of WLM-managed batch initiators will potentially change velocity values. See “Migration Considerations for Velocity” on page 151 for more information.

You do not need to adjust these values until you are ready to turn on WLM batch initiator management.

Note that the following bullets apply to your system, as well.
  - **If you are migrating from OS/390 Release 5 or earlier:**

Certain types of work, when overachieving their goals, may now have their resources “capped” in order to give discretionary work a better chance to run. See “Migration Considerations for Discretionary Goal Management” on page 152 for more information.

You can adjust these goals at any time, either before or after observing the effects of discretionary goal management.

Note that the following bullets may apply to your system, as well.
  - **If you are migrating from OS/390 Release 6 or earlier:**

Dynamic alias management may have migration implications if it is not yet available on other systems in a sysplex, or if it is available but disabled for a particular device. See “Migration Considerations for Dynamic Alias Management” on page 152 for more information.

Note that the following bullet may apply to your system, as well.
  - **If you are migrating from OS/390 Release 7 or earlier:**

It is recommended that multisystem enclaves be used only in a sysplex where all systems are at the OS/390 Release 8 level or higher. See “Migration Considerations for Multisystem Enclaves” on page 152 for more information.
  - **If you are migrating from OS/390 Release 9 or earlier:**

On OS/390 Release 10, you should be aware of new options available to help system administrators protect critical work, and how these options may affect other work. See Chapter 12, “Defining Special Protection Options for Critical Work” on page 89 for more information.

Also, be aware that if you increase your service definition to LEVEL011, your service definition will be subject to a more rigorous verification. Errors may

be flagged if you attempt to either save or install your service definition. The application will allow you to save the service definition, but will prevent you from installing it until the errors are corrected.

If you use the Release 10 application to modify a service definition below LEVEL011, these same conditions will instead appear as warnings when you save the service definition. In this case, it will *not* prevent you from installing the service definition.

For more information, see the discussion of LEVEL011 in “Service Definition Functionality Levels, CDS Format Levels, and WLM Application Levels” on page 140.

• **If you are migrating from OS/390 Release 10 or earlier:**

If you plan on using the Intelligent Resource Director functions available on z/OS Release 1, you may need to define a specific coupling facility structure named SYSZWLM\_XXXXYYYY in the CFRM policy. In addition, if you plan on using channel subsystem priority queuing, you’ll need to turn on I/O priority management, as described in “Specifying I/O Priority Management” on page 84.

See Chapter 18, “The Intelligent Resource Director” on page 159 for more information on these and other steps you may need to take in order to make the Intelligent Resource Director work in your installation.

**2. Install z/OS Release 4 on one system.**

If you are running with mixed MVS releases on a sysplex, you need to install the PTFs for the appropriate compatibility APARs on all downlevel systems to enable different levels of workload management to coexist until you can upgrade the entire sysplex to the new release. Note that this is a cumulative list. An MVS/SP 5.1 system, for example, would need all of the APARs listed here.

Release	Install APARs:
MVS/SP 5.1	OW08866, OW20913, OW25831, OW30930, OW43856
MVS/SP 5.2, OS/390 R1, R2	OW20913, OW25831, OW30930, OW43856
OS/390 R3	OW25831, OW30930, OW43856
OS/390 R4, R5	OW30930, OW43856
OS/390 R6, R7, R8, R9	OW43856
OS/390 R10	None
z/OS R1	None

**Note:** The WLM service definition is stored in ISPF tables. When a new release adds certain specifications to the service definition, structural changes to the ISPF tables are required. When you open a service definition whose table structure is different (older) than the one currently used by the WLM application, the WLM application automatically updates the service definition structure. Once this occurs, the service definition cannot be read by older levels of the WLM application unless the compatibility APARs are installed.

The following releases changed the ISPF table structure:

- MVS 5.1
- MVS 5.2
- OS/390 Release 3
- OS/390 Release 4
- OS/390 Release 6

- OS/390 Release 10

If you ship the service definition PDS between systems (for example, via a TSO transmit), you need to verify that the sending and receiving systems are at compatible levels.

For complete information on installing z/OS, please see *z/OS and z/OS.e Planning for Installation*.

**3. If you are migrating from OS/390 Release 3 or earlier, you need to reallocate the WLM couple data set.**

Because of changes to the structure of the WLM couple data set, you must use a CDS format level that is appropriate to your version of the administrative application. At present there are three CDS format levels (as shown in Table 11 on page 143). CDS format level 3 is required for use with the OS/390 Release 10 version of the WLM administrative application. If your WLM couple data set was allocated by OS/390 Release 3 or below (CDS format level 1 or 2), you must reallocate it.

Use the following checklist:

- a. Size the existing WLM couple data set** to determine how much space to allocate for the new one.  
For more information, “Calculate the Size of the WLM Couple Data Set” on page 145.
- b. Allocate a new WLM couple data set** using the information obtained in Step 3, and adding space if planning to exploit new functions.  
For more information, see “Allocate a WLM Couple Data Set” on page 145.
- c. Make the new WLM couple data set available**, replacing the old one.  
For more information, see “Make a Newly Formatted Couple Data Set Available to the Sysplex” on page 150.

**If you are migrating from OS/390 Release 4 or Release 5, you *may* need to reallocate the WLM couple data set**, based upon the following potential scenario:

If you use the new description fields when defining classification rules or classification groups (a new function in OS/390 Release 6), each field will make the service definition slightly larger. If your WLM couple data set is currently just barely large enough to accommodate your service definition, and you add enough of these new fields, it is possible that your service definition will become too large for your WLM couple data set. If this happens, the WLM application will issue a message, IWMAM047, WLM couple data set is too small to hold the service definition.

If you know that you’ll be adding many new description fields, you may wish to reallocate the WLM couple data set to prevent this potential problem.

**4. Start the WLM application**

For more information, see “Start the Application and Enter/Edit the Service Definition” on page 139. Please note the discussion of service definition functionality levels in “Service Definition Functionality Levels, CDS Format Levels, and WLM Application Levels” on page 140. Once you choose to use a new functionality level, from that point on you must always use a level of the WLM application that is compatible with that functionality level. It is recommended that you do not use the new functions (which automatically updates the functionality level) until all those systems are upgraded to the new release.

---

## Migration Activities

The following sections provide more detail for certain migration activities referenced out of the migration checklists. To determine if you need to perform an activity, refer to the preceding checklists.

### Restricting Access to the WLM Service Definition

Before you create a WLM service definition, you should determine who needs access to it, and the kind of functions each person needs to perform. The installation's capacity planners, systems programmers that analyze workloads and the system's performance, system operators, service administrator, and help desk support may all need access to the service definition information.

There are two levels of access you need to consider:

- Access to a partitioned data set
- Access to the WLM couple data set

#### Restricting Access to A Partitioned Data Set

Any user with access to the administrative application can create their own "practice" service definitions in their own partitioned data set. The partitioned data set containing the actual service definition that will be installed into the WLM couple data set should be protected. Use a data set profile, just as you would for any other data set. Give READ access only to those people who should be able to view the service definition in the partitioned data set, and UPDATE access to those people who should be able to create or modify the service definition in the partitioned data set.

#### Restricting Access to the WLM Couple Data Set

Once you have determined who needs access to the WLM couple data set itself, define the kind of access authority required, as follows:

**READ** With READ access, the user can extract a service definition from the WLM couple data set.

#### UPDATE

With UPDATE access, the user can:

- Do all the functions available for READ access
- Install a service definition to a WLM couple data set
- Activate a service policy.

To control access to the WLM couple data set, use RDEFINE to add a profile to the RACF database. Then use PERMIT to permit or deny access to the RACF profile. Do not forget to issue the SETROPTS REFRESH command after the PERMIT command to refresh the RACF data base and activate the changes you have made.

#### Example of RDEFINE for the WLM couple data set

```
RDEFINE FACILITY MVSADMIN.WLM.POLICY UACC(NONE) NOTIFY(user)
```

Where:

#### **user**

Indicates the user that should be notified of unauthorized access attempts to the database.

### Example of PERMIT for the WLM couple data set

```
PERMIT MVSADMIN.WLM.POLICY CLASS(FACILITY)
                                ID(user)
                                ACCESS(READ)

PERMIT MVSADMIN.WLM.POLICY CLASS(FACILITY)
                                ID(user)
                                ACCESS(UPDATE)
```

Where:

**user**

Indicates the user or user group that needs access to the WLM couple data set.

**ACCESS**

Indicates the type of access, either READ or UPDATE.

## Start the Application and Enter/Edit the Service Definition

Start the new release's WLM administrative application, and enter your service definition.

If you want to keep the previous release's WLM application, each version of the application must have unique library names. So, when you install the new release, make sure you keep the previous release application on the system under a unique name. You can rename your libraries by using an exit as described in Appendix A, "Customizing the WLM ISPF Application" on page 217.

When you enter the service definition, you can keep the service definition in an MVS partitioned data set until you are ready to install the service definition into the WLM couple data set. If you are migrating to use a new version of the WLM application, always save the partitioned data set created by the previous version of the application for backup purposes.

### Before Starting the WLM Application

To start the WLM application, you use the TSO/E REXX exec IWMARIN0. The exec concatenates (via LIBDEF and ALTLIB), the following libraries necessary to run the application:

SYS1.SBLSCLI0	Application REXX code
SYS1.SBLSKEL0	Application skeletons
SYS1.SBLSPNL0	Application panels
SYS1.SBLSTBL0	Application keylists and commands
SYS1.SBLSMSG0	Application messages

The exec also allocates some MVS partitioned data sets for the service definition using TSO ALLOCATE, and then invokes the WLM panels. If you have different data set conventions for your IPCS/WLM libraries, or if you use storage managed data sets, you should use the WLM application exits IWMAREX1 and IWMAREX2. For more information about how to code the exits, see Appendix A, "Customizing the WLM ISPF Application" on page 217.

## Start the WLM Application

To start the application, specify:

```
ex 'SYS1.SBLSCLI0(IWMARIN0)'
```

For more information about IWMARIN0, and for examples on how to start the application specifying the WLM exits, see Appendix A, “Customizing the WLM ISPF Application” on page 217.

## Enter/Edit the Service Definition

From the worksheets in Chapter 15, “Worksheets for Defining a Service Definition” on page 125, type in the service definition (or else edit the existing service definition), and then save it a partitioned data set (PDS). For help on using the application, see Chapter 19, “Using the WLM ISPF Application” on page 171.

Note that if you change the service definition to use certain functions of the new release, you may not be able to use the service definition on the previous release.

## Service Definition Functionality Levels, CDS Format Levels, and WLM Application Levels

A service definition has a **functionality level** for each release as shown in Table 10:

Table 10. Functionality Levels for Service Definition

Release	Functionality Level
MVS/ESA SP 5.1	LEVEL001
MVS/ESA SP 5.2	LEVEL002
OS/390 R3	LEVEL003
OS/390 R4/R5	LEVEL004
OS/390 R6	LEVEL006 or LEVEL007 (LEVEL007 available only with APAR OW33509 installed)
OS/390 R7	LEVEL007 or LEVEL008 (LEVEL008 available only with APAR OW39854 installed)
OS/390 R8	LEVEL007 or LEVEL008 (LEVEL008 available only with APAR OW39854 installed)
OS/390 R9	LEVEL008
OS/390 R10 and z/OS Release 1	LEVEL011
z/OS Release 2, 3, 4	LEVEL013

**Note:** LEVEL005, LEVEL009, LEVEL010, and LEVEL012 are reserved.

If you do not use any of the new functions for a new release, then the functionality level does not change, even if you are using the service definition on a new release. When you install the service definition, the system checks whether you have used any of the new functions, and sets the functionality level. For example, if you created your service definition on MVS/ESA SP 5.1, then its functionality level is LEVEL001. If you installed this service definition from a 5.2 system but did not use any of the new functions, then its functionality level remains LEVEL001.

You should use the new functions when you are comfortable running the new release on your sysplex. Once you use the new functions and increase the functionality level, then you may not be able to use the service definition on a lower level system. For example, you cannot extract a LEVEL002 service definition

from a MVS/ESA SP 5.1 system. You also cannot activate a policy in a LEVEL002 service definition from a MVS/ESA SP 5.1 system.

The following MVS/ESA SP 5.2 functions increase the service definition level to LEVEL002:

- Classification rules using:
  - Collection name
  - Correlation information
  - Connection type
  - Package name
  - Plan name
- Classification groups for:
  - Connection type
  - Package name
  - Plan name
- Classification rule extensions
- Service definition extensions

The following OS/390 Release 3 functions increase the service definition level to LEVEL003:

- Application environment defined
- Use of classification qualifiers:
  - Procedure name
  - Perform
  - Perform group
- More than 255 report classes
- Setting the I/O Priority Management service option to YES

The following OS/390 Release 4 functions increase the service definition level to LEVEL004:

- Scheduling environment or scheduling environment resources defined
- Use of Priority classification qualifier
- Use of SYSTEM and SYSSTC service class names within classification rules

For OS/390 Release 5, the service definition level remains at LEVEL004.

**Note:** LEVEL005 is reserved.

The following OS/390 Release 6 functions increase the service definition level to LEVEL006:

- Use of the description field in classification rules or classification groups

The following functions, available on either OS/390 Release 6 with APAR OW33509 installed, or on OS/390 Release 7 and higher, increase the service definition level to LEVEL007:

- Use of the PC (process name) classification qualifier

The following functions, available on either OS/390 Release 7 or 8 with APAR OW39854 installed, or on OS/390 Release 9, increase the service definition level to LEVEL008:

- Changing the Dynamic Alias Management service option from NO (the default) to YES.

**Note:** LEVEL0009 and LEVEL010 are reserved.

The following functions, available on OS/390 Release 10 and z/OS Release 1, increase the service definition level to LEVEL011:

- Using the CPU protection, storage protection, or exemption from transaction server management options
- Use of classification qualifiers:
  - Scheduling environment name
  - Subsystem collection name
  - Sysplex name
  - System name
  - System name group

**Note:** As soon as you increase your service definition level to LEVEL011 or above, your service definition will be subject to a more rigorous verification. The following errors will now be flagged if you attempt to either save or install your service definition:

- Classifying both CICS/IMS transactions and any other (non-CICS/IMS) work into the same service class. The application will return either message IWMAM911E (if you create a rule for a non-CICS/IMS subsystem type which references a CICS/IMS service class), or IWMAM912E (if you specify a CICS/IMS service class as the default service class for a non-CICS/IMS subsystem).
- In a multiperiod service class, going from one period to the next with an increase in the importance level. For example, a period 1 goal with an importance level of 4 and a period 2 goal with an importance level of 3 would be flagged. (The application returns message IWMAM910E to identify the out-of-order importance.) The importance level must always either decrease or remain the same.

The application will allow you to save the service definition, but will prevent you from installing it until the errors are corrected. You can use the “Verify service definition” option on the Definition Menu, as described in “Using the Menu Bar on the Definition Menu” on page 178, to ensure that your service definition does not contain any of these errors.

If you use the Release 10 application to modify a service definition below LEVEL011, these same conditions will instead be flagged as warnings (IWMAM910W, IWMAM911W, and IWMAM912W) when you save the service definition. In this case, it will *not* prevent you from installing the service definition.

The following function, available on z/OS Release 2, increases the service definition level to LEVEL013:

- New subsystem SYSH

Table 12 on page 144 summarizes the functions that you are allowed to perform from a system depending on the functionality level of the service definition. This table assumes that the WLM administrative application is at the same level as the MVS system where it is being used. In general, the WLM application must be at the same level as the MVS system it runs on in order to access the WLM couple data set (using the extract, install, or activate functions). The table illustrates the following points:

- The WLM application can handle any service definition at a functionality level equal to or lower than the one that corresponds to the WLM application release level.

- If you wish to install a service definition and activate a service policy with the z/OS Release 2 WLM application, you must use a WLM CDS allocated by the format utility from OS/390 Release 4 or above, otherwise known as CDS format level 3, as shown below in Table 11.

*Table 11. WLM Couple Data Set Format Levels*

<b>CDS Format Level</b>	<b>Description</b>
<b>1</b>	Original format, as built in MVS/SP Release 5.1. This is the format level for MVS/SP Release 5.1 and 5.2.
<b>2</b>	Format updated in OS/390 Release 3, with addition of application environments. This is the format level for OS/390 Release 3 only.
<b>3</b>	Format updated in OS/390 Release 4, with addition of scheduling environments. This is the format level for OS/390 Release 4 through Release 10, and z/OS Release 1 and Release 2.

Table 12. Workload Management Compatibility

Functionality Level of Service Definition	WLM Function To Be Performed	WLM Level: MVS/SP 5.1	WLM Level: MVS/SP 5.2	WLM Level: OS/390 R3	WLM Level: OS/390 R4/R5	WLM Level: OS/390 R6	WLM Level: OS/390 R6+/R7/R8 *	WLM Level: OS/390 R7+/R8+/R9 **	WLM Level: OS/390 R10, z/OS R1	WLM Level: z/OS R2,R3,R4
LEVEL001	Open/Save PDS Extract from CDS Install/Activate CDS 1 (pre-R3 format) Install/Activate CDS 2 (R3 format) Install/Activate CDS 3 (R4 and above format)	YES YES YES YES YES	YES YES YES YES YES	YES YES NO YES YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES
LEVEL002	Open/Save PDS Extract from CDS Install/Activate CDS 1 (pre-R3 format) Install/Activate CDS 2 (R3 format) Install/Activate CDS 3 (R4 and above format)	NO NO NO NO NO	YES YES YES YES YES	YES YES NO YES YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES
LEVEL003	Open/Save PDS Extract from CDS Install/Activate CDS 1 (pre-R3 format) Install/Activate CDS 2 (R3 format) Install/Activate CDS 3 (R4 and above format)	NO NO NO NO NO	NO NO NO NO NO	YES YES NO YES YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES
LEVEL004	Open/Save PDS Extract from CDS Install/Activate CDS 1 or 2 Install/Activate CDS 3 (R4 and above format)	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES
LEVEL006	Open/Save PDS Extract from CDS Install/Activate CDS 1 or 2 Install/Activate CDS 3 (R4 and above format)	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES
LEVEL007 *	Open/Save PDS Extract from CDS Install/Activate CDS 1 or 2 Install/Activate CDS 3 (R4 and above format)	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES
LEVEL008 **	Open/Save PDS Extract from CDS Install/Activate CDS 1 or 2 Install/Activate CDS 3 (R4 and above format)	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO YES	YES YES NO NO YES	YES YES NO NO YES	YES YES NO NO YES
LEVEL011	Open/Save PDS Extract from CDS Install/Activate CDS 1 or 2 Install/Activate CDS 3 (R4 and above format)	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO YES	YES YES NO NO YES	YES YES NO NO YES
LEVEL013	Open/Save PDS Extract from CDS Install/Activate CDS 1 or 2 Install/Activate CDS 3 (R4 and above format)	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO	NO NO NO NO NO
PDS	ISPF partitioned data set containing a saved service definition									
CDS	WLM couple data set — number denotes CDS format level, as described in Table 11 on page 143									
* LEVEL007 is available on OS/390 Release 6 with APAR OW33509 installed, or on OS/390 Release 7 or 8.										
** LEVEL008 is available on OS/390 Release 7 or 8 with APAR OW39854 installed, or on OS/390 Release 9.										

## Calculate the Size of the WLM Couple Data Set

When you reallocate a WLM couple data set, it must be at least as large as the current one or you will not be able to make it available to the sysplex. It is recommended that you increase the size of your couple data set to allow you to exploit new functions, even if you do not plan to exploit new functions immediately.

**If you are using OS/390 Release 3 or later and your couple data set was allocated at OS/390 Release 3 or later:**

Use the WLM ISPF Application Utilities option (see “Utilities” on page 179) with the “Allocate couple data set using CDS values” option to determine your couple data set size. If you already have a record of the values, you can skip this step; just ensure you allocate the WLM couple data set with the same or larger values as the current one.

## Allocate a WLM Couple Data Set

You need to define a WLM couple data set for storing the service definition information. If you are running a sysplex with mixed release levels, you should format the WLM couple data set from the highest level system. This allows you to use the current level of the WLM application. You can continue to use the downlevel WLM application on downlevel systems provided that you do not attempt to exploit new function.

To allocate the WLM couple data set you can either use the facility provided in the WLM application, or you can run an XCF utility. For each case, you need to estimate how many workload management objects you are storing on the WLM couple data set. You must provide an approximate number of the following:

- Policies in your service definition
- Workloads in your service definition
- Service classes in your service definition.

The values you define are converted to space requirements for the WLM couple data set being allocated. The total space is *not* strictly partitioned according to these values. For most of these values, you can consider the total space to be one large pool. (For the actual layout, see Table 13 on page 148.)

If you specify 50 service classes, for instance, you are simply requesting that the space required to accommodate 50 service classes be added to the couple data set. You have NOT limited yourself to 50 service classes in the service definition. Although note that if you DO define more than 50 service classes, you will use up space that was allocated for something else.

You should define an alternate WLM couple data set (similar to the sysplex alternate couple dataset) for recovery purposes. You can define an alternate WLM couple data set using the same method (either the ISPF application or the XCF utility), but specifying a different data set.

To allocate a WLM couple data set using the ISPF application, choose Utilities from the Definition Menu.

To allocate a WLM couple data set using the XCF utility, you can follow some JCL provided in the IWMFTCDS member of SYS1.SAMPLIB.

## Sample JCL to Allocate a WLM Couple Data Set

To allocate a WLM couple data set, use the sample JCL and fill in the following information:

```
//FMTDCDS JOB MSGLEVEL=(1,1)
//STEP1 EXEC PGM=IXCL1DSU
//STEPLIB DD DSN=SYS1.MIGLIB,DISP=SHR
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
        DEFINEDS SYSPLEX( PLEX1)
                DSN(SYS1.WLMCDS01) VOLSER(TEMPAK)
                MAXSYSTEM(32)
                CATALOG
        DATA TYPE(WLM)
                ITEM NAME(POLICY) NUMBER(10)
                ITEM NAME(WORKLOAD) NUMBER(35)
                ITEM NAME(SRVCLASS) NUMBER(30)
                ITEM NAME(SVDEFEXT) NUMBER(5)
                ITEM NAME(SVDCREXT) NUMBER(5)
                ITEM NAME(APPLENV) NUMBER(50)
                ITEM NAME(SVAEAEXT) NUMBER(5)
                ITEM NAME(SCHENV) NUMBER(50)
                ITEM NAME(SVSEAEXT) NUMBER(5)
```

Where:

### **SYSPLEX**

The name of your sysplex as it appears in your COUPLExx parmlib member.

### **DSN**

The name you are calling your WLM couple data set

### **VOLSER**

A volume that you have access to. If you are using DFSMS, you do not need to specify a VOLSER.

### **TYPE**

The type of function for which this data set is allocated. For a service definition, the type is WLM.

### **ITEM NAME(POLICY) NUMBER( )**

Specifies that an increment of space large enough to accommodate the specified number of policies be allocated in the WLM couple data set. (Default=5, Minimum=1, Maximum=99)

### **ITEM NAME(WORKLOAD) NUMBER( )**

Specifies that an increment of space large enough to accommodate the specified number of workloads be allocated in the WLM couple data set. (Default=32, Minimum=1, Maximum=999)

### **ITEM NAME(SRVCLASS) NUMBER( )**

Specifies that an increment of space large enough to accommodate the specified number of service classes be allocated in the WLM couple data set. (Default=128, Minimum=1, Maximum=999)

**Note:** As of OS/390 Release 3, WLM allows no more than 100 service classes to be defined in a service definition. The default, however, remains at the pre-OS/390 Release 3 value of 128. This will set aside as much space as you will ever need for service classes, as well as a little extra for other WLM objects.

If you are migrating from a pre-OS/390 Release 3 system and you previously reserved space for more than 128 service classes, do the following:

- If you intend to use SETXCF PSWITCH to make the new WLM couple data set active, continue to reserve space for the same number of service classes as you did in the earlier release. SETXCF PSWITCH will not allow a smaller WLM couple data set to be added as an alternate couple data set.
- If you intend to IPL the sysplex with the new WLM couple data set and install the service definition from a saved ISPF data set, then reserve space for as many service classes as you ever expect to define.

**ITEM NAME(SVDEFEXT) NUMBER( )**

Specifies that an exact amount of space (in K bytes) for extension areas to the WLM Service Definition (IWMSVDEF) be allocated in the WLM couple data set. (Default=0, Minimum=0, Maximum=8092) *See note below.*

**ITEM NAME(SVDCREXT) NUMBER( )**

Specifies that an exact amount of space (in K bytes) for extension areas to the WLM Service Definition Classification Rules (IWMSVDCR) be allocated in the WLM couple data set. (Default=0, Minimum=0, Maximum=8092) *See note below.*

**ITEM NAME(APPLENV) NUMBER( )**

Specifies that an increment of space large enough to accommodate the specified number of application environments be allocated in the WLM couple data set. (Default=100, Minimum=1, Maximum=999)

**ITEM NAME(SVAEAEXT) NUMBER( )**

Specifies that an exact amount of space (in K bytes) for extension areas to the WLM Service Definition Application Environment Area (IWMSVAEA) be allocated in the WLM couple data set. (Default=0, Minimum=0, Maximum=8092)

**ITEM NAME(SCHENV) NUMBER( )**

Specifies that an increment of space large enough to accommodate the specified number of scheduling environments be allocated in the WLM couple data set. (Default=100, Minimum=1, Maximum=999)

**ITEM NAME(SVSEAEXT) NUMBER( )**

Specifies that an exact amount of space (in K bytes) for extension areas to the WLM Service Definition Scheduling Environment Area (IWMSVSEA) be allocated in the WLM couple data set. (Default=0, Minimum=0, Maximum=8092)

If you encounter a problem during processing, make sure you take a dump by adding the following to your JCL and re-submit.

```
//SYSABEND DD SYSOUT=*
```

**Note:** The intended users of SVDEFEXT, SVDCREXT, SVAEAEXT, and SVSEAEXT are system management product vendors who wish to include some of their own unique information about customer workload definitions along with the WLM definitions. The WLM interfaces allow these extensions to accompany the service class definitions, report class definitions, or even classification rules. The amount of extra information is specific to each product that exploits these interfaces. That product's documentation should tell the customer how to set SVDEFEXT, SVDCREXT, SVAEAEXT, and SVSEAEXT to ensure that there is sufficient space available in the WLM couple data set to hold the extra information. For more information, see the "Adding Program-Specific Extensions to a Service Definition" section of the "Using the Administrative Application Services" chapter in *z/OS MVS Programming: Workload Management Services*.

## Increasing the Size of the WLM Couple Data Set

You must use a series of SETXCF commands to add a new, larger couple data set as the primary WLM couple data set. During this processing you may encounter a message from XCF (IXC250I) indicating your new couple data set is too small. The message indicates which subrecords had insufficient space. If this occurs you must reallocate the new couple data set with a larger size.

### If you are using OS/390 Release 3 or later and your couple data set was allocated at OS/390 Release 3 or later:

From the Definition Menu in the ISPF application, choose Utilities. In the Utilities pull-down, choose "4. Allocate couple data set." Record the values that you see in that panel. Then choose "5 Allocate couple data set using CDS values." Record the values that you see in that panel. Now compare the two sets of values and choose the *highest* from each category. For example, if the two panels showed the following values:

Allocate couple data set	Allocate couple data set using CDS values
Service policies . . . _5	Service policies . . . 10
Workloads . . . . . _40	Workloads . . . . . _35
Service classes . . . _35	Service classes . . . _30
Application environments . . . . 100	Application environments . . . . _50
Scheduling environments . . . . _80	Scheduling environments . . . . _50

Then you should use the highest values in each category, in this case 10, 40, 35, 100, and 80, to allocate your new couple data set.

### If you are using a pre-OS/390 Release 3 system or your couple data set was *not* allocated at OS/390 Release 3 or later:

To get initial values based on your current couple data set, see "Calculate the Size of the WLM Couple Data Set" on page 145. Table 13 shows each subrecord as identified in the XCF message and which WLM parameters used with the IXCL1DSU utility affect that subrecord. For each subrecord it is recommended to increase the value of the first keyword listed in the table. If that is insufficient, then increase the value of the next keyword listed, and repeat until you have enough space for the subrecord.

Table 13. Mapping of WLM Couple Data Set Subrecords to IXCL1DSU Keywords

WLM Couple Data Set Subrecord	Applicable IXCL1DSU Keywords	Release Where Introduced
IWMSVDEF	SRVCLASS POLICY WORKLOAD SVDEFEXT	5.1
IWMSVDCR	SRVCLASS SVDCREXT	5.1
IWMSVAEA	APPLENV SVAEAEXT	OS/390 R3
IWMSVSEA	SCHENV SVSEAEXT	OS/390 R4
IWMSVPOL	SRVCLASS WORKLOAD SVDEFEXT	5.1

Table 13. Mapping of WLM Couple Data Set Subrecords to IXCL1DSU Keywords (continued)

WLM Couple Data Set Subrecord	Applicable IXCL1DSU Keywords	Release Where Introduced
IWMSVPCR	SVDCREXT	5.1
IWMSVPAE	APPLENV SVAEAEXT	OS/390 R3
IWMSVPSE	SCHENV SVSEAEXT	OS/390 R4

If you used the sizer utility (IWMARSZ) to determine initial item values, you can use the following steps to help identify the problem:

- Run IWMARSZ with the DEBUG option against the current primary WLM couple data set.
- Run IWMARSZ with the DEBUG option against the new WLM couple data set that you are trying to add.
- Compare the lengths of the records until you find a discrepancy, then adjust the keywords according to Table 13 on page 148.

## Make a WLM Couple Data Set Available to the Sysplex for the First Time

This section applies when creating a service definition for the first time. If you already have a service definition and want to make a re-allocated WLM couple data set available, see “Make a Newly Formatted Couple Data Set Available to the Sysplex” on page 150.

To make your WLM couple data set available to the sysplex, you must either:

- Update your COUPLExx parmlib member to include the data set name and volume of your WLM couple data set, and re-IPL. Use this option if you have not yet IPLed in a sysplex.
- Issue the SETXCF command, if you are already IPLed in a sysplex. You must still update your COUPLExx member for subsequent IPLs.

### Using the SETXCF Command

To make the WLM couple data set available to the sysplex, you can use the SETXCF command. Remember that you still need to update your COUPLExx member as shown in “Updating the COUPLExx Member” on page 150 so that any subsequent IPLs will automatically pick up the WLM couple data sets.

For more information about using the SETXCF command, see *z/OS MVS System Commands*.

#### Examples of SETXCF Command

- To make a primary WLM couple data set called SYS1.WLMCDS01 residing on volume TEMP01 available to the sysplex, enter the following command:  

```
SETXCF COUPLE,TYPE=WLM,PCOUPLE=(SYS1.WLMCDS01,TEMP01)
```
- To make an alternate WLM couple data set called SYS1.WLMCDS02 residing on volume TEMP02 available to the sysplex, enter the following command:  

```
SETXCF COUPLE,TYPE=WLM,ACOUPLE=(SYS1.WLMCDS02,TEMP02)
```

## Updating the COUPLExx Member

To make the WLM couple data set available for use in the sysplex, you need to update the DATA keyword in the COUPLExx parmlib member, and IPL so that the member is in use. For more information about updating the COUPLExx member, see *z/OS MVS Setting Up a Sysplex*.

### Example of Updating COUPLExx Member

```
DATA      TYPE(WLM)
          PCOUPLE(SYS1.WLMCDS01,TEMP01)
          ACOUPLE(SYS1.WLMCDS02,TEMP02)
```

Where:

#### TYPE

The function type, WLM.

#### PCOUPLE(*dataset.name,volume*)

The WLM couple data set name, and the volume it resides on.

#### ACOUPLE(*dataset.name,volume*)

The alternate WLM couple data set name, and the volume it resides on. If you do not have an alternate WLM couple data set, then delete this keyword.

Specify the modified COUPLExx member on your next IPL.

## Make a Newly Formatted Couple Data Set Available to the Sysplex

This section applies when you already have WLM couple data sets, have just allocated new WLM couple data sets, and want to make them available to the sysplex.

There are different reasons why you would want to do this. One reason is when you would like to increase the size of the existing WLM couple data set. Another reason is migrating your WLM couple data set to the current format (OS/390 Release 4 and later).

You must use a series of SETXCF command to switch from the currently active primary and alternate couple data sets to the new couple data sets. All systems in the sysplex then operate with the newly allocated data set.

If you're making a newly formatted WLM couple data sets available to the sysplex, you can continue to use an older WLM application to modify, install and activate your service definition (as long as new functions are not exploited), or you can switch to using the z/OS Release 2 WLM application.

For more information on compatibility of release levels, WLM application levels, couple data set formats, and functionality levels, see "Service Definition Functionality Levels, CDS Format Levels, and WLM Application Levels" on page 140.

### Example of Making Re-allocated Couple Data Sets Available

1. Allocate two new couple data sets as described in “Allocate a WLM Couple Data Set” on page 145. For this example, it is assumed you want a primary and an alternate couple data set, and that the names of the new data sets are SYS1.WLMP residing on volume SYS001 and SYS1.WLMA residing on volume SYS002.

2. Make SYS1.WLMP the alternate using the command:

```
SETXCF COUPLE,TYPE=WLM,ACOUPLE=(SYS1.WLMP,SYS001)
```

As part of this processing, SETXCF copies the contents of the current primary WLM couple data set to SYS1.WLMP which now is the new alternate.

3. Switch SYS1.WLMP to primary using the command:

```
SETXCF COUPLE,TYPE=WLM,PSWITCH
```

4. Now make SYS1.WLMA the new alternate using the command:

```
SETXCF COUPLE,TYPE=WLM,ACOUPLE=(SYS1.WLMA,SYS002)
```

As in Step 1 above, this causes the contents of the new primary WLM couple data set SYS1.WLMP to be copied to the new alternate SYS1.WLMA.

## Migration Considerations for Velocity

Initiation delays cause the velocity value to decrease. Recalculate and adjust your velocity goals accordingly. See “Velocity Formula” on page 42 for information on calculating velocity.

Before migrating to WLM batch management, you can estimate the new velocity goal for a service class as follows:

**Note:** All jobs with the same service class should be migrated together to WLM-managed job classes.

- Run the jobs under normal circumstances
- Examine the initiation delay data:
  - In the IWMWRCAA data area if you are using the workload reporting services:

#### RCAETOTDQ

Total delay samples, including initiation delay

#### RCAETOTU

Total using samples

- In the SMF Type 72 record, Subtype 3, if you are using RMF:

#### R723CTDQ

Total delay samples, including initiation delay

#### R723CTOU

Total using samples

Include the initiation delay in the velocity formula for an estimate of the new, lower velocity. Plugging this delay data into the velocity formula gives you:

$$\frac{\text{RCAETOTU}}{\text{RCAETOTU} + \text{RCAETOTDQ}} \times 100$$

Or:

$$\frac{\text{R732CTOU}}{\text{R723CTOU} + \text{R723CTDQ}} \times 100$$

RMF will do this calculation for you — look for the **INIT MGMT** field in the RMF Monitor I workload activity report (on the line that begins “VELOCITY MIGRATION:”). See “Adjusting Velocity Goals Based on Samples Included in Velocity Calculation” on page 45 for more information.

If you had originally given a velocity goal to a service class period *only* because TYPRUN=HOLD time was included in response time goals, you can now give that service class period a response time goal because the TYPRUN=HOLD time is no longer included in the response time. In this case, you no longer need to recalibrate the velocity goal since it has been replaced with the response time goal.

## Migration Considerations for Discretionary Goal Management

In OS/390 Release 6 and later, certain types of work, *when overachieving their goals*, potentially will have their resources “capped” in order to give discretionary work a better chance to run. Specifically, work that is *not* part of a resource group and has one of the following two types of goals will be eligible for this resource donation:

- A velocity goal of 30 or less
- A response time goal of over one minute.

Work that is eligible for resource donation may be affected in OS/390 Release 6 and later if this work has been significantly overachieving its goals. If you have eligible work that must overachieve its goals to provide the required level of service, adjust the goals to more accurately reflect the work’s true requirements.

## Migration Considerations for Dynamic Alias Management

On OS/390 Release 7 or 8 with APAR OW39854 installed, or on OS/390 Release 9, dynamic alias management is supported for parallel access volumes. With dynamic alias management, WLM can automatically perform alias address reassignments to help work meet its goals and to minimize IOS queueing.

It is recommended *not* to use dynamic alias management for a device unless all systems sharing that device have dynamic alias management enabled. Otherwise, WLM will be attempting to manage alias assignments without taking into account the activity from the non-participating systems.

See “Specifying Dynamic Alias Management” on page 85 for more information.

## Migration Considerations for Multisystem Enclaves

Before using multisystem enclaves, an installation needs to define a specific coupling facility structure named SYSZWLM\_WORKUNIT in the CFRM policy. See Chapter 17, “Defining a Coupling Facility Structure for Multisystem Enclave Support” on page 155 for more information.

In addition, the CFRM couple data set needs to be formatted to support system-managed rebuild, which can only be done on systems at an OS/390 Release 8 level or higher. Any systems at a level below OS/390 Release 8 would be unable to access the CFRM couple data set, and therefore would be unable to connect to a coupling facility structure. It is recommended that multisystem enclaves be used only in a sysplex where all systems are at the OS/390 Release 8 level or higher.

Programs that use data from the SMF 30 record may need to be updated in conjunction with multisystem enclave support. The enclave owner's SMF 30 record has new fields containing the CPU time accumulated by all of its split transactions, for *all* systems on which they executed.

For more detailed information on multisystem enclaves, see the "Creating and Using Enclaves" in *z/OS MVS Programming: Workload Management Services*.

## **Migration Considerations for Protection of Critical Work**

On OS/390 Release 10 and above, you should be aware of new options available to help system administrators protect critical work, and how these options may affect other work.

These options include:

- Long-term storage protection
- Long-term CPU protection
- Exemption from management as a transaction server

See Chapter 12, "Defining Special Protection Options for Critical Work" on page 89 for more information on these options.



---

## Chapter 17. Defining a Coupling Facility Structure for Multisystem Enclave Support

Some work managers split large transactions across multiple systems in a parallel sysplex, improving the transaction's overall response time. These work managers can use multisystem enclaves to provide consistent management and reporting for these types of transactions.

Among the benefits of using multisystem enclaves:

- All parts of a split transaction are managed to the same service class. If the service class has multiple periods, the CPU usage of the entire transaction is used to switch periods.
- The enclave owner's SMF 30 record includes CPU time accumulated by all of its split transactions, for *all* systems on which they executed.

Before using multisystem enclaves, an installation needs to define a specific coupling facility structure named `SYSZWLM_WORKUNIT` in the CFRM policy. Once the CFRM policy with this structure definition is activated, WLM will automatically connect to the structure, enabling the use of multisystem enclaves.

This chapter shows how to define the `SYSZWLM_WORKUNIT` structure, a prerequisite to the use of multisystem enclaves. For more information on defining coupling facilities, see *z/OS MVS Setting Up a Sysplex*.

In addition, the CFRM couple data set needs to be formatted to support system-managed rebuild, which can only be done on systems at an OS/390 Release 8 level or higher. Any systems at a level below OS/390 Release 8 would be unable to access the CFRM couple data set, and therefore would be unable to connect to a coupling facility structure. It is recommended that multisystem enclaves be used only in a sysplex where all systems are at the OS/390 Release 8 level or higher.

Programs that use data from the SMF 30 record may need to be updated in conjunction with multisystem enclave support. The enclave owner's SMF 30 record has new fields containing the CPU time accumulated by all of its split transactions, for *all* systems on which they executed.

For more detailed information on multisystem enclaves, see the "Creating and Using Enclaves" in *z/OS MVS Programming: Workload Management Services*.

---

### Defining the Coupling Facility

It may be difficult to size the `SYSZWLM_WORKUNIT` structure at first, as there is no sure way to know exactly how many parallel units-of-work may exist at any given time. The best option is take a best guess at the initial and maximum sizes and then alter the structure size based on performance and/or change in demand.

If the structure's maximum size is defined too low, work managers will experience failures when they try to export enclaves. It is the work manager's responsibility to respond to such a failure. The work requests may instead be run locally (increasing the response time), or the work requests may fail.

The best way to estimate the storage size needed is to use the CFSIZER tool, which can be accessed through the following web page:

<http://www.ibm.com/s390/cfsizer>

Alternately, there are formulas in *PR/SM Planning Guide* to help estimate the storage size needed. As shown in Table 14, the TDEC value is the estimated number of concurrently executing parallel units-of-work. Use the TDEC estimate along with the other values explicitly given in the table, as follows:

Table 14. Values to Use in Storage Estimation Formulas

Value	Description	Specify:
TDEC	Total directory entry count — the maximum number of concurrently executing parallel units-of-work	Best Estimate
TDAEC	Total data area element count	TDEC X 2
MSC	Maximum number of storage classes	1
MCC	Maximum number of castout classes	1
MDAS	Maximum number of data area elements associated with a directory entry	32
DAEX	Data area element characteristic	3
AAI	Adjunct assignment indicator	0
R_de/R_data	Directory to data ratio	1/2

Once you have estimated the initial and maximum sizes for the SYSZWLM\_WORKUNIT structure, define the structure as described in *z/OS MVS Setting Up a Sysplex*. Keep in mind the following:

- WLM requests a coupling facility with “default” connectivity.
- Non-volatility is not required.
- The coupling facility control code must be at CFLEVEL 9 or higher.

The following sample JCL shows the definition of a SYSZWLM\_WORKUNIT structure:

```
//POLICYX JOB
//STEP1 EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=A
//SYSIN DD *

DATA TYPE(CFRM) REPORT(YES)

DEFINE POLICY NAME(POLICY1) REPLACE(YES)

CF NAME(FACIL01)
TYPE(123456)
MFG(IBM)
PLANT(02)
SEQUENCE(123456789012)
PARTITION(1)
CPCID(00)
SIDE(0)
DUMPSPACE(2000)

CF NAME(FACIL02)
TYPE(123456)
MFG(IBM)
PLANT(02)
SEQUENCE(123456789012)
PARTITION(2)
CPCID(00)
SIDE(1)
DUMPSPACE(2000)

STRUCTURE NAME(SYSZWLM_WORKUNIT)
SIZE(2000)
INITSIZE(1000)
PREFLIST(FACIL02,FACIL01)
```

---

## Shutting Down the Coupling Facility

If it becomes necessary to shut down a coupling facility containing the SYSZWLM\_WORKUNIT structure (either to apply maintenance or to reconfigure), there are two options:

- If another coupling facility has enough storage available, use the XES system-managed rebuild function to rebuild the SYSZWLM\_WORKUNIT structure into another coupling facility. See *z/OS MVS System Commands* for more information.
- If there is no other coupling facility into which the SYSZWLM\_WORKUNIT structure can be rebuilt, the structure will be deleted when its coupling facility is shut down and therefore multisystem enclave support will be disabled (as described in “Coupling Facility Failures”).

An installation should take the appropriate steps to quiesce any active work which may be using multisystem enclaves before shutting down the coupling facility containing the SYSZWLM\_WORKUNIT structure.

---

## Coupling Facility Failures

If the coupling facility containing the SYSZWLM\_WORKUNIT structure fails, or if the structure itself fails, then all existing multisystem enclaves will be lost. It is the work manager’s responsibility to respond to such a failure. The work manager may fail the work requests, or it may process them without using multisystem enclaves.

If another coupling facility is available, WLM will automatically create a new (empty) SYSZWLM\_WORKUNIT structure in it. New multisystem enclaves can now be created for new work requests.

If the original coupling facility is still intact, but the *link* fails, then the use of multisystem enclaves is temporarily disabled. Again, it is the work manager’s responsibility to respond to this situation, either failing the work requests, or processing them without using multisystem enclaves. When the link is restored, then the use of multisystem enclaves can continue.

---

## Chapter 18. The Intelligent Resource Director

The **Intelligent Resource Director (IRD)** extends the concept of goal-oriented resource management by allowing you to group system images that are resident on the same physical server running in LPAR mode, and in the same parallel sysplex, into an "LPAR cluster." This gives Workload Management the ability to manage processor and channel subsystem resources, not just in one single image but across the entire cluster of system images. The following figure shows one LPAR cluster in one central processor complex (CPC)

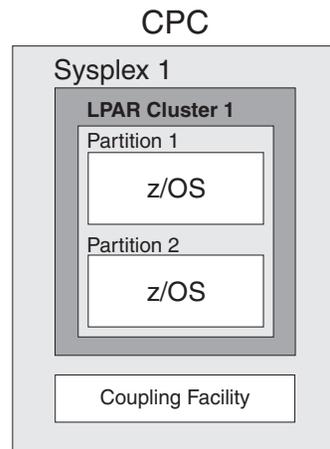


Figure 19. One LPAR Cluster on One CPC

A CPC can have multiple LPAR clusters supporting different parallel sysplexes, and a parallel sysplex can in turn comprise multiple LPAR clusters in different CPCs. This is illustrated in the following figure, in which two sysplexes across two CPCs are grouped into four LPAR clusters:

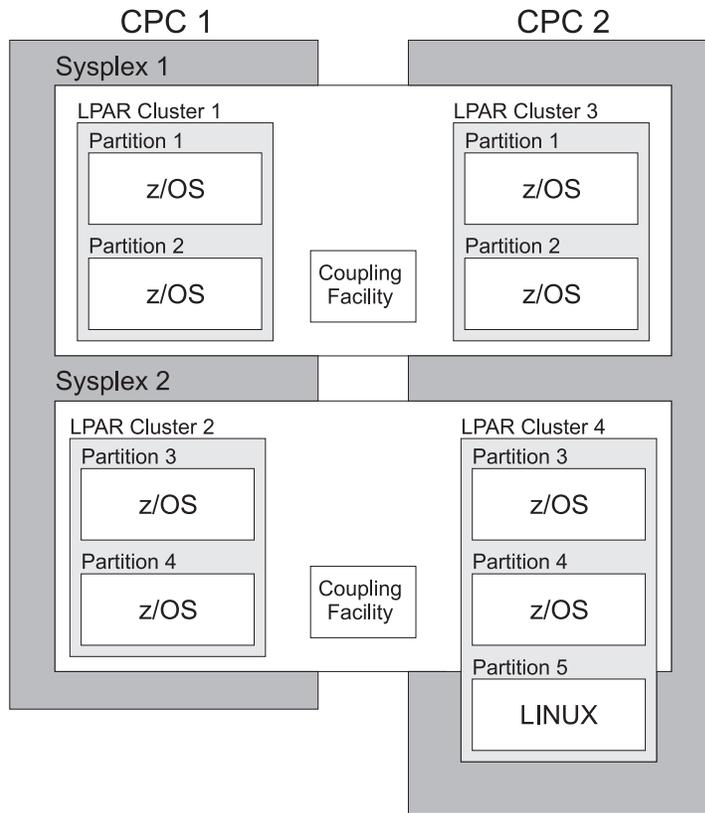


Figure 20. Four LPAR Clusters on Two CPCs

WLM manages a parallel sysplex by directing work to the available resources. With the Intelligent Resource Director, WLM can additionally move resources within an LPAR cluster to the work. Processor resources are automatically moved to the partitions with the greatest need, based on the business goals of the workloads they are running. Channels are automatically moved to the I/O control units with the greatest need, based on the business goals of the workloads using them.

If a z/OS system was running in a partition managed by the Intelligent Resource Director, and you intend to IPL any non z/OS system or OS/390 system in the same partition, you need to deactivate and reactivate the partition. This is necessary because the Intelligent Resource Director might have taken offline CPUs when the z/OS partition was running. If you IPL a non z/OS operating system without deactivating and reactivating the partition, these offline CPUs will not come online.

The LINUX partition is not part of the parallel sysplex, but WLM manages it as if it were part of it.

The three functions that make up the Intelligent Resource Director are as follows:

- LPAR CPU Management
- Dynamic Channel Path Management
- Channel Subsystem Priority Queuing.

**Note:** Dynamic channel path management and channel subsystem priority queuing also are functional in single-system environments, such as a CPC

operating in basic mode, or a z/OS system operating as a monoplex. If a z/OS system is operating as a monoplex, each z/OS system must be assigned a different sysplex name.

---

## LPAR CPU Management

LPAR CPU Management allows dynamic adjustment of processor resources across logical partitions in the same LPAR cluster. WLM achieves this with two different mechanisms:

- **LPAR Weight Management**

When you divide your central processing complex into separate logical partitions, each partition is assigned its own “LPAR weight,” which corresponds to the percentage of overall processing power that’s guaranteed to the work in that partition. Previously, if the workload shifted to the extent that more processing power was needed in a particular partition, the weights had to be changed manually. With LPAR Weight Management, you give each logical partition an initial LPAR weight, along with a minimum and maximum weight. WLM will then dynamically balance these weights to best meet the goals of the work in the partitions, with no human intervention. Note that the total weight of the cluster as a whole will remain constant, so LPARs outside the cluster are unaffected.

LPAR Weight Management takes effect when two or more LPARs in the cluster are running CPU-constrained, and the CPC’s shared physical CPs are fully utilized. When just one LPAR is using all or most of the CPC because other LPARs are idle, then LPAR Weight Management will have no effect on resource distribution.

If there are more than one non z/OS partitions in a LPAR cluster, the Intelligent Resource Director manages the weight of these partitions by exchanging weight among these partitions. LPAR weight of any non z/OS partition will not be reset to its initial weight when the z/OS partition is reset or deactivated. Non z/OS partitions weight can be reset through the service element panel in this situation.

- **VARY CPU Management**

VARY CPU Management works for z/OS partitions only. It works hand-in-hand with LPAR Weight Management. As the LPAR weights change, the number of online logical CPUs are also changed to maintain the best match between logical CPU speed and physical CPU speed. Optimizing the number of logical CPUs benefits workloads that have large amounts of work done under single tasks, and minimizes LPAR overhead for all workloads.

LPAR CPU Management requires zSeries servers in z/Architecture mode. z/OS images require a CF level 9 coupling facility structure.

LINUX for S/390 or z/Series requires kernel 2.4 or higher. General purpose CPUs are supported, but Integrated Facility for LINUX (IFL) CPUs are not supported.

---

## Dynamic Channel Path Management

Prior to Dynamic Channel Path Management, all channel paths to I/O control units had to be statically defined. In the event of a significant shift in workload, the channel path definitions would have to be reevaluated, manually updated via HCD, and activated or POR’ed into the configuration. Dynamic Channel Path Management lets Workload Management dynamically move channel paths through the ESCON Director from one I/O control unit to another, in response to changes in the workload requirements. By defining a number of channel paths as “managed,” they become eligible for this dynamic assignment. By moving more

bandwidth to the important work that needs it, your DASD I/O resources are used more efficiently. This may decrease the number of channel paths you need in the first place, and could improve availability — in the event of a hardware failure, another channel could be dynamically moved over to handle the work requests.

Dynamic Channel Path Management operates in two modes: balance mode and goal mode. In balance mode, Dynamic Channel Path Management will attempt to equalize performance across all of the managed control units. In goal mode, which is available only when WLM is operating in goal mode on *all* systems in an LPAR cluster, WLM will still attempt to equalize performance, as in balance mode. In addition, when work is failing to meet its performance goals due to I/O delays, WLM will take additional steps to manage the channel bandwidth accordingly, so that important work meets its goals.

Dynamic Channel Path Management requires z/OS and a zSeries server in z/Architecture mode, running in either Basic or LPAR mode. If a system image running Dynamic Channel Path Management in LPAR mode is defined as being part of a multisystem sysplex, it also requires a CF level 9 coupling facility structure, even if it is the only image currently running on the CPC.

See the WLM Web site at:

<http://www.ibm.com/servers/eserver/zseries/zos/wlm/>

for the latest list of devices that support Dynamic Channel Path Management:

---

## Channel Subsystem Priority Queuing

Channel Subsystem Priority Queuing is an extension of the existing concept of I/O priority queuing. Previously, I/O requests were handled by the channel subsystem on a first-in, first-out basis. This could at times cause high priority work to be delayed behind low priority work. With Channel Subsystem Priority Queuing, if important work is missing its goals due to I/O contention on channels shared with other work, it will be given a higher channel subsystem I/O priority than the less important work. This function goes hand in hand with the Dynamic Channel Path Management described above — as additional channel paths are moved to control units to help an important workload meet goals, Channel Subsystem Priority Queuing ensures that the important workload receives the additional bandwidth before less important workloads that happen to be using the same channel.

WLM sets the priorities using the following basic scheme:

- System related work is given the highest priority
- High importance work missing goals is given a higher priority than other work
- Work meeting goals is managed so that light I/O users will have a higher priority than heavy I/O users
- Discretionary work is given the lowest priority in the system.

Channel Subsystem Priority Queuing requires z/OS and a zSeries server in z/Architecture mode. It works in both basic mode and LPAR mode, and does not require a coupling facility structure.

## Example

To illustrate how the Intelligent Resource Director works in a mixed workload environment, here's an example:

- You have three workloads running on one server:
  - **Online Transactions**, your most important workload. This runs only during the day shift.
  - **Data Mining**, which has a medium importance. This is always running, and will consume as much resource as you give it.
  - **Batch**, which is your lowest importance work. Like data mining, it is always running, and will consume as much resource as you give it.
- In this example, the server is divided into two logical partitions (and both partitions are in the same sysplex):
  - **Partition 1** runs both the online transactions and the batch work, as they happen to share the same database.
  - **Partition 2** runs the data mining work.

### Day Shift

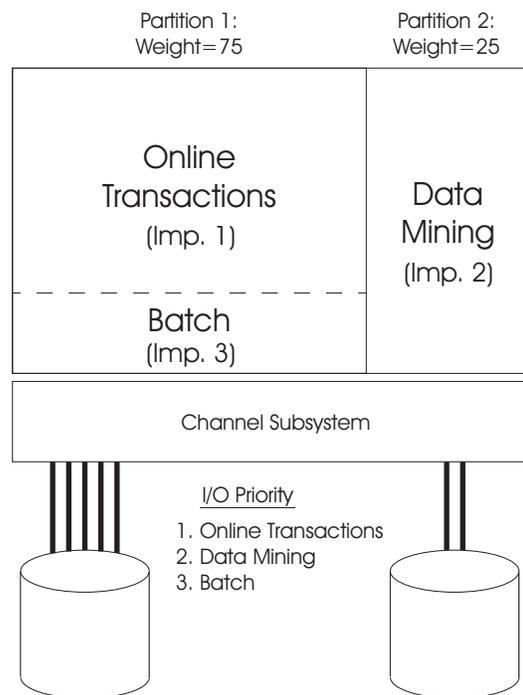


Figure 21. Intelligent Resource Director Example – Day Shift

This first figure shows a day shift configuration. As the online transaction workload is the most important, Partition 1 is given a high enough weight to ensure that the online transaction work does not miss its goals due to CPU delay. Within the partition, the existing Workload Management function is making sure that the online transaction work is meeting its goals before giving any CPU resource to the batch work.

The DASD used by the online transaction work is given enough channel bandwidth to ensure that channel path delays do not cause the work to miss its goals. The channel subsystem I/O priority ensures that online transaction I/O

requests are handled first. Even though the batch work is running in Partition 1 (with the increased partition weight and channel bandwidth), the data mining I/O requests will still take precedence over the batch I/O requests if the data mining work is not meeting its goals.

## Night Shift

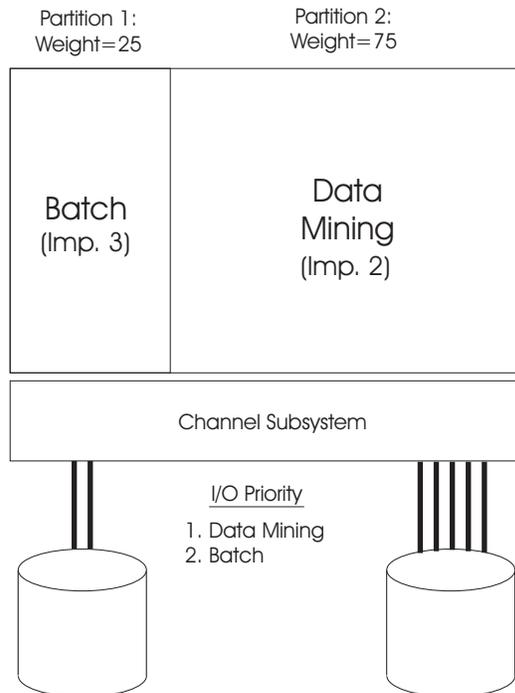


Figure 22. Intelligent Resource Director Example – Night Shift

This figure shows the night shift, when there are no more online transactions. If the partition weights had remained the same, then the batch work would be consuming most of the CPU resource, and using most of the I/O bandwidth, even though the more important data mining work may still be missing its goals. LPAR CPU Management automatically adjusts to this change in workload, adjusting the partition weights accordingly. Now the data mining work will receive the CPU resource it needs to meet its goals. Similarly, dynamic channel path management will move most of the I/O bandwidth back to the data mining work.

## Making the Intelligent Resource Director Work

There are several tasks you need to perform to make the Intelligent Resource Director work in your installation:

- Defining the SYSZWLM\_XXXXYYYY Coupling Facility Structure
- Enabling LPAR CPU Management
- Enabling Dynamic Channel Path Management
- Enabling Channel Subsystem Priority Queuing.

## Defining the SYSZWLM\_xxxxxyyyy Coupling Facility Structure

Before using either LPAR CPU Management or Dynamic Channel Path Management in a multisystem sysplex, you'll need to define a specific coupling facility structure named SYSZWLM\_xxxxxyyyy in the CFRM policy for each LPAR cluster.

In each SYSZWLM\_xxxxxyyyy structure, the 8-character xxxxyyyy field represents a portion of the CPU ID for the CPC. To obtain this 8-character field, issue the D M=CPU command, as follows:

```
SYS1 d m=cpu
SYS1 IEE174I 13.51.42 DISPLAY M 884
PROCESSOR STATUS
ID CPU SERIAL
1 + 1315122064
```

Use the *last* eight characters in the returned CPU ID. So in this example, the coupling facility structure would be named SYSZWLM\_15122064.

Define the structure as in the following sample JCL code:

```
//POLICYX JOB
//STEP1 EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=A
//SYSIN DD *
DATA TYPE(CFRM) REPORT(YES)
DEFINE POLICY NAME(POLICY1) REPLACE(YES)
...
STRUCTURE NAME(SYSZWLM_15122064)
MINSIZE(10000) /* 1K UNITS */
INITSIZE(10000)
SIZE(10000)
PREFLIST(FACIL01,FACIL02)
```

**Note:** In this sample code, a fixed structure size of 10 MB is used. This would be sufficient for most installations. Alternately, you could define a smaller structure size (it is recommended that it be at least 6MB) and use auto alter to allow the structure to increase in size as necessary.

Activate the new CFRM policy using the SETXCF START command, as follows:

```
SETXCF START,POL,POLNAME=POLICY1,TYPE=CFRM
```

Once the CFRM policy is activated, WLM will automatically connect to the structure, enabling the use of the Intelligent Resource Director.

Keep in mind the following:

- WLM requests a coupling facility with “default” connectivity.
- Non-volatility is not required.
- The coupling facility control code must be at CFLEVEL 9 or higher.

If it becomes necessary to shut down a coupling facility containing the SYSZWLM\_xxxxxyyyy structure (either to apply maintenance or to reconfigure), there are two options:

- If another coupling facility has enough storage available, use the XES system-managed rebuild function to rebuild the SYSZWLM\_xxxxxyyyy structure into another coupling facility.
- If there is no other coupling facility into which the SYSZWLM\_xxxxxyyyy structure can be rebuilt, the structure will be deleted when its coupling facility is

shut down and therefore LPAR CPU Management or Dynamic Channel Path Management in a multisystem sysplex will be disabled.

If the coupling facility containing the SYSZWLM\_XXXXYYYY structure fails, or if the structure itself fails, then all LPAR CPU Management or Dynamic Channel Path Management in a multisystem sysplex is stopped. If another coupling facility is available, WLM will automatically create a new (empty) SYSZWLM\_XXXXYYYY structure in it. These functions can now resume.

If the original coupling facility is still intact, but the *link* fails, then these functions will be temporarily disabled. When the link is restored, they will resume.

## Enabling LPAR CPU Management

Once you have defined the SYSZWLM\_XXXXYYYY coupling facility structure, the remaining actions you need to take to enable LPAR CPU Management will all occur on the hardware management console (HMC).

- In the Primary Support Element Workplace, make sure that “Not dedicated central processors” is selected in the activation profile of each logical partition. This ensures that the logical partition will use shared CPs instead of dedicated CPs. On the same panel, specify the initial and reserved number of logical CPs you wish to have available to each logical partition. Production partitions which potentially need access to the full power of the CPC should be defined with the maximum number of logical CPs (equal to the number of shared physical CPs). Other partitions can be defined with a fewer number of logical CPs if they need less power, or if you wish to specifically restrict them to less power. At a minimum, the number of logical CPs should be sufficient to achieve the partition’s maximum weight.

For some logical partitions, you may not need to make any of the above changes, as the settings are already correct. If you do need to make any of the above changes, note that you’ll need to deactivate and then reactivate the logical partition for the changes to take effect. For the remaining actions, you will *not* need to do this after making changes.

- For each logical partition that will participate in LPAR Weight Management, do the following:
  - Make sure that “Initial Capping” is turned off. WLM cannot manage the weight of a logical partition that is capped.
  - Enter the initial processing weight. This becomes the logical partition’s weight when it is first IPLed.
  - Enter the minimum and maximum weights. These set the lower and upper limits for the weights that WLM will assign to the logical partition.
  - Check the “WLM Managed” box. This is the final step in activating LPAR Weight Management.

Once you have made the above changes for z/OS images, you do not have to do anything else to activate VARY CPU Management. You can disable VARY CPU Management, if you wish, by adding the keyword VARYCPU=NO to the IEAOPTxx parmlib member and then issuing the SET OPT=xx command to activate the change. The scope of this command is by single system. If there are multiple systems in the LPAR cluster, then the other systems will continue to use VARY CPU Management. You can set VARYCPU=YES to return the system to VARY CPU Management. Note that if any CPs were taken offline by an operator, those CPs will need to be configured back online before then can again be managed by VARY CPU Management.

**Note:** If VARY CPU Management is enabled, then any program that specifies the CPUMASK keyword on the IEAMSCHD macro, or directly updates the SRBCPAFF field in an SRB or the TCBAFFN field in a TCB, must be modified to use the IWMCPAFN macro. It is not necessary to use the IWMCPAFN macro when CPU affinity is set in the program properties table in the SCHEDxx parmlib member. It is also not necessary to use it for programs that use cryptographic coprocessors. (WLM will not vary offline processors that have cryptographic coprocessors associated with them.) See *z/OS MVS Programming: Workload Management Services* for more information on the IWMCPAFN macro.

## Enabling Non z/OS CPU Management

In addition to the actions you have already taken to enable LPAR CPU Management, you need to do the following to enable non z/OS CPU Management:

- Before activating the LINUX partition, specify the CP management cluster name under the Options tab of the activation profile. This causes LPAR to group the LINUX partition with the appropriate sysplex.
- For the LINUX partition the system name must be set in order to activate LPAR CPU management. WLM needs the LINUX system name in order to handle the LINUX partition. This also applies if you only use the PX qualifier. The sysplex name specified in the service element panel must be the same as the z/OS sysplex name. The system name must be unique in relation to the other system names in the same sysplex.

Set the system name by running the

**insmod hwc\_cpi.o system\_name=sysname** command. **hwc\_cpi.o** is currently being shipped in the directory: */lib/modules/x.y.z/kernel/drivers/s390/char/*, where x.y.z is the kernel version. If you change the LINUX system name you must re-IPL to make it known to LPAR CPU Management.

VSE/ESA and z/VM set the system name during IPL.

- Goals for non z/OS partitions are specified in the WLM service definition. You can define velocity goals, but no discretionary goals. Multiple periods are not supported.

## Enabling Dynamic Channel Path Management

When using Dynamic Channel Path Management, your CPC can be running in either basic or LPAR mode. In basic mode, a coupling facility structure is not needed, as there will be no logical partitions grouped together in an LPAR cluster. In LPAR mode, a coupling facility structure is required if you wish to use dynamic channel path management in any logical partition containing a system that is a member of a multisystem complex (even if the system image is the only member of that sysplex on this CPC). You do not need a coupling facility structure if all the logical partitions are running in XCFLOCAL or MONOPLEX mode.

The IBM Redbook *z/OS Intelligent Resource Director* (SG24-5952) provides extensive guidance on choosing the appropriate channels and control units for Dynamic Channel Path Management. Once you have selected them, there are two specific HCD definitions that will need to be changed:

- **Channel Definitions**

When defining (or modifying) a channel, you must specify YES in the “Managed” field. You must also specify a dynamic switch (in this case, the ESCON Director) to which the channel is attached. It is recommended that you also specify the entry switch ID and entry port so that HCD can do consistency checking.

If you are running in LPAR mode, you must also define the name of the sysplex to which a logical partition must belong in order to have this channel in its configuration. Specify this name in the "I/O Cluster" field. Finally, you must also define this channel as being shared. Note that unlike traditional shared channels which potentially can be shared by all logical partitions on a CPC, managed channels can only be shared by logical partitions in the same LPAR cluster.

**Note:** Ensure that each LPAR cluster name (the sysplex name that is associated with the LPAR cluster, as described above) is uniquely named across the entire CPC. As managed channels have an affinity to a specific LPAR cluster, non-unique names would create problems with the scope of control.

- **Control Unit Definitions**

Whereas non-managed channel paths (otherwise called "static" channel paths) are defined in the traditional way, via the CHPID number, a managed path is defined by specifying a double asterisk ("\*\*"). The number of double asterisks you specify will limit the number of managed channel paths per LPAR cluster. The total number of non-managed and managed channel paths, per CPC, cannot exceed 8.

**Important**

You must define at least one non-managed channel path (which must be defined as shared) per control unit.

The control unit must be attached to a switch (again, in this case an ESCON Director) which in turn must be attached to managed channels.

After changing the HCD definitions, the remaining actions to enable Dynamic Channel Path Management occur in the Primary Support Element Workplace, as follows:

- For basic or LPAR mode, ensure that the CPC's reset profile is enabled for dynamic I/O.
- For an LPAR cluster environment, ensure that all partitions in the cluster are authorized to control the I/O configuration.
- In addition, ensure that the automatic input/output (I/O) interface reset option is enabled in the CPC's reset profile. This will allow Dynamic Channel Path Management to continue functioning in the event that one participating system image fails.

If at any time you wish to disable Dynamic Channel Path Management, issue the SETIOS DCM=OFF command. The SETIOS DCM=ON command will turn it back on. Issue SETIOS DCM=REFRESH to refresh the control unit model table (for instance, to include a new IOST $nnn$  load module provided by the control unit's manufacturer).

**Note:** After issuing the SETIOS DCM=OFF command, your I/O configuration may now be unable to handle your workload needs, as it will now be in whatever state Dynamic Channel Path Management left it before being disabled. You may need to activate a new I/O configuration that will meet your workload needs across the entire LPAR cluster.

## Enabling Channel Subsystem Priority Queuing

To enable Channel Subsystem Priority Queuing, you'll need to do the following:

- In the WLM ISPF application, make sure that "I/O Priority Queuing" is set to YES on the "Service Coefficients/Service Definition" panel.
- If your CPC is partitioned, click on the "Change LPAR I/O Priority Queuing" icon in the Primary Support Element Workplace. You'll see a list of all of the logical partitions. Define the range of I/O priorities that will be used by each image, specifying the minimum and maximum I/O priority numbers. It is recommended that you use a range of eight values (for example 8 to 15), as this will correspond to the number of values in the range that WLM uses. While it is not enforced, it is assumed that you'll set the same range of I/O priorities for all images in the same LPAR cluster. You can, however, prioritize multiple LPAR clusters on the same CPC by setting different ranges for each LPAR cluster.

If you have a partition running a system other than z/OS (for instance OS/390, VM, TPF, VSE) specify an appropriate default priority by setting the minimum and maximum to the same number. For example, if you have a VM partition running work that is equal in importance to the discretionary work in a z/OS partition, you could set the VM partition's range to 8–8, and the z/OS partition's range to 8–15. If you have critical OLTP applications running in a partition, on the other hand, you could set the range for that partition to 16–16, and the z/OS partition's range to 8–15. In this way, the OLTP work would always have a higher priority than the z/OS work.

- In the Primary Support Element Workplace, click on the "Enable I/O Priority Queuing" icon. You'll see a simple panel that reads "Global input/output (I/O) priority queuing" — click on the "Enable" box. Channel Subsystem Priority Queuing is now enabled for the entire CPC.

## Handling Downlevel Releases

An OS/390 system can participate in an LPAR cluster with partitions running z/OS, but with the following restrictions:

- The partition will always run with its initial processing weight. Even as the weights of the other partitions are adjusted in response to changes in workload, the OS/390 partition's weight will remain constant.
- The partition will never participate in Dynamic Channel Path Management. All of the z/OS partitions in the cluster can participate, but any non-z/OS partitions will be ignored.
- The work in this partition can only be assigned one constant I/O priority, as described in "Enabling Channel Subsystem Priority Queuing".

---

## For More Information

- For more detailed information on the Intelligent Resource Director, see the IBM Redbook, *z/OS Intelligent Resource Director*.
- For more information on defining coupling facilities, see *z/OS MVS Setting Up a Sysplex*.
- For more information on PR/SM, see the *PR/SM Planning Guide*.
- For more information on the Primary Support Element Workplace, see the *Support Element Operations Guide*.
- For more information on HCD, see the *z/OS HCD User's Guide*.
- For more information on RMF, see the *z/OS RMF User's Guide*.
- For more information on the IEAOPTxx parmlib member, see the *z/OS MVS Initialization and Tuning Reference*.

- For more information on any of the system commands mentioned here, see *z/OS MVS System Commands*.

---

## Chapter 19. Using the WLM ISPF Application

This chapter explains how to use the ISPF application. It explains the functions that are available, and how you can navigate through the panels.

---

### Before You Start

Before you go into the application, you should have already created a service definition using the worksheets provided in Chapter 15, “Worksheets for Defining a Service Definition” on page 125.

Once you have created at least one service policy and your classification rules on the worksheets, you are ready to start using the ISPF application. Your service policies and classification rules make up a service definition. You can store a service definition in the following kinds of data sets:

#### **WLM couple data set**

In order for all systems in a sysplex to process with an active service policy, they must all be able to access a service policy. They all access the policy from a WLM couple data set. To use workload management, you must allocate a WLM couple data set, define it to the sysplex, and install your service definition onto it. You can allocate the WLM couple data set from the application. Only one service definition can be installed on the WLM couple data set.

#### **MVS partitioned data set (PDS)**

You do not need to preallocate the data sets. You specify a data set name, and the application allocates it for you. You can save one service definition per MVS PDS.

#### **Notes:**

1. If you use customized data sets in your installation, or if you use DFSMS, you can use WLM application exits IWMAREX1 and IWMAREX2 to specify those changes. See Appendix A, “Customizing the WLM ISPF Application” on page 217 for how to code the exits.
2. The data set *userid.WLM.SAVExx* (where *userid* is the TSO ID running the application and *xx* is some numeric value such as SAVE01) is allocated by the WLM application for recovery and is deleted by WLM upon exiting the application. This naming convention should therefore *not* be used for a new service definition.

---

### Panel Areas and How to Use Them

Most panels have a menu bar, action field, status line, scrollable area, function key area, and command line. You tell the application what actions to perform by making choices or typing information on a panel.

In this topic, examples of panels and pop-ups are shown to help familiarize you with the product. The examples closely match what you see on your terminal, but in some cases the spacing or function key settings may not exactly match what you see on your terminal.

## Using the Menu Bar

A menu bar at the top of every panel shows the actions you can take on that panel. Press F10, or the Home key on some terminals, to move the cursor to the beginning of the menu bar from any position on a panel.

Figure 23 shows an example of the menu bar on the Definition menu.



Figure 23. Menu Bar on the Definition Menu

To select an action, use the Tab or cursor movement keys to position the cursor on your choice, then press ENTER.

When you select an option on the menu bar, WLM displays a pull-down with choices related to the option you selected. While a pull-down is displayed, only the actions in the pull-down or on the menu bar are available. For example, if you select a pull-down and the option you want is not listed in it, you can select another pull-down on the menu bar.

To select an option in a pull-down, type the number of your choice in the action field and press ENTER. You can also use the cursor movement keys to position the cursor on your choice, then press ENTER. Figure 24 shows an example of the pull-down choices on the File option on the Definition Menu.

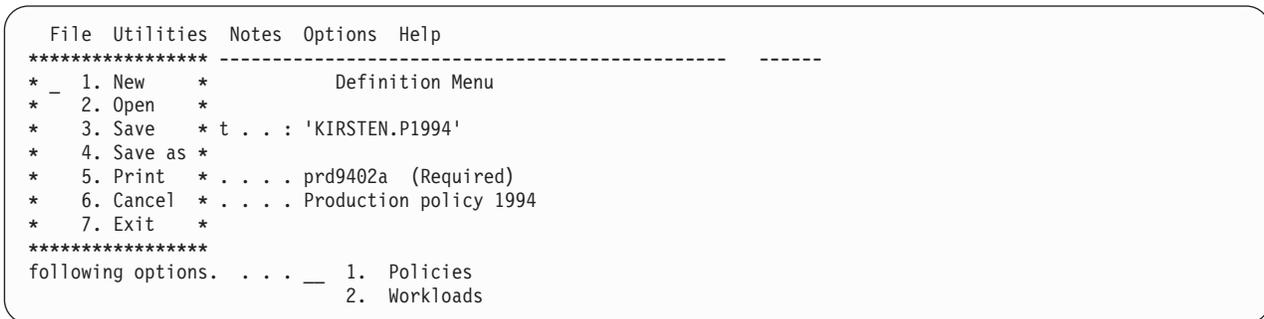


Figure 24. Definition Menu File Choices

### Using the Menu Bar on Selection Lists

The menu bar is a bit different on selection lists. You type a slash next to the name of the object you want to work with in the action field. Then, move to the menu bar, select an option, and press ENTER. You then choose the desired action from the pull-down.

For example, from the service class selection list, choose the STC\_1 service class, move to the menu bar on the Service-Class option, and press ENTER. Then, type 3 in the menu pull-down, and press ENTER.

```

Service-Class View Notes Options Help
*****-----
* 3 1. Create *      Service Class Selection List      ROW 14 TO 26 OF 29
* 2. Copy *
* 3. Modify *
* 4. Browse * Create, 2=Copy, 3=Modify, 4=Browse, 5=Print, 6=Delete,
* 5. Print * Menu Bar
* 6. Delete *
* 7. Exit *
*****
Description                                Workload
--- MDLIMSR  IMS Model response              TRNIMS
--- PRDCICS  MSA application                        PRDCICS
--- PRDCICSM Prod IMS non-response             PRDCICS
--- PRDIMSNR IMS response                      PRDIMS
/  STC_1    Highest Priority stc                STC
--- STC_2    High Priority stc                 STC

```

Figure 25. Service Class Selection List

### Using the Status Line

Some, but not all panels have a status line. A status line is displayed on the right side of a panel or pop-up beneath the title. The status line indicates the number of items or lines currently displayed in a list or topic, and the total size of and your current location in that list or topic.

For example, a service class selection list status line such as **Row 1 to 8 of 16** states that displayed in the panel is a list of eight of the 16 service classes contained in the service definition.

### Using the Scrollable Area

On selection list type panels, there is a scrollable area that contains a list or text. In a service class selection list you see a list of service classes. The status line shows how many service classes there are in the list. You scroll backward or forward to move through the list. Selection lists can also be pop-ups, such as selecting a workload from a list to associate it with a service class.

### Using the Menu Bar on a Selection List

On selection lists, you can select the object you want to work with before you select actions you want to perform from the menu bar.

To mark items for selection in a scrollable area, type a slash (/) over the underscore in front of the listed choice. When you select an menu bar item, the application performs the action for the marked items only.

Figure 26 on page 174 shows an example of the scrollable area on a service class selection list. The first two service classes are marked with a / for actions on the menu bar.

```

Service-Class View Notes Options Help
-----
                Service Class Selection List                ROW 13 TO 24 OF 29

Action Codes: 1=Create, 2=Copy, 3=Modify, 4=Browse, 5=Print, 6=Delete,
              /=Menu Bar

Action  Class      Description                      Workload
/      PRDIMS     IMS Production                    PRDIMS
/      PRDCICS    PRDCICS                          PRDCICS
_      PRDCICSM   MSA application                   PRDCICS

```

Figure 26. Service Class Selection List

### Using the Action Field

The action field is where you specify the action to take. On the definition menu, the action field is where you specify the workload management object that you want to work with. The action codes are standard in the application--except on a few selection lists, such as the subsystem type selection list in classification rules, and the service policy selection list. Figure 27 shows the action codes and action field on the subsystem type selection list.

**Note:** The menu bar pull-down choices for the file match the action codes on the selection lists. So you can choose the method according to your preference.

```

Subsystem-Type View Notes Options Help
-----
                Subsystem Type Selection List for Rules        ROW 1 TO 10 OF 10
Command ==>> _____

Action Codes: 1=Create, 2=Copy, 3=Modify, 4=Browse, 5=Print, 6=Delete,
              /=Menu Bar

Action  Type      Description                      -----Class-----
_      ASCH      IBM-defined USAA almost modified  Service Report
_      CICS      IBM-defined USAA modified         PRDCICS

```

Figure 27. Action field on the Subsystem Type Selection List

### Using the Command Line

The command line is displayed according to the currently active user profile. You can issue the following commands from the command line:

- =value**  
Repeats the previous command.
- BACKWARD**  
Scrolls backward.
- DOWN**  
Scrolls forward.
- EXIT** Exits the panel.
- FORWARD**  
Scrolls forward.
- FKA ON|OFF**  
Determines whether to display the function key area.

**HELP** Displays the help panel for the displayed panel.

**KEYLIST**

Displays the keylist utility where you can adjust PF key settings.

**PFSHOW**

Shows the PF key settings.

**RETRIEVE**

Displays any previous command.

**UP** Scrolls backward.

## Using the Function Keys

The function key area at the bottom of each panel or pop-up displays actions that you can complete by pressing a function key. When a pop-up is displayed, you can press only the function keys listed in that pop-up, not the keys listed at the bottom of the panel.

Standard actions are assigned to function keys 1 through 12 and are repeated for function keys 13 through 24. The function key assignments can vary slightly, depending on options selected during installation. If you want to customize the key settings for your installation, you can use the KEYLIST utility. See “Customizing the Keylists” on page 224 for more information about how to customize the key settings.

The WLM application displays function keys 1 to 12. If you want to display function keys 13 to 24, or see the standard function key settings, see *ISPF Dialog Management Guide and Reference*.

Figure 28 shows a sample of the function keys on the Service Class Selection list.

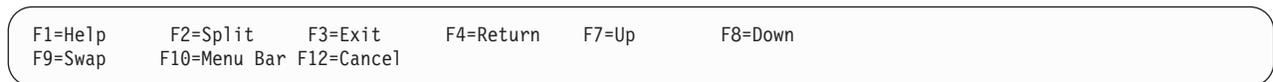


Figure 28. Function Key Area

---

## Starting the WLM Application

The application is shipped in the IPCS library. When you start the application, the system needs to concatenate the WLM/IPCS data sets, allocate some data sets, and then invoke the WLM panels.

To start the application, specify:

```
ex 'SYS1.SBLSCLI0(IWMARIN0)'
```

If you have different data set conventions for your IPCS data sets, or if you use storage managed data sets, you should use the WLM application exits IWMAREX1 and IWMAREX2. For more information about IWMARIN0 and how to customize the application with the WLM exits, see Appendix A, “Customizing the WLM ISPF Application” on page 217.

---

## Now You're Started

Upon entry to the interface, the WLM logo panel is displayed. Press ENTER to continue, and the application displays the *Specify Definition* pop-up as shown in Figure 29. You should then specify which service definition you want to work with. You can work with one service definition at a time in the application.

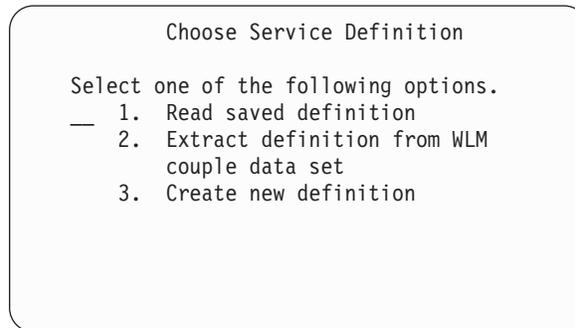


Figure 29. Choose Service Definition Pop-Up

If this is your first time in the application, choose option 3, *Create a new service definition*. This option brings you to the definition menu, where you can define your service definition. Once you created a new service definition, upon exiting the application, you can either:

- Save the service definition in a PDS  
The application prompts you for a data set name, allocates the PDS, and saves the service definition.
- Install the service definition on the WLM couple data set  
The install option puts the service definition currently displayed in the application out on the WLM couple data set. When installed, any changes you made to the service definition are available when you activate a policy. “Installing and Extracting a Service Definition” on page 209 explains how to install a service definition.
- Discard the service definition.

If you want to edit a service definition previously defined and stored in an MVS PDS, choose option 1 *Read saved definition*.

If you want to work with the service definition on the WLM couple data set, choose option 2 *Extract definition from WLM couple data set*. You must have previously installed a service definition on the WLM couple data set. “Installing and Extracting a Service Definition” on page 209 explains how to extract a service definition.

---

## Using the Definition Menu

The definition menu is the central place for entering your service definition. When you set up a service definition, you must enter a service definition name and optionally, a description on the Definition Menu.

Figure 30 on page 177 shows a sample Definition Menu with the service definition name and a description filled in.

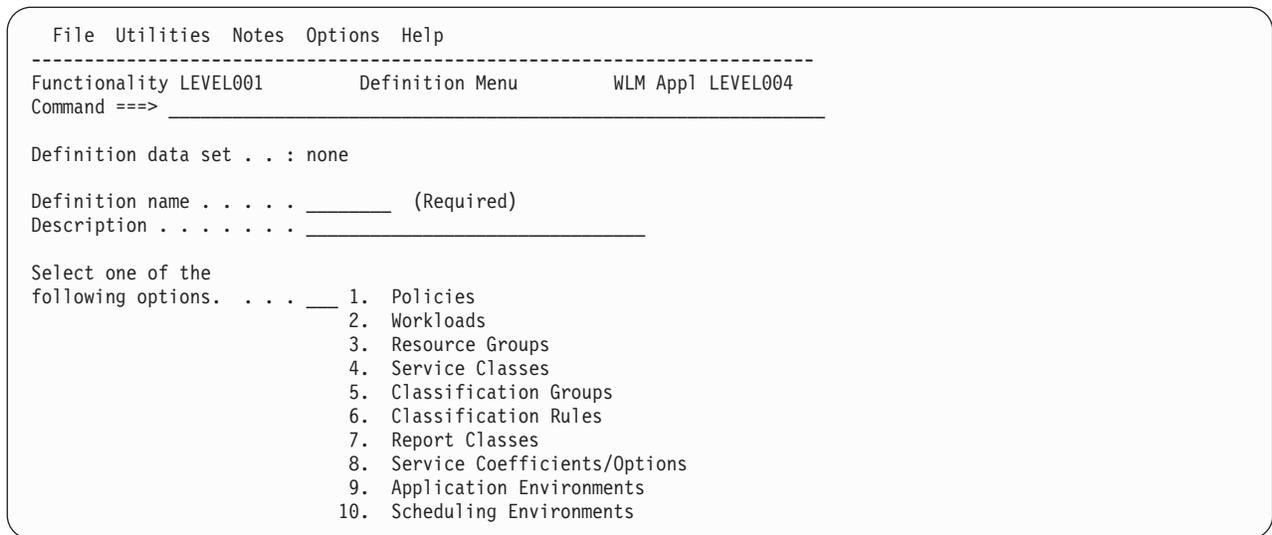


Figure 30. Definition Menu

When you define your service definition for the first time, you should define it in the following order:

**1. Policies**

A policy consists of a name, a description, and policy over-rides. The first time you set up a service definition, define a policy name and description. If you do not have a business need to change your goals, you can run with one service policy, without any policy overrides.

You use a policy override only if you have a business need to change a goal for a certain time, such as for the weekend, or for nighttime. You can define your policy overrides once you have defined you service classes.

**2. Workloads**

A workload logically consists of a group of one or more service classes. You associate a workload with a service class in the Service Class panel. Enter your workloads before creating your service classes.

**3. Resource groups (optional)**

A resource group is a minimum or maximum amount of processing capacity. You associate a resource group with a service class in the Service Class panel. Enter resource groups before creating your service classes.

**4. Service classes**

A service class is a group of work with similar performance goals, resource requirements, or business importance. You make the association with a workload and a resource group in the service class panel. You associate a service class with incoming work in the classification rules panel. Enter service classes before creating classification rules.

**Policy overrides**

Once you have created a service class, you can create a policy override. You specify the policy override by selecting Service Policies from the Definition Menu, and then specifying the action code for Override service class or Override resource group.

**5. Classification groups (optional)**

You use groups to simplify classification. You associate a classification group with a service class in the classification rules panel. If you intend to use them,

create groups before creating classification rules. See Chapter 9, “Defining Classification Rules” on page 49 for descriptions of group qualifiers.

## 6. Classification rules

Classification rules assign incoming work to service classes. Before you create your classification rules, you must understand which subsystem’s work is represented in each of your service classes.

When you choose the option Classification Rules, you go to the Subsystem Type Selection List for Rules. This selection list is primed with all of the IBM-Supplied subsystem types. They are reserved names.

## 7. Report classes (optional)

A report class is a group of work for which you want reporting data. You do not have to define report classes before assigning them to work in classification rules. You can create them from within the classification rules menu.

## 8. Service coefficients/options

Service coefficients define the weight to be applied to one type of service over another in the calculation of service rates. You can enter new values for the CPU, IOC, MSO, and SRB service coefficients.

See “Service Definition Coefficients” on page 83 for more information.

There are two additional options on this panel:

- **I/O Priority Management:** The default is **no**, meaning that I/O priorities will be the same as dispatching priorities. Specifying **yes** means I/O priorities should be managed separately from dispatching priorities, according to the goals of the work. See “Specifying I/O Priority Management” on page 84 for more information.
- **Dynamic Alias Management:** The default is **no**, meaning that dynamic alias management is disabled for the entire sysplex. Specifying **yes** will cause workload management to dynamically reassign parallel access volume aliases to help work meet its goals and to minimize IOS queueing. See “Specifying Dynamic Alias Management” on page 85 for more information.

## 9. Application Environments

An application environment is a group of application functions invoked by request and executed in server address spaces. You can have workload management start and stop these server address spaces automatically, or do this manually or through automation. You define the application environment, an optional procedure name for starting the server address spaces, and any start parameters needed for the start procedure.

## 10. Scheduling Environments

A scheduling environment is a list of resource names along with their required states. By associating incoming work with a scheduling environment, you ensure that that work is assigned to a system only if that system satisfies all of the requirements. You define the scheduling environment, listing all of the resource names and required states that are contained within. You also define the resource names themselves.

## Using the Menu Bar on the Definition Menu

The menu bar on the Definition Menu has some functions not accessible from any other panel in the application. From the menu bar, you:

- Verify a service definition
- Allocate a WLM couple data set
- Install a service definition on the WLM couple data set
- Activate a service policy

Table 15 shows the options available from the menu bar on the Definition Menu. The section below explains each of the options.

Table 15. Menu Bar Options on the Definition Menu

File	Utilities	Notes	Options	Help
New	Install definition	Edit notepad	Process ISPF list dataset	General Help
Open	Extract definition			Keys Help
Save	Activate service policy			Using Help
Save as	Allocate couple data set			Tutorial
Print	Allocate couple data set using CDS values			About...
Print as GML	Validate definition			
Cancel				
Exit				

## File

**New** Use new to define a new service definition.

**Open** Use open to read a previously defined service definition. The Read saved definition panel is displayed, where you can specify the data set name.

**Save** Use save to save the currently displayed service definition.

### Save as

Use Save as to save the currently displayed service definition in a PDS. The Save to... panel is displayed where you can specify the data set name. You do not need to preallocate the data set. If the data set does not exist, the application displays the Create Data Set? panel where you can continue with the PDS create.

**Print** Use Print to print the complete service definition to the ISPF list data set. Use the Options menu bar option to process the ISPF list dataset. This option requires no formatting step.

### Print as GML

Use Print as GML for a more readable, tabular display of service definition objects and values. This option creates a source data set with GML starter set tags imbedded. The data set must be allocated as variable block with logical record length 255. This data set can then be formatted with the SCRIPT/VS processor.

### Cancel

Use cancel to cancel any actions performed. Cancel is the same as using the cancel PF key.

**Exit** Use exit to exit from the definition menu and the application. Exit is the same as using the exit PF key.

## Utilities

### Install definition

Use this option to install the service definition onto the WLM couple data set. Installing the service definition makes any changes available for policy activation.

### Extract definition

Use this option to extract the service definition previously installed on the WLM couple data set.

### Activate service policy

Use this option to activate a policy. When you select this option, the application displays a list of the service policies defined in the service

definition currently installed on the WLM couple data set. You activate the service policy by selecting it from the list.

**Note:** If you have just made changes for a service definition, make sure you install it to have the changes take effect.

### **Allocate couple data set**

Use this option to allocate both your primary and alternate WLM couple data sets. This option is for users who are doing one of the following:

- allocating a WLM couple data set for the first time
- migrating to an OS/390 R3 or later release where the current WLM couple data set was allocated with a pre-OS/390 R3 release.

In this case you should run the IWMARSZ utility to determine the size of the current WLM couple data set. For more information on how and when to run IWMARSZ, see "Calculate the Size of the WLM Couple Data Set" on page 145.

All other users should use the option "Allocate couple data set using CDS values" described below. To make the WLM couple data set available for use in the sysplex, you must update your COUPLExx parmlib member and issue the SETXCF command.

### **Allocate couple data set using CDS values**

Use this option to allocate both your primary and alternate WLM couple data sets based on your existing WLM couple data set size. Use this option if your current WLM couple data set was allocated using OS/390 R3 or higher. The application displays the current size values on the panel.

To make the WLM couple data set available for use in the sysplex, you must update your COUPLExx parmlib member and issue the SETXCF command.

### **Validate definition**

Use this option to verify that your service definition is free from certain errors that would otherwise be flagged when you attempt to save or install the service definition.

## **Notes**

### **Edit notepad**

Use this option to create and edit a notepad. You can use the notepad to keep track of changes made to all parts of a service definition.

## **Options**

### **Process ISPF list dataset**

Use this to process the list data set if you have previously done a Print.

## **Help**

### **General Help**

Use general help for information about the panel currently displayed.

### **Keys Help**

Use keys help for information about the using PF keys.

### **Using Help**

Use this option for information about how to get help while using the WLM application.

### Tutorial

Use the tutorial option for information about how to use the panels. This options provides context-specific examples and scenarios.

### About...

This option provides information about the copyright and license.

---

## Working with Service Policies

When you choose the Policy option for the first time, the application displays the Create a Service Policy Menu. Figure 31 shows a sample panel.

```
Service-Policy  Notes  Options  Help
-----
                          Create a Service Policy

Enter or change the following information:

Service Policy Name . . . . . _____ (Required)
Description . . . . . _____

Command ==> _____
```

Figure 31. Create a Service Policy

Once you have created a service policy, any other time you choose the policy option from the definition menu, the application displays a policy selection list. From here, you can modify your policy description, print and browse your service policies, and define your service policy overrides. Figure 32 shows a Service Policy Selection List.

```
Service-Policy  View  Notes  Options  Help
-----
                          Service Policy Selection List          ROW 1 TO 3 OF 3

Action Codes: 1=Create, 2=Copy, 3=Modify, 4=Browse, 5=Print, 6=Delete,
              7=Override Service Classes, 8=Override Resource Groups,
              /=Menu Bar

Action  Name      Description                               User      Date
---    ---      ---
---    HOLIDAY    Policy for shut-Down holidays             KIRSTEN   1996/12/04
---    WEEKDAY    Policy for Mon - Friday                   KIRSTEN   1996/12/04
---    WEEKEND    Policy for Fri - Sun                      KIRSTEN   1996/12/04
***** BOTTOM OF DATA *****

Command ==> _____
```

Figure 32. Service Policy Selection List

## Working with Workloads

When you choose the Workload option for the first time, the application displays the Create a Workload panel. Figure 33 shows a Create a Workload panel.

```
Workload  Notes  Options  Help
-----
                          Create a Workload

Enter or change the following information:

Workload Name . . . . . _____ (Required)
Description . . . . . _____

Command ==> _____
```

Figure 33. Create a Workload

You associate the workload with a service class in the service class panel.

Once you have created a workload, any other time you choose the workload option from the definition menu, the application displays a workload selection list. The Workload Selection List is similar to the Policy Selection List. From here, you can modify your workload description, print, and browse your workloads. Figure 34 shows a Workload Selection List.

```
Workload  View  Notes  Options  Help
-----
                          Workload Selection List          ROW 1 TO 9 OF 9

Action Codes: 1=Create, 2=Copy, 3=Modify, 4=Browse, 5=Print, 6=Delete,
              /=Menu Bar

Action  Name      Description          ---Last Change---
      User      Date
-----
  ---  APPC          All batch           KIRSTEN  1996/10/07
  ---  BATCH         All batch           KIRSTEN  1996/10/07
  ---  APIMS         IMS application     KIRSTEN  1996/10/07
  ---  PRDCICS       CICS Production     ADRIAN   1996/10/17
  ---  PRDIMS       IMS Production      ADRIAN   1996/10/17
  ---  STC          Started tasks       KIRSTEN  1996/10/07
  ---  TRNCICS      CICS training       KIRSTEN  1996/10/07
  ---  TRNIMS       IMS training        GEORGEN  1996/10/07
  ---  TSO         All TSO             GEORGEN  1996/10/07
***** BOTTOM OF DATA *****

Command ==> _____
```

Figure 34. Workload Selection List

---

## Working with Resource Groups

To define a resource group, you choose option 3 on the Definition menu. You define a name, a description (optional), a minimum and/or maximum capacity. You associate the resource group with a service class in the service class panel. Figure 35 shows a Create a Resource Group panel.

```
Resource-Group  Notes  Options  Help
-----
                          Create a Resource Group

Enter or change the following information:

Resource Group Name . . . . . _____ (required)
Description . . . . . _____

Minimum Capacity . . . . . _____
Maximum Capacity . . . . . _____

Command ==> _____
```

Figure 35. Create a Resource Group

As with a workload, once you have created a resource group, any other time you choose the resource group option from the definition menu, the application displays a selection list. From here, you can modify your resource group description, as well as print, and browse it.

---

## Working with Service Classes

Once you have defined your workloads and resource groups, you can define your service classes. Choose the Service Class option on the definition menu. Figure 36 on page 184 shows a Create a Service Class panel. You must assign a workload to the service class in the Workload field.

```

Service-Class  Notes  Options  Help
-----
                          Create a Service Class                          ROW 1 TO 1 OF 1

Service Class Name . . . . . _____ (Required)
Description . . . . . _____
Workload Name . . . . . _____ (name or ?)
Base Resource Group . . . . . _____ (name or ?)
Cpu Critical . . . . . NO _____ (YES or NO)

Specify BASE GOAL information. Action Codes: I=Insert new period,
E=Edit period, D=Delete period.

      ---Period---  -----Goal-----
Action # Duration  Imp.  Description
  _i
***** BOTTOM OF DATA *****

Command ==> _____

```

Figure 36. Create a Service Class

Use the “Cpu Critical” field to specify CPU protection for critical regions.

**Important**

The use of these options limits WLM’s ability to manage the system. This may affect system performance and/or reduce the system’s overall throughput.

### Defining Goals

To enter the goal information, enter an **i** in the Action field, as shown in Figure 36. The goal selection pop-up is displayed, as shown in Figure 37. From this panel, select the type of goal you want to assign to the service class.

```

Choose a goal type for period 1

1_ 1. Average response time
   2. Response time with percentile
   3. Execution velocity
   4. Discretionary

```

Figure 37. Choose a Goal Type Pop-Up

When you choose option 1, average response time, the application displays the average response time goal pop-up. There is a different pop-up for each goal type where you can fill in the information for the goal. If you are defining a single period goal, then you should not fill in a duration. If you are defining multiple periods then you must fill in a duration. Figure 38 on page 185 shows an average response time goal pop-up.

```

Average response time goal
Enter a response time of up to 24 hours for period 1

Hours . . . . . _ (0-24)
Minutes . . . . . _ (0-99)
Seconds . . . . . _5_ (0-9999)

Importance . . . 2 (1=highest, 5=lowest)
Duration . . . _____ (1-999,999,999, or
                           none for last period)

```

Figure 38. Average Response Time Goal Pop-Up

When you press EXIT, you return to the create a service class panel with the goal information filled in, as shown in Figure 39.

```

Service-Class  Notes  Options  Help
-----
                          Create a Service Class                      ROW 1 TO 2 OF 2

Service Class Name . . . . . QUACK      (Required)
Description . . . . . _____
Workload Name . . . . . APPC          (name or ?)
Base Resource Group . . . . . _____ (name or ?)
Cpu Critical . . . . . NO             (YES or NO)

Specify BASE GOAL information. Action Codes: I=Insert new period,
E=Edit period, D=Delete period.

      ---Period---  -----Goal-----
Action # Duration  Imp.  Description
-----
_      1           2    Average response time of 00:00:05.000
***** BOTTOM OF DATA *****

```

Figure 39. Create a Service Class Panel

## Using Action Codes on Service Class Panels

You use action codes on this panel to define and edit goals. Figure 40 shows the edit codes available for service class.

```

Specify BASE GOAL information. Action Codes: I=Insert new period,
E=Edit period, D=Delete period.

      ---Period---  -----Goal-----
Action # Duration  Imp.  Description
-----
_      1           5    80% complete within 00:30:00.000
***** BOTTOM OF DATA *****

```

Figure 40. Action Codes for Goal

### I=Insert new period

Use I to define a new period. The application adds a line below. If you have multiple periods, then a duration is required on the previous period. Use action code E to edit the previous period. You'll go through the windows with the goals information filled in, and you can add a duration.

### E=Edit period

Use E to edit a period.

### D=Delete period

Use D to delete a period. If you have defined multiple periods for a service class, remember that you do not define a duration for the last period.

---

## Defining Service Policy Overrides

You define all service policy overrides from the service policy selection list. You can define three kinds of service policy overrides:

- Override service class goals
- Override resource group assignment
- Override resource group attributes.

To override a service class goal, choose either the action code or the menu bar option to Override Service Classes. Figure 41 shows a service policy selection list where you have chose to override the service classes for the weekend policy.

```
Service-Policy View Notes Options Help
-----
Service Policy Selection List          ROW 1 TO 3 OF 3

Action Codes: 1=Create, 2=Copy, 3=Modify, 4=Browse, 5=Print, 6=Delete,
              7=Override Service Classes, 8=Override Resource Groups,
              /=Menu Bar

Action  Name      Description              ---Last Change---
-----  -
_       HOLIDAY  Policy for shut-Down holidays  KIRSTEN  1996/12/04
_       WEEKDAY  Policy for Mon - Friday        KIRSTEN  1996/12/04
_7      WEEKEND  Policy for Fri - Sun           KIRSTEN  1996/12/04
***** BOTTOM OF DATA *****

Command ==> _____
```

Figure 41. Service Policy Selection List

The application displays the Override Service Class Selection List. This is the list of all of your defined service classes, similar to the Service Class Selection list, except it has an extra field called *Overridden Goal* which indicates whether or not the service class goal or resource group assignment was overridden for that policy. Figure 42 on page 187 shows an Override Service Class Selection List. Select a service class whose goal or resource group assignment you want to change and specify the action code or menu bar option to override the service class.

```

Override View Notes Options Help
-----
IWMAP4A          Override Service Class Selection List          ROW 1 TO 8 OF 14

Service Policy Name . . . . : BACKUPS

Action Codes: 3=Override Service Class, 4=Browse, 5=Print,
              6=Restore Base attributes, /=Menu Bar

Action  Service  Overridden
      Class  Goal      Description
---    -
---    BATCHJ  NO        Class J work
---    BATCHTST NO        Test
---    BATCHX  NO        Class X work
---    CICSFAST NO       Fast CICS work
---    CICSSLOW NO      Slow CICS work
3_    HOTBATCH NO       Hot batch
---    IMSNRESP NO      IMS non response
---    IMSRESP NO       IMS response

Command ==> _____

```

Figure 42. Override Service Class Selection List

When you choose the override service class option, the application displays the Override Attributes for a Service Class panel. Figure 43 shows an Override Attributes for a Service Class panel. To override the goal, use the same codes to edit the goal as you do on the Create or Modify a Service Class panel. You can also change the “Cpu Critical” setting, just as you do on the Create or Modify a Service Class panel. To change the resource group assignment of the service class, either enter in a resource group name, or put a ? in the resource group field to select a resource group from the selection list.

If you want to remove a service from a resource group, blank out the name.

```

Service-Class Xref Notes Options Help
-----
                Override attributes for a Service Class          ROW 1 TO 2 OF 2

Service Policy Name . . . . : STANDARD
Service Class Name . . . . : BAT_0

Override the following information:

Resource Group . . . . . _____ (name or ?)
Cpu Critical . . . . . NO          (YES or NO)

Action Codes: I=Insert new period, E=Edit period, D=Delete period.

      ---Period---  -----Goal-----
Action # Duration  Imp. Description
---    -
---    1           2   Execution velocity of 70
***** BOTTOM OF DATA *****

Command ==> _____

```

Figure 43. Override Attributes for a Service Class

Once you have edited your goal or changed the resource group assignment, press EXIT, and you return to the Override Service Class Selection List. The *Overridden Goal* field for that service class now says YES.

## Working with Classification Rules

Classification of work depends on having the rules defined for the correct subsystem type. When you choose the *Classification Rules* option from the *Definition Menu*, you go to the *Subsystem Type Selection List for Rules* panel. This panel initially contains the reserved names of the IBM-supplied subsystem types.

Although you may want to change the description of the subsystem types, you should not delete any of the entries provided by IBM unless your installation does not plan to ever use them. If your installation later does need to use them, they can be manually added at that time.

Figure 44 shows the subsystem type selection panel.

```
Subsystem-Type View Notes Options Help
-----
IWMAP7A          Subsystem Type Selection List for Rules          ROW 1 TO 12 OF 12

Action Codes: 1=Create, 2=Copy, 3=Modify, 4=Browse, 5=Print, 6=Delete,
              /=Menu Bar

Action  Type      Description                      -----Class-----
      ASCH      Use Modify to enter YOUR rules    Service Report
      CICS      CICS interactive
      DB2       Use Modify to enter YOUR rules
      DDF       Use Modify to enter YOUR rules
      _3      IMS       Use Modify to enter YOUR rules
      IWEB      Use Modify to enter YOUR rules
      JES       batch
      LSFM      Use Modify to enter YOUR rules
      MQ        Use Modify to enter YOUR rules
      OMVS      Use Modify to enter YOUR rules
      SOM       Use Modify to enter YOUR rules
      STC       Use Modify to enter YOUR rules
      TSO       TSO
***** Bottom of data *****

Command ==> _____
```

Figure 44. Subsystem Type Selection List for Rules

To create your rules, use the Modify option (3) to create rules for the IBM-supplied subsystem types. For example, use Modify on subsystem type IMS as show in Figure 44 to create the rules for your IMS work. Figure 45 on page 189 shows the *Modify Rules for the Subsystem Type* panel.

```

Subsystem-Type Xref Notes Options Help
-----
Modify Rules for the Subsystem Type          Row 1 to 2 of
Command ==> _____ SCROLL ==> PAG

Subsystem Type . : CICS          Fold qualifier names?  Y (Y or N)
Description . . . IBM-defined subsystem type

Action codes:  A=After    C=Copy      M=Move    I=Insert rule
                B=Before  D=Delete row R=Repeat  IS=Insert Sub-rule
                                           More==>

Action  -----Qualifier-----
-----Type      Name      Start
_____ 1 SI      IMSTRN*  _____
_____ 2 TC      15      _____
_____ 1 SI      IMSMDL*  _____
_____ 2 TC      15      _____
_____ 1 TC      15      _____

                DEFAULTS:  Service  Report
                PRDIMSR  _____
                TRNIMSR  _____
                TRNIMSNR _____
                MDLIMSR  _____
                MDLIMSNR _____
                PRDIMSNR _____

Command ==> _____

```

Figure 45. Modify Rules for the Subsystem Type

Note that the “Fold qualifier names” option, set to the default Y, means that the qualifier names will be folded to upper case as soon as you type them in and then press Enter. If you set this option to N, then the qualifier names will remain in the case they are typed in. Leave this option set to Y unless you know that you need mixed case qualifier names in your classification rules.

While creating rules, you can scroll right (PF11) to fill in the description fields. Figure 46 shows the *Modify Rules for the Subsystem Type* panel, scrolled right to the description fields.

```

Subsystem-Type Xref Notes Options Help
-----
Modify Rules for the Subsystem Type          Row 1 to 2 of
Command ==> _____ SCROLL ==> PAG

Subsystem Type . : CICS          Fold qualifier names?  Y (Y or N)
Description . . . IBM-defined subsystem type

Action codes:  A=After    C=Copy      M=Move    I=Insert rule
                B=Before  D=Delete row R=Repeat  IS=Insert Sub-rule
                                           <=== More ===>

Action  -----Qualifier-----
-----Type      Name      Start
_____ 1 SI      IMSTRN*  _____
_____ 2 TC      15      _____
_____ 1 SI      IMSDL*  _____
_____ 2 TC      15      _____
_____ 1 TC      15      _____

                -----Description-----
                (Descriptions here... _____
                Any valid characters, _____
                up to 32 in length.) _____

Command ==> _____

```

Figure 46. Modify Rules for the Subsystem Type, Scrolled Right to Description Fields

For ASCH, CICS, IMS, JES, OMVS, STC, and TSO work, you can scroll right (PF11) yet again to fill in the “Storage Critical” field. For JES and STC work only, you can also fill in the “Manage Regions Using Goals Of:” field, with a value of either **TRANSACTION**, in which case the region will be managed to the transaction

response time goals, or **REGION**, in which case the region is managed to the goal of the service class assigned to the region. Note that for ASCH, CICS, IMS, and OMVS work, the “Manage Regions Using Goals Of:” field will read N/A. See Chapter 12, “Defining Special Protection Options for Critical Work” on page 89 for more information on using these fields.

**Important**

The use of these options limits WLM’s ability to manage the system. This may affect system performance and/or reduce the system’s overall throughput.

Figure 47 shows the *Modify Rules for the Subsystem Type* panel, scrolled right to the “Storage Critical” and “Manage Regions Using Goals Of:” fields.

```

Subsystem-Type Xref Notes Options Help
-----
Command ==> Modify Rules for the Subsystem Type Row 1 to 2 of
SCROLL ==> PAG

Subsystem Type . : STC Fold qualifier names? Y (Y or N)
Description . . . IBM-defined subsystem type

Action codes: A=After C=Copy M=Move I=Insert rule
               B=Before D=Delete row R=Repeat IS=Insert Sub-rule
               <=== More

Action  Type  Name  Start  Storage  Manage Region
-----  -  -  -  -  -  -
   1  TN  COMBL*  ___  NO  TRANSACTION
   2  UI  COMBLD  ___  NO  TRANSACTION
   2  UI  COMFTP  ___  YES  TRANSACTION
   1  UI  SRVLIB  ___  YES  TRANSACTION
   1  PF  99  ___  YES  REGION
   1  TC  B  ___  NO  REGION
   1  TC  C  ___  NO  REGION

Command ==>

```

Figure 47. *Modify Rules for the Subsystem Type, Scrolled Right to Storage Critical and Region Management Fields*

Keep your classification rules as simple as possible. Use subsystem defaults to cover most cases, and list the exceptions. Remember that the rules are order sensitive, and the first matched rule applies. The Start field applies only to the accounting information and subsystem parameter qualifiers.

If you have a non-IBM supplied subsystem type, you can use Create to create your subsystem type. You can also modify the qualifier selection list to include only the qualifier types that apply to your subsystem type.

For more information about defining the order of rules, and how to use the Start field, see Chapter 9, “Defining Classification Rules” on page 49.

## Using Action Codes on the Modify Rules Panel

You use action codes on this panel to create and change your classification rules.

```

Subsystem-Type Xref Notes Options Help
-----
Command ==> Modify Rules for the Subsystem Type Row 1 to 2 of 2
SCROLL ==> PAGE

Subsystem Type . : CICS Fold qualifier names? Y (Y or N)
Description . . . IBM-defined subsystem type

Action codes: A=After C=Copy M=Move I=Insert rule
              B=Before D=Delete row R=Repeat IS=Insert Sub-rule

```

Figure 48. Action Codes for Classification Rules

In the Qualifier Type field, enter the type of qualifier that you are defining. Qualifier type can be one of the following:

- AI Accounting Information
- CI Correlation Information
- CN Collection Name
- CT Connection Type
- CTG Connection Type Group
- LU LU Name
- LUG LU Name Group
- NET Net ID
- NETG Net ID Group
- PC Process Name
- PF Perform
- PFG Perform Group
- PK Package Name
- PKG Package Name Group
- PN Plan Name
- PNG Plan Name Group
- PR Procedure Name
- PRI Priority
- PX Sysplex Name
- SE Scheduling Environment Name
- SI Subsystem Instance
- SIG Subsystem Instance Group
- SPM Subsystem Parameter
- SSC Subsystem Collection Name
- SY System Name
- SYG System Name Group
- TC Transaction Class
- TCG Transaction Class Group
- TN Transaction Name
- TNG Transaction Name Group
- UI Userid
- UIG Userid Group

You can either enter in the qualifier type, or select it from a selection list. To select it from a list, enter a ? in the Qualifier Type field. "Using Selection Lists for Classification Rules" on page 192 explains more about using selection lists.

Then type in the name of the qualifier in the Name field, and the service class in the Service Class field.

You can use the action codes on this panel to enter in additional classification rules. The action codes are:

**I=Insert rule**

To create a rule at the same level as the rule you type the action code next to.

**IS=Insert Sub-rule**

To create a next level rule or a nest. IS=Insert Sub-rule specifies a rule under the previous rule; that is, a rule at the next indented level.

**R=Repeat**

To copy a rule and its sub-rules, if any. The copied rule is placed directly below the one you are copying. If a rule has sub-rules; that is, it is a family, the entire family is copied.

**D=Delete row**

To delete a rule.

**A=After**

Use with C=Copy and M=Move to specify that the copied/moved rule is to go after this rule.

**B=Before**

Use with C=Copy and M=Move to specify that the copied/moved rule is to go before this rule.

**C=Copy**

To copy a rule and its sub-rules, if any. Use B=Before or A=After with C=Copy to specify where to copy the rule. The rule is copied at the same level as the rule you are placing it Before or After. For example, a level 3 rule becomes a level 2 rule if you copy it before or after a level 2 rule.

When a rule changes levels as a result of a Copy, the levels of any sub-rules that were copied with it are bumped up or down accordingly. The rule family is copied at the same level as the Before or After rule.

**M=Move**

To move a rule and its sub-rules, if any. Use B=Before or A=After with M=Move to specify where to move the rule. The rule is moved according to the way you do a copy, by placing it Before or After.

## Using Selection Lists for Classification Rules

For a list of applicable qualifiers for the subsystem type, put a ? in the qualifier type field. You can select a qualifier from the list by putting a / next to the one you want to use.

For a list of the defined service classes put a ? in the service class field. You can select a service class from the list by typing a / next to the one you want to use.

You can define a report class by typing the report class name in the report class field. A pop-up panel lets you confirm the create. For a list of the defined report classes, put a ? in the report class field.

## Creating a Subsystem Type for Rules

Use the Create option only if your installation has its own subsystem type, or a vendor subsystem type that supports workload management. You must check that product's documentation for its reserved subsystem type name.

## Deleting a Subsystem Type for Rules

Use the delete option on the *Subsystem Type Selection List for Rules* only to remove an IBM-supplied subsystem type that your installation does not have, or does not plan on using. The application displays a pop-up for confirmation of the delete.

## Working with Classification Groups

If you have a long list of work that you want to use in a classification rule, you can create a group. You can create groups for the following qualifier types:

- Connection Type
- LU Name
- Net ID
- Package Name
- Plan Name
- Perform
- Subsystem Instance
- System Name
- Transaction Class
- Transaction Name
- Userid Groups

For example, you may want to create a transaction name group to use in classification rules for your started tasks. Figure 49 shows a *Create a Group* panel. You can use wild-card and masking notation in qualifier groups. Note that there is also room next to each classification group for a description.

```
Group Xref Notes Options Help
-----
                          Create a Group                      ROW 1 TO 10 OF 17

Enter or change the following information:

Qualifier type . . . . . : Transaction Name
Group name . . . . . : STC_GR1
Description . . . . . : Low Priority STC
Fold qualifier names? . . . . Y (Y or N)

Qualifier Name  Description
UCC7           (Descriptions here... _____)
WSF2*         Any valid characters, _____
PHOENIX       up to 32 in length.) _____
NVDM          _____
BMC*          _____
DBUS*         _____
DFHSM*        _____
EMAIL         _____
NETEX        _____
```

Figure 49. Create a Group

Then, you go to modify the STC subsystem on the *Subsystem Type Selection List for Rules* panel, and reference the group in the rules. Choose transaction name group for the qualifier type, and enter the name of the group. You can also type a ? in the qualifier name field for a list of the defined groups.

Figure 50 shows a *Modify Rules for STC Subsystem* panel. On this panel, you specify TNG in the qualifier type field. Then you can enter the transaction group name STC\_GR1 in the qualifier name field.

```

Subsystem-Type Xref Notes Options Help
-----
                Modify Rules for the Subsystem Type          Row 1 to 7 of 8

Subsystem Type . . . . . : STC
Description . . . . . IBM-defined subsystem type
Fold qualifier names? . . . . Y (Y or N)

Enter one or more action codes: A=After B=Before C=Copy D=Delete
M=Move I=Insert rule IS=Insert Sub-rule R=Repeat

Action      -----Qualifier-----          -----Class-----
Action  Type      Name      Start          Service      Report
-----  -
      1  TNG      STC_GR1  _____  STC_5        _____
      _____  STC_1        _____

                DEFAULTS:

```

Figure 50. Modify Rules for STC Subsystem

## Working with Report Classes

- You can create a report class by either:
- Selecting the report classes option on the Definition Menu
  - Typing in a name on the report class field of the *Modify Rules for the Subsystem Type* panel.

When you select the report classes option on the Definition menu, you go to the Create a Report Class panel. From there, you enter the report class name and description. The next time you choose the option from the Definition Menu, the application displays Report Class selection list.

To create a report class from the Modify Rules for Subsystem panel, enter a name in the report class field. Figure 51 shows a Modify Rules for the CICS Subsystem panel. If you type in a name, for example CICSHOT, in the report class field, the application displays a Create Report Class pop-up panel. Figure 52 on page 195 shows a Create a Report Class pop-up.

```

Subsystem-Type Xref Notes Options Help
-----
                Modify Rules for the Subsystem Type          Row 1 to 6 of
Command ==> _____ SCROLL ==> PAG
Subsystem Type . : CICS      Fold qualifier names? Y (Y or N)
Description . . . IBM-defined USAA modified

Action codes: A=After      C=Copy      M=Move      I=Insert rule
              B=Before      D=Delete row R=Repeat  IS=Insert Sub-rule

Action      -----Qualifier-----          -----Class-----
Action  Type      Name      Start          Service      Report
-----  -
      1  TN      MSA*      _____  PRDCICS      CICSHOT_
      1  SI      CICSTRN*  _____  PRDCICSM      _____
      _____  TRNCICS      _____

***** BOTTOM OF DATA *****

```

Figure 51. Creating a Report Class on the Modify Rules Panel

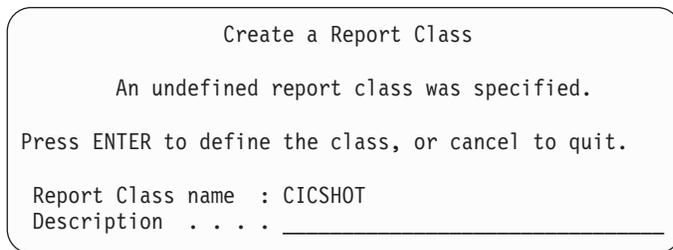


Figure 52. Create a Report Class Confirmation Panel

When you press ENTER, the report class is created, and is added to the report class selection list. You can type a ? in the report class field on the Modify Rules for a Subsystem Type panel for a selection list of report classes.

---

## Working with Service Coefficients and Options

On the Service Coefficient/Service Definition Option panel, you can specify the service definition coefficients, whether you want workload management to manage your I/O priorities, and whether you want workload management to manage your parallel access volume alias addresses.

Chapter 11, “Defining Service Coefficients and Options” on page 83 provides some advice on how and when to adjust your coefficients.

To find out what service definition coefficients you are currently running with, check your RMF Monitor I Workload Activity Report. This report lists the service definition coefficients your installation is currently using.

If you want workload management to manage I/O priorities for you, specify **YES** for *I/O priority management*. When you specify **YES**, workload management manages your I/O priorities and includes I/O usings and delays in its execution velocity calculation. If you want I/O priorities to be managed the same as in pre-OS/390 R3 releases, specify **NO** for *I/O priority management* which is the default. When you specify **NO**, workload management sets the I/O priority to be the same as the dispatching priority, and I/O usings and delays are not included in the execution velocity calculation.

If you want workload management to manage parallel access volume alias addresses for you, specify **YES** for *Dynamic alias management*. When you specify **YES**, workload management dynamically reassigns alias addresses from one base to another to help work meet its goals and to minimize IOS queueing. See “Specifying Dynamic Alias Management” on page 85 for more information on this global setting and its relationship to the **WLMPAV=** setting on each individual device. When you specify **NO**, which is the default, dynamic alias management is globally disabled for the entire sysplex. Systems will still use the aliases assigned to the base devices, but there will be no automatic reassignment of aliases based on goals or queueing.

Figure 53 on page 196 shows the Service Coefficient/Service Definition Options panel.

Coefficients/Options	Notes	Options	Help
-----			
Service Coefficient/Service Definition Options			
Enter or change the Service Coefficients:			
CPU	.....	_____	(0.0-99.9)
IOC	.....	_____	(0.0-99.9)
MSO	.....	_____	(0.0000-99.9999)
SRB	.....	_____	(0.0-99.9)
Enter or change the service definition options:			
I/O priority management	.....	NO	(Yes or No)
Dynamic alias management	.....	NO	(Yes or No)
Command ==> _____			

Figure 53. Service Coefficients Panel

## Working with Application Environments

Define an application environment so that workload management can assist in managing work that runs in server address spaces. You can specify that workload management should dynamically start and stop server address spaces to process work running in the application environment. As of OS/390 R3 and later, you can define application environments for the DB2, SOM, and IWEB IBM-supplied subsystem types. Refer to subsystem reference information for guidance on how to specify application environments for the subsystem.

Figure 54 shows the Application Environments panel.

Application-Environment	Notes	Options	Help
-----			
Create an Application Environment			
Command ==> _____			
Application Environment	.....	_____	Required
Description	.....	_____	
Subsystem Type	.....	_____	Required
Procedure Name	.....	_____	
Start Parameters	.....	_____	
Limit on starting server address spaces for a subsystem instance:			
1	1.	No limit	
	2.	Single address space per system	
	3.	Single address space per sysplex	

Figure 54. Application Environments

## Working with Scheduling Environments

Use the scheduling environment panels to define scheduling environments and their lists of resource names and required states, and to define the individual resource names themselves.

## Creating a New Scheduling Environment

If you have not yet defined any scheduling environments, Figure 55 shows the first panel you will see.

```
IWMAPA7                Decide what to create?
Command ==> _____

No scheduling environments exist.  Would you like to create a
scheduling environment or list resources for scheduling
environments?

Select one of the following options.
_  1.  Create Scheduling Environment
_  2.  Create Resource(s)
```

Figure 55. *Decide What to Create Panel*

From here you can either enter a **1** to create a scheduling environment, or a **2** to create resources. (You might want to create all of your resources, for instance, before creating your first scheduling environment. This is covered in “Creating a New Resource Directly from the Main Panel” on page 201.)

**Note:** Until you have created at least one scheduling environment (no matter how many resources you create), you will always see this initial panel.

Figure 56 shows the main Scheduling Environment Selection List panel. Assuming that you have at least one scheduling environment already created, this is the first panel you see when you work with scheduling environments.

```
Scheduling-Environments  Notes  Options  Resources  Help
-----
IWMAPAA                Scheduling Environment Selection List          Row 1 to 2 of 2
Command ==> _____

Action Codes: 1=Create, 2=Copy, 3=Modify, 4=Browse, 5=Print, 6=Delete,
              /=Menu Bar

Action  Scheduling Environment Name      Description
  1_    IMSPRIME                             Prime IMS Processing
  _     CICS_SPEC                             Special CICS Processing
```

Figure 56. *Scheduling Environment Selection List Panel*

By entering a **1** on any Action line, you will go to the Create a Scheduling Environment panel, as shown in Figure 57 on page 198. (You can also get there by putting the cursor under the Scheduling-Environment field at the top of the screen, and then entering a **1** in the pop-up box.)

```

Scheduling-Environments  Notes  Options  Help
-----
IWMAPAD                Create a Scheduling Environment          Row 1 to 1 of 1
Command ==> _____

Scheduling Environment Name  DB2LATE          Required
Description . . . . . Offshift DB2 Processing_____

Action Codes: A=Add D=Delete

Action  Resource Name      Required
a_      _____          State      Resource Description

```

Figure 57. Create a Scheduling Environment Panel

Enter the scheduling environment name and description, and then enter an **A** on the Action line to go to the next panel, Resource Definition List, as seen in Figure 58.

```

Resources  Notes  Options  XREF  Help
-----
IWMAPAS                Resource Definition List          Row 1 to 4 of 4
Command ==> _____

Selection For Scheduling Environment DB2LATE

Action Codes: A=Add S=Select X=XREF /=Menu Bar

Action  Resource Name      In Use  Resource Description
_       CICS001              YES     CICS Subsystem
S_      DB2A                  YES     DB2 Subsystem
_       IMS1                  YES     IMS Subsystem
S_      PRIMETIME            YES     Peak Business Hours

```

Figure 58. Resource Definition List Panel

This panel shows *all* of the resource definitions in the service definition. The **YES** in the In Use field simply means that the resource name is currently a part of some scheduling environment. (It does *not* mean that it is in use by the scheduling environment you are defining when you get to this screen, in this case DB2LATE.)

You can select as many resources as you like. In this case, we'll select DB2A and PRIMETIME. Both have already been defined (although only PRIMETIME is currently being used by another scheduling environment). To make the resource name part of the scheduling environment, put an **S** on the Action line next to it. (Alternately, you can put a **/** next to the resource names and then put the cursor under the Resources field at the top of the screen and hit ENTER. Then select 2 from the options in the pop-up box.)

If you have many resources, you can use the LOCATE primary command to scroll the display to a particular resource name.

**Note:** You can create new resources from this panel, as well. See "Creating a New Resource While Working with a Scheduling Environment" on page 203. This will take you back to the Create a Scheduling Environment panel, as shown in Figure 59 on page 199.

```

Scheduling-Environments  Notes  Options  Help
-----
IWMAPAD                Create a Scheduling Environment          Row 1 to 2 of 2
Command ==>> _____

Scheduling Environment Name  DB2LATE_____ Required
Description . . . . . Offshift DB2 Processing_____

Action Codes: A=Add D=Delete

Action  Resource Name      Required
-----  -
   DB2A          ON_____ DB2 Subsystem
   PRIMETIME     OFF_____ Peak Business Hours

```

Figure 59. Create a Scheduling Environment Panel

The DB2A and PRIMETIME names and descriptions are automatically brought back to this panel. Now you must specify either **ON** or **OFF** under the Required State field. In this example, we have chosen ON for DB2A and OFF for PRIMETIME (meaning that the DB2A resource state must be set to ON and the PRIMETIME resource state must be set to OFF for this scheduling environment to be satisfied).

### Modifying a Scheduling Environment

Once a scheduling environment is created, you can modify it at any time, either changing the required states of resources, adding new resources to the scheduling environment, or deleting resources from the scheduling environment.

Now that DB2LATE has been created in the previous section, here’s how we would modify it. From the main Scheduling Environment Selection List panel, enter a **3** in the Action line next to the scheduling environment you want to change, as shown in Figure 60. (Alternately, you can enter a / next to it, put your cursor under the Scheduling-Environments field at the top of the screen, and then select 3 from the pop-up box.)

If you have many scheduling environments, you can use the LOCATE primary command to scroll the display to a particular name.

```

Scheduling-Environments  Notes  Options  Resources  Help
-----
IWMAPAA                Scheduling Environment Selection List      Row 1 to 3 of 3
Command ==>> _____

Action Codes: 1=Create, 2=Copy, 3=Modify, 4=Browse, 5=Print, 6=Delete,
              /=Menu Bar

Action  Scheduling Environment Name      Description
-----  -
   3_   DB2LATE                          Offshift DB2 Processing
   ___  IMSPRIME                         Prime IMS Processing
   ___  CICS_SPEC                       Special CICS Processing

```

Figure 60. Scheduling Environment Selection List Panel

This will take you to the Modify a Scheduling Environment panel. Here, you can add and delete resources, and change required states. In Figure 61 on page 200, we could delete DB2A, and change the required state for PRIMETIME from OFF to ON. (To add new resources from this panel, see “Creating a New Resource While Working with a Scheduling Environment” on page 203.) When you press enter,

you'll get a confirmation message, telling you to press EXIT to save your changes, or CANCEL to discard them.

```

Scheduling-Environments  Notes  Options  Help
-----
IWMAPAJ                Modify a Scheduling Environment          Row 1 to 2 of 2
Command ==>>> _____

Scheduling Environment Name  DB2LATE_____ Required
Description . . . . . Offshift DB2 Processing_____

Action Codes: A=Add D=Delete

Action  Resource Name      Required
State   Resource Description
d_      DB2A                    ON_____ DB2 Subsystem
_       PRIMETIME             ON_____ Peak Business Hours

```

Figure 61. Modify a Scheduling Environment Panel

**Important!** When you “delete” a resource name from this panel, you are simply deleting it *from the scheduling environment*. The resource name will still exist on the Resource Definitions List panel. To delete the resource definition itself, see “Deleting a Resource” on page 206.

## Copying a Scheduling Environment

To copy a scheduling environment, from the main Scheduling Environment Selection List panel (as shown in Figure 60 on page 199) enter a **2** on the Action line next to the name. This will bring you to the Copy a Scheduling Environment panel, as shown in Figure 62.

```

IWMAPAC                Copy A Scheduling Environment          Row 1 to 2 of 2
Command ==>>> _____

Scheduling Environment Name  NEWSCHENV_____ Required
Description . . . . . _____

Action Codes: A=Add D=Delete

Action  Resource Name      Required
State   Resource Description
_       DB2A                    ON_____ DB2 Subsystem
_       PRIMETIME             OFF_____ Peak Business Hours

```

Figure 62. Copy a Scheduling Environment Panel

You must give the new scheduling environment a new, unique name. From this screen, you can modify the new scheduling environment, as shown in “Modifying a Scheduling Environment” on page 199.

The copy function is useful when you are working with scheduling environments with large numbers of resource names. If you wish to create a new scheduling environment that is similar to an existing one, you can simply copy the original and then make whatever changes are necessary.

## Browsing a Scheduling Environment

To browse a scheduling environment, from the main Scheduling Environment Selection List panel (as shown in Figure 60 on page 199) enter a **4** on the Action line next to the name. You'll see a browse screen as shown in Figure 63 on page 201.

```

IWMAPUB                               Line 00000000 Col 001 072
Command ==>                            SCROLL ==> PAGE
***** Top of Data *****

Scheduling Environment Name. . DB2LATE
Description. . . . . Offshift DB2 Processing

Resource Name      State  Resource Description
-----
DB2A               ON    DB2 Subsystem
PRIMETIME         OFF   Peak Business Hours
***** Bottom of Data *****

```

Figure 63. Browse a Scheduling Environment Panel

## Printing a Scheduling Environment

To print a scheduling environment, from the main Scheduling Environment Selection List panel (as shown in Figure 60 on page 199) enter a **5** on the Action line next to the name. The output will be written to your ISPF list dataset.

## Deleting a Scheduling Environment

To delete a scheduling environment, from the main Scheduling Environment Selection List panel (as shown in Figure 60 on page 199) enter a **6** on the Action line next to the name. This will bring you to a confirmation screen, as shown in Figure 64.

```

-----
IWMAPAX                Delete A Scheduling Environment
Command ==> _____

Scheduling Environment Name : DB2LATE
Description . . . . . : Offshift DB2 Processing
-----

Confirm the deletion request for the above scheduling environment.

Response . . . . . yes (Yes or No)

```

Figure 64. Delete a Scheduling Environment Panel

Note that all of the resource definitions that were part of this scheduling environment *still exist*. They just aren't members of this particular scheduling environment anymore.

## Creating a New Resource

In the example shown in Figure 58 on page 198, as we were creating a new scheduling environment (DB2LATE), we chose resource names (DB2A and PRIMETIME) that already existed. What if we had needed to add a new resource name called VECTOR? As you saw on that Resource Definition List panel, VECTOR did not yet exist.

There are two ways to add a new resource definition, as follows:

### Creating a New Resource Directly from the Main Panel

From the main Scheduling Environment Selection List panel, put the cursor under the Resources field at the top of the screen and hit ENTER. When the pop-up box

appears, enter a 1, as shown in Figure 65. (In this pop-up, 1 is the only option.)

```

Scheduling-Environments  Notes  Options  Resources  Help
-----
Scheduling Environment * 1 1. Process Resources * to 3 of 3
Command ==> _____ *****

Action Codes: 1=Create, 2=Copy, 3=Modify, 4=Browse, 5=Print, 6=Delete,
              /=Menu Bar

Action  Scheduling Environment Name  Description
--     -
--     CICS_SPEC                     Special CICS Processing
--     DB2LATE                       Offshift DB2 Processing
--     IMSPRIME                      Prime IMS Processing

```

Figure 65. Scheduling Environment Selection List Panel

This will take you directly to the Resource Definition List panel, where where you can add new resources. Enter an A on any Action line, as shown in Figure 66. (Alternately, you can put a / next to the resource names and then put the cursor under the Resources field at the top of the screen and hit ENTER. Then select 1 from the options in the pop-up box.)

```

Resources  Notes  Options  XREF  Help
-----
IWMAPAR   Resource Definition List   Row 1 to 4 of 4
Command ==> _____

Action Codes: A=Add D=Delete X=XREF /=Menu Bar

Action  Resource Name  In Use  Resource Description
--     -
--     CICS001      YES    CICS Subsystem
--     DB2A        YES    DB2 Subsystem
--     IMS1        YES    IMS Subsystem
a_     PRIMETIME      YES    Peak Business Hours

```

Figure 66. Resource Definition List Panel

This would take you to the Define Resource panel, where you would enter the information for the new resource name, as shown in Figure 67.

```

IWMAPAE   Define Resource
Command ==> _____

Resource name . . . . . VECTOR _____ Required
Resource Description . . . . . Vector Processor _____

```

Figure 67. Define Resource Panel

When we return to the Resource Definition List panel, the new resource name is there, as shown in Figure 68 on page 203.

```

Resources  Notes  Options  XREF  Help
-----
IWMAPAR          Resource Definition List          Row 1 to 5 of 5
Command ==>>> _____

Action Codes: A=Add  D=Delete  X=XREF  /=Menu Bar

Action  Resource Name      In Use  Resource Description
--      -
--      CICS001                YES     CICS Subsystem
--      DB2A                  YES     DB2 Subsystem
--      IMS1                 YES     IMS Subsystem
--      PRIMETIME           YES     Peak Business Hours
--      VECTOR              YES     Vector Processor

```

Figure 68. Resource Definition List Panel

Note that the panel automatically alphabetizes the resource names. Note also that VECTOR is not yet shown as In Use. We would need to go create or modify a scheduling environment and then select VECTOR.

### Creating a New Resource While Working with a Scheduling Environment

While you are in the middle of creating or modifying a scheduling environment, you can take this shortcut method to create new resources. In the above case where you wanted to create the new VECTOR resource and add it to DB2LATE, this method would be the logical choice.

From either the Create a Scheduling Environment panel or the Modify a Scheduling Environment panel (they look very similar), enter an A on any Action line, as shown in Figure 69.

```

Scheduling-Environments  Notes  Options  Help
-----
IWMAPAJ          Modify a Scheduling Environment          Row 1 to 2 of 2
Command ==>>> _____

Scheduling Environment Name  DB2LATE_____ Required
Description  . . . . . Offshift DB2 Processing_____

Action Codes: A=Add  D=Delete

Action  Resource Name      Required State  Resource Description
--      -
--      DB2A              ON_____ DB2 Subsystem
a_      PRIMETIME         OFF_____ Peak Business Hours

```

Figure 69. Modify a Scheduling Environment Panel

This would take you to the Resource Definition List panel, where you can enter an A on any Action line, as shown in Figure 70 on page 204. (Alternately, you can put a / next to the resource names and then put the cursor under the Resources field at the top of the screen and hit ENTER. Then select 1 from the options in the pop-up box.)

```

Resources  Notes  Options  XREF  Help
-----
IWMAPAS           Resource Definition List           Row 1 to 4 of 4
Command ==>>> _____

Selection For Scheduling Environment DB2LATE

Action Codes: A=Add  S=Select  X=XREF  /=Menu Bar

Action  Resource Name  In Use  Resource Description
--      CICS001         YES    CICS Subsystem
--      DB2A           YES    DB2 Subsystem
--      IMS1         YES    IMS Subsystem
a_     PRIMETIME       YES    Peak Business Hours

```

Figure 70. Resource Definition List Panel

As in the method described in “Creating a New Resource Directly from the Main Panel” on page 201, from here you’ll go to the Define Resource panel (as was shown in Figure 67 on page 202) to enter the new resource name and description.

When you come back to the Resource Definition List panel, you can now select the new VECTOR resource to be part of the DB2LATE scheduling environment, as shown in Figure 71.

```

Resources  Notes  Options  XREF  Help
-----
IWMAPAS           Resource Definition List           Row 1 to 5 of 5
Command ==>>> _____

Selection For Scheduling Environment DB2LATE

Action Codes: A=Add  S=Select  X=XREF  /=Menu Bar

Action  Resource Name  In Use  Resource Description
--      CICS001         YES    CICS Subsystem
--      DB2A           YES    DB2 Subsystem
--      IMS1         YES    IMS Subsystem
--      PRIMETIME       YES    Peak Business Hours
s_     VECTOR         YES    Vector Processor

```

Figure 71. Resource Definition List Panel

Having done that and exiting, you come back to the Modify a Scheduling Environment panel, as shown in Figure 72.

```

Scheduling-Environments  Notes  Options  Help
-----
IWMAPAJ           Modify a Scheduling Environment           Row 1 to 3 of 3
Command ==>>> _____

Scheduling Environment Name  DB2LATE_____ Required
Description . . . . . Offshift DB2 Processing_____

Action Codes: A=Add  D=Delete

Action  Resource Name  Required State  Resource Description
--      DB2A           ON_____ DB2 Subsystem
--      PRIMETIME    OFF_____ Peak Business Hours
--      VECTOR       ON_____  Vector Processor

```

Figure 72. Modify a Scheduling Environment Panel

Now the VECTOR resource name is part of the DB2LATE scheduling environment. As always, you need to specify either ON or OFF under the Required State. In this case, we chose ON.

## Summary of the Two Methods for Creating New Resources

You can create new resources in two ways:

- Directly from the main panel (“Creating a New Resource Directly from the Main Panel” on page 201). This method is useful when you want to create resources without actively working on one particular scheduling environment. Also, note that when you go to the Resource Definition List panel with this method, the Action Codes look like this:

```
Action Codes: A=Add D=Delete X=XREF /=Menu Bar
```

There is no S=Select, because you are not actively working on any scheduling environment (and therefore there is no place for you to select resources to).

There is a D=Delete action code. This will be important when you are deleting resources, as in “Deleting a Resource” on page 206.

- While creating or modifying scheduling environments (“Creating a New Resource While Working with a Scheduling Environment” on page 203). This method is useful when you want to create resources while in the midst of creating or modifying a scheduling environment. Note that when you go to the Resource Definition List panel with this method, the Action Codes (and the line above it) look like this:

```
Selection For Scheduling Environment DB2LATE
```

```
Action Codes: A=Add S=Select X=XREF /=Menu Bar
```

The name of scheduling environment that you are currently working with appears on this panel, and there is a S=Select among the action codes. You can select the new resource to be a part of the scheduling environment immediately after creating it.

There is no D=Delete action code, because you *cannot* delete resources when you come the Resource Definition List panel this way. See “Deleting a Resource” on page 206.

## Showing All Cross-References for a Resource Definition

You can check which specific scheduling environments use a given resource definition by entering an X in the action column next to the resource name, as shown in Figure 73 on page 206. (You can use either path to get to the Resource Definition List path, as discussed in “Summary of the Two Methods for Creating New Resources”.)

```

Resources  Notes  Options  XREF  Help
-----
IWMAPAS          Resource Definition List          Row 1 to 4 of 4
Command ==>>> _____

Selection For Scheduling Environment DB2LATE

Action Codes: A=Add  S=Select  X=XREF  /=Menu Bar

Action  Resource Name      In Use  Resource Description
--      -
  X_    CICS001                YES     CICS Subsystem
  _     IMS1                   YES     IMS Subsystem
  _     PRIMETIME              YES     Peak Business Hours
  _     VECTOR                 YES     Vector Processor

```

Figure 73. Resource Definition List Panel

This will take you to the Resource Cross-Reference of Scheduling Environments panel, as shown in Figure 74.

```

IWMAPAY          Resource Cross-Reference Of Scheduling Environments Row 1 to 2 of 2
Command ==>>> _____

Resource . . . . . : PRIMETIME
Description . . . . . : Peak Business Hours

Scheduling Environment Name  Description
IMSPRIME                    Prime IMS Processing
DB2LATE                     Offshift DB2 Processing

```

Figure 74. Resource Cross-Reference Of Scheduling Environments Panel

### Deleting a Resource

Unlike adding a resource definition, there is only one path you can take to delete a resource definition: from the main Scheduling Environment Selection List panel (as discussed in “Summary of the Two Methods for Creating New Resources” on page 205).

Put the cursor under the Resources field at the top of the screen and hit ENTER. When the pop-up box appears, enter a 1, as was shown in Figure 65 on page 202. (In this pop-up, 1 is the only option.)

This will take you directly to the Resource Definition List panel. Enter a D on the Action line next to the resource definition you wish to delete, as shown in Figure 75.

```

Resources  Notes  Options  XREF  Help
-----
IWMAPAR          Resource Definition List          Row 1 to 5 of 5
Command ==>>> _____

Action Codes: A=Add  D=Delete  X=XREF  /=Menu Bar

Action  Resource Name      In Use  Resource Description
--      -
  _     CICS001                YES     CICS Subsystem
  _     DB2A                  YES     DB2 Subsystem
  _     IMS1                   YES     IMS Subsystem
  _     PRIMETIME              YES     Peak Business Hours
  D_    VECTOR                 YES     Vector Processor

```

Figure 75. Resource Definition List Panel

Note that you can only delete a resource definition that is not currently in use by any scheduling environment. If you attempted to delete the PRIMETIME resource in Figure 66 on page 202, you would receive an error message stating the resource cannot be deleted because a scheduling environment is using it. You must first check the cross-references for that resource definition (as in “Showing All Cross-References for a Resource Definition” on page 205) and then go into every scheduling environment that uses it to delete those references.

---

## Coordinating Updates to a Service Definition

Since you can keep your service definition in either the WLM couple data set or in PDSs, it is possible to have multiple, uncoordinated updates. You should decide on a process for updating your service definition.

If you decide to work from the WLM couple data set, installing and extracting the service definition when you need to edit it, you should keep a copy in a PDS. The WLM application allows only one user to update a PDS at a time. If you do try to edit a PDS that is in use, you are locked out, and get a message saying the data set is in use.

If you select the Install option, and someone has done an install since you’ve extracted, the application displays a pop-up panel asking for confirmation. This way you can prevent inadvertent over-writes of the service definition. Figure 76 shows the Overwrite warning panel.

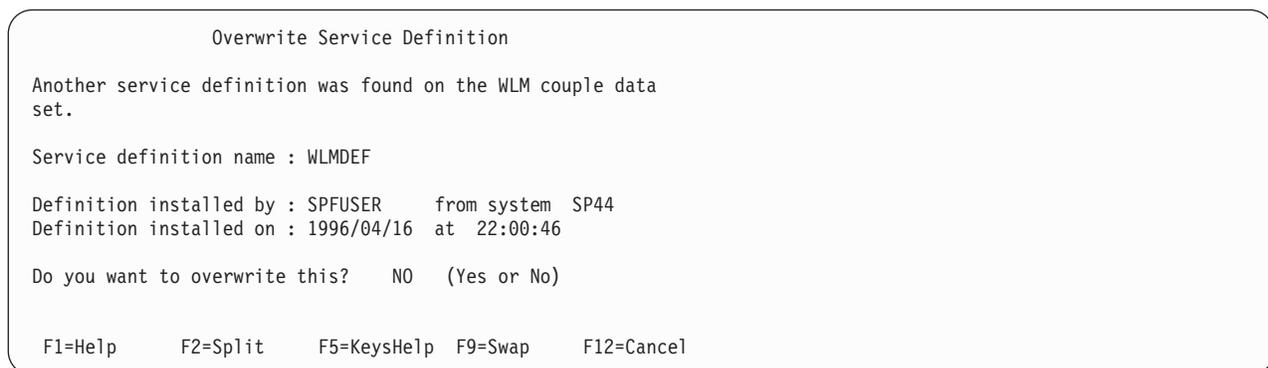


Figure 76. Overwrite service definition

## Using the WLM Couple Data Set

The service definition installed on the WLM couple data set is the *installed* service definition. The policy that you activate must be in the installed service definition.

Of course, before you can install your service definition on the WLM couple data set, you must first allocate one, and define it to the sysplex. You can allocate a WLM couple data set through a function in the WLM application, or you can use a utility. For information on how to allocate using the utility, see “Allocate a WLM Couple Data Set” on page 145.

### Allocating a WLM Couple Data Set

For availability purposes, you should allocate both a primary and an alternate WLM couple data set.

If your current WLM couple data set was allocated on OS/390 R3 or later, you can allocate a new primary and alternate WLM couple data set from the application using the *Allocate couple data set using CDS values* function on the Utilities menu bar

option of the Definition Menu. Figure 77 shows the allocate panel that uses CDS values. (If your current couple data set was allocated on a pre-OS/390 R3 system, do not use this function. Use the other allocation function called *Allocate a couple data set.*)

```

Allocate couple data set for WLM using CDS values
Command ==> _____

Sysplex name   SYSPLEX1 (Required)
Data set name  'GAILW.PCOUPLE.P1' (Required)

-----

Size parameters (optional):
Service policies . . 10 (1-99)
Workloads . . . . . 35_ (1-999)
Service classes . . 30_ (1-100)
Application
environments . . . . 50_ (1-999)
Scheduling
environments . . . . 50_ (1-999)
SVDEF extensions . . 0__ (0-8092)
SVDCR extensions . . 0__ (0-8092)
SVAEA extensions . . 0__ (0-8092)
SVSEA extensions . . 0__ (0-8092)

Storage parameters:
Storage class . . . . A
Management class . . _____

or

Volume . . . . . _____
Catalog data set?  _ (Y or N)

```

Figure 77. Allocate Couple Data Set Using CDS Values

This function primes the size parameters with values from the current couple data set. This way, you can ensure that the new couple data set is at least as large as the current one. If the new one is smaller than the current one, you will not be able to use SETXCF to make this data set available to the sysplex. (You will receive a series of messages starting with IXC255I UNABLE TO USE DATA SET.) You may need to increase one or more of these primed values if you have added new objects to your service definition since the last time you installed it.

If there is no current WLM couple data set, or you are migrating and your current couple data set was allocated on OS/390 R2 or earlier, use the *Allocate couple data set* function on the Utilities menu. This function fills in the size parameters based on the service definition you are about to install. It gets the values from the ISPF PDS containing the service definition. Figure 78 on page 209 shows the allocate panel that uses service definition values.

```

Allocate couple data set for WLM

Sysplex name   SYSPLEX2 (Required)
Data set name  'KIRSTEN.PCOUPLE.P1' (Required)

-----
-----

Size parameters (optional):
Service policies . . 10_ (1-99)
Workloads . . . . . 35_ (1-999)
Service classes . . 30_ (1-100)
Application
environments . . . . 50_ (1-999)
Scheduling
environments . . . . 50_ (1-999)

Storage parameters:
Storage class . . . . A
Management class . . _____
or
Volume . . . . . _____
Catalog data set? _ (Y or N)

```

Figure 78. Allocate Couple Data Set Using Service Definition Values

If you have defined your service definition, the application fills in the number of policies, workloads, service classes, application environments, and scheduling environments in that definition. These actual values are the best estimate of what this service definition requires when installed on the couple data set. If you define too large a number, it can use up a lot of space!

If you are replacing couple data sets already in use, the new data sets must be at least as large as the existing ones. If the existing data sets were allocated using pre-OS/390 R3 releases, you need to run a special sizing utility to determine their size. See “Calculate the Size of the WLM Couple Data Set” on page 145 for instructions.

If the existing data sets were allocated using OS/390 R3, use the method shown below. This method ensures that the new couple data set is at least as large as the current one, and, in most cases, can accommodate any new objects you might have added to the service definition.

1. First use the *Allocate couple data set* option as a quick way of displaying the values required for the service definition you have just created or modified. Write down these values and exit the panel without allocating a couple data set.
2. Then use the *Allocate couple dataset using CDS values* function to obtain a panel primed with the current couple data set values.
3. Take the maximum of the values from steps 1 and 2 and enter these values on the panel. Complete the couple data set allocation from this panel.

Once you have allocated the WLM couple data sets, you need to make them available for use in the sysplex. If you are making couple data sets available for the first time, see “Make a WLM Couple Data Set Available to the Sysplex for the First Time” on page 149. If you are replacing couple data sets already in use, see “Make a Newly Formatted Couple Data Set Available to the Sysplex” on page 150.

### Installing and Extracting a Service Definition

You can install and extract the service definition using the *install* and *extract* functions on the Utilities menu bar option of the Definition Menu.

Installing the service definition over-writes any service definition previously installed on the WLM couple data set.

When you extract the service definition, a copy remains on the WLM couple data set until the next install.

## Using MVS PDSs

If you have multiple people updating the service definition, or if you want to work with more than one service definition, you can work in multiple MVS PDSs. For example, if you want to work on the service definition for next year's SLA, you can keep it in a PDS.

If you specify *Read saved definition*, you can choose from a list of data sets already created. The list is made up of all data sets your userid has worked with (the list is built based on use).

---

## Restricting Access to Your Service Definition

To restrict access to the service definition stored in both the PDS and the WLM couple data set, you can use the same resource control facility as you do for a PDS. For information about how to restrict access to the WLM application functions, see "Restricting Access to the WLM Service Definition" on page 138.

---

## Activating a Service Policy

You can activate a service policy from the ISPF application, as well as from the operator console. To activate a policy from the operator console, you issue the VARY WLM command. For more information about issuing the VARY command, see "Migration Activities" on page 138. To activate a policy from the application, choose the utilities option from the menu bar on the Definition menu.

When you choose the activate option, the application extracts the service definition information from the WLM couple data set, and displays the *Policy Selection List*, as shown in Figure 79 on page 211. From the list, you select the service policy that you want to activate.

If you have been editing your service definition and want your changes to take effect, you have to install the changed service definition, and then activate the service policy.

```

Policy Selection List          ROW 1 TO 3 OF 3

The following is the current Service Definition installed on the WLM
couple data set.

Name . . . . : SLA1993

Installed by : USERID    from system SYSTEM
Installed on : 1996/04/01 at 11:30

Select the policy to be activated with "/"

Sel Name      Description                                ---Last Change---
  - BACKUPS   During weekly backups                    User      Date
  - DAYTIME   Policy from 7:00am to 5:00 pm                NORTH    1996/05/11
  - NIGHT     Late night batch window                            NORTH    1996/05/11
***** BOTTOM OF DATA *****

Command ==> _____

```

Figure 79. Activating a Service Policy

---

## Printing in the Application

The print function on the selection lists in the WLM application prints the information into an ISPF list data set. You can process the ISPF list data set just as you would any other ISPF list data set. For more information about using list data sets, see the *ISPF Dialog Management Guide and Reference*.

The “Print as GML” function under the File action bar choice of the Definition menu prints the complete service definition to a data set in GML format. The data set can be formatted using SCRIPT/VIS. For more information, see the *DCF SCRIPT/VIS User’s Guide*.

**Note:** To use SCRIPT/VIS, you must have Document Composition Facility (DCF) Version 3 or later installed.

---

## Browsing Definitions

The browse function is available on selection lists in the application, and is similar to the print function. Browse prints information into a temporary data set, and displays it.

Figure 80 on page 212 shows the browse on the BAT\_T service class from the service class selection list.

```

Browse                               Line 00000000 Col 001 072
***** TOP OF DATA *****
* Service Class BAT_T - Batch class T

Created by user BERKEL on 1996/10/07 at 11:29:05
Base last updated by user BERKEL on 1996/10/07 at 11:29:05

Base goal:

# Duration Imp Goal description
- - - - -
1          5 80% complete within 00:30:00.000

***** BOTTOM OF DATA *****

Command ==>

```

Figure 80. Browse Function from the Service Class Selection List

## Using XREF Function to View Service Definition Relationships

To view the relationship between any two definitions (service class, workload resource group etc), you can use the XREF function. From any Modify panel in the application, you can use the XREF function to determine an definition's relationship to another.

For example, on the Modify Service Class panel, you may want to know whether a service policy overrides the defined goal. You use the XREF function to check whether its goals are overridden in a service policy. From the Modify Service Class panel, you can also check which subsystem types reference that service class in the classification rules. If there are some that do reference it, the application displays a pop-up selection list where you can browse the Rules referencing that service class. Figure 81 shows an example of the pop-up displayed for an Xref by Subsystem for a service class.

```

Service Class Reference ROW 1 TO 1 OF 1

These subsystem types refer to the service class.

Action Codes: 4=Browse

Action Name Description
_         STC  IBM-defined subsystem type
***** BOTTOM OF DATA *****

```

Figure 81. Service Class Subsystem Xref

## WLM Application Messages

---

**IWMAM040 Unexpected error, RC=xx,  
RSN=xxxxxxx.**

**Explanation:** The application has detected an unexpected error.

In the message text:

xx        The return code.

xxxxxxx  
          The hexadecimal reason code.

**Source:** Workload manager (WLM)

**System Action:** The requested operation is not performed.

**Service Administrator Response:** Search problem reporting data bases for a fix for the problem. If no fix exists, contact the IBM Support Center. Provide the text of this message.

---

**IWMAM041 WLM couple data set is unavailable.**

**Explanation:** Workload manager (WLM) was unable to find the WLM couple data set. It is possible that it does not exist. Or it may exist, but it has not been defined to the sysplex.

**Source:** Workload manager (WLM)

**System Action:** The requested operation is not performed.

**Service Administrator Response:** Either:

- Make the couple data set available to the sysplex by updating the DATA keyword of the COUPLExx parmlib member or issuing the SETXCF command.
- Allocate the WLM couple data set.

---

**IWMAM042 Extract failed, no service definition was found on the WLM couple data set.**

**Explanation:** Workload manager (WLM) was unable to find a service definition on the WLM couple data set.

**Source:** Workload manager (WLM)

**System Action:** The extract is not performed.

**Service Administrator Response:** Ensure that the service definition has been properly installed on the WLM couple data set.

---

**IWMAM043 Unable to obtain storage,  
RSN=xxxxxxx.**

**Explanation:** Workload manager (WLM) was unable to obtain enough storage to complete the operation.

In the message text:

xxxxxxx

The hexadecimal reason code.

**Source:** Workload manager (WLM)

**System Action:** The requested operation is not performed.

**Service Administrator Response:** Increase the region size and repeat the operation. If this does not help, search problem reporting data bases for a fix for the problem. If no fix exists, contact the IBM Support Center. Provide the text of this message.

---

**IWMAM044 Install failed, service definition is not valid. Validation reason code: xxxx,  
Validation offset: yyyy**

**Explanation:** The service definition you are trying to install is not valid.

xxxx is the validation reason code. yyyy is the validation offset (a hex offset into the iwmszzz data structures where the problem exists).

**Source:** Workload manager (WLM)

**System Action:** The install is not performed.

**Service Administrator Response:** For further explanation of this error, see "Appendix B: Application Validation Reason Codes" in *z/OS MVS Programming: Workload Management Services*. Search problem reporting data bases for a fix for the problem. If no fix exists, contact the IBM Support Center. Provide the text of this message.

---

**IWMAM046 Errors were found during validation of the service definition. The install has failed.**

**Explanation:** The errors listed on the previous panel (shown when you attempted to install) have prevented the installation of the service definition.

**Source:** Workload manager (WLM)

**System Action:** The install is not performed.

**Service Administrator Response:** Correct the errors and retry. To capture a list of the errors, go back to the error panel by either attempting to install the service definition again or by validating the service definition (use the "Validate definition" utility as described in "Using the Menu Bar on the Definition Menu" on page 178). From that panel, you can capture a list of the errors by selecting "Save listing" on the "File" menu bar option.

---

**IWMAM050 Exceeded the maximum number of attempts to read the service definition.**

**Explanation:** A failure in reading the service definition from the WLM couple data set has occurred repeatedly.

**Source:** Workload manager (WLM)

**System Action:** The requested operation is not performed.

**Service Administrator Response:** Search problem reporting data bases for a fix for the problem. If no fix exists, contact the IBM Support Center. Provide the text of this message.

---

**IWMAM051 Access was denied to the WLM couple data set.**

**Explanation:** The user does not have appropriate RACF authority to the WLM couple data set.

**Source:** Workload manager (WLM)

**System Action:** The requested operation is not performed.

**Service Administrator Response:** Verify that the user should be authorized and have the RACF administrator give the user appropriate access. For information about how to restrict access to the couple data set, see "Restricting Access to the WLM Service Definition" on page 138.

---

**IWMAM052 The service definition functionality level (LEVELxxx) is not compatible with the WLM ISPF application level (LEVELyyy). To extract a service definition or to activate a policy, the WLM ISPF application and the MVS system must be at the same level as the service definition.**

**Explanation:** The service definition in the WLM couple data set uses functions that are not compatible with this level of the WLM ISPF application

**Source:** Workload manager (WLM)

**System Action:** The requested operation is not performed.

**Service Administrator Response:** Use the WLM ISPF application on a system that is compatible with the functionality level of the service definition.

---

**IWMAM054 Failure in ISPF: ISPF error information.**

**Explanation:** An error occurred using the application.

**Source:** Workload manager (WLM)

**System Action:** The requested operation is not performed.

**Service Administrator Response:** The ISPF error

information may provide a clue as to how to overcome the problem. If not, search problem reporting data bases for a fix for the problem. If no fix exists, contact the IBM Support Center. Provide the text of this message.

---

**IWMAM055 Extract failed, service definition is not valid. Validation reason code: xxxx, Validation offset: yyyy**

**Explanation:** The service definition on the WLM couple data set is not valid. The data set may be corrupted.

xxxx is the validation reason code. yyyy is the validation offset.

**Source:** Workload manager (WLM)

**System Action:** The extract is not performed.

**Service Administrator Response:** If the data set is corrupted, try restoring it from back-ups.

---

**IWMAM058 Install failed. WLM couple data set has not been reallocated for use with this OS/390 release. Reallocate the WLM couple data set. Refer to the migration chapter in "MVS Planning: Workload Management."**

**Explanation:** The WLM couple data set upon which you are attempting to install a service definition is not formatted for your service definition.

**Source:** Workload manager (WLM)

**System Action:** The system does not install the service definition.

**Service Administrator Response:** Re-allocate the WLM couple data set for the service definition.

---

**IWMAM072 The service definition was not read due to a mismatch between the service definition PDS (LEVELxxx) and the WLM ISPF application (LEVELyyy). To read the service definition PDS, restart the WLM ISPF application at level LEVELxxx or higher.**

**Explanation:** The service definition in the PDS uses functions that are not compatible with this level of the WLM ISPF application.

**Source:** Workload manager (WLM)

**System Action:** The requested operation is not performed.

**Service Administrator Response:** Use the WLM ISPF application on a system that is compatible with the functionality level of the service definition.

---

**IWMAM077 Unable to use *datasetname* for service definition data, member *membername* has an unrecognized format.**

**Explanation:** WLM cannot use the *datasetname* definition data as a WLM service definition. WLM has determined that member *membername* in the PDS contains information that is not recognized by WLM. The incorrect member contains key and name information that does not match what WLM expects. Note that the member shown is the first member that was found to be invalid. It is possible that others are invalid, or that the entire dataset is corrupted.

**Source:** Workload manager (WLM)

**System Action:** The requested operation is not performed.

**Service Administrator Response:** If you are using a downlevel level of the WLM ISPF, check to see that the appropriate compatibility APARs are installed on the downlevel system. (See “Migrating to a New z/OS Release With an Existing Service Definition” on page 135.) Otherwise, the data set may be corrupted. If you cannot see an obvious problem in the data set, contact the IBM Support Center.



---

## Appendix A. Customizing the WLM ISPF Application

This appendix explains how to customize your WLM application to:

- Customize the WLM application libraries.  
If you have re-named or changed your IPCS/WLM library names, then you can use the IWMAREX1 exit to specify your library names.
- Customize the WLM application data sets.  
If you would like to allocate the application data sets with your storage management policies, use IWMAREX2.
- Add the WLM application as an option on your ISPF menu.  
If you plan to use the WLM application frequently, you can add the WLM application as an option on your ISPF application.
- Move pop-up windows.
- Customize the keylists.

---

### Introduction

#### Specifying the Exits

To start the WLM application, you use a TSO/E REXX exec IWMARIN0. IWMARIN0 concatenates the IPCS/WLM data sets, allocates some data sets required for a service definition, and invokes the application panels. If EXITS is specified on the EXEC statement, IWMARIN0 uses the exits in the specified data set.

Table 16 shows the return codes from IWMARIN0.

Table 16. Return Codes from IWMARIN0

Return Code	Explanation
4	Not in ISPF. Application cannot be started.
8	Unexpected keyword(parameter) on WLM application invocation. Unexpected keyword(parameter) is: <i>keyword</i> .  <i>keyword</i> represents the keyword(parameter) from the command invocation.
12	Unexpected error occurred when calling installation exit IWMAREX1. One of the following: <ul style="list-style-type: none"><li>• Installation exit IWMAREX1 must exist in <i>data-set name</i> for the WLM application to run. <i>data-set name</i> represents the data set that must contain the installation exit. Check to make sure the data set contains IWMAREX1.</li><li>• Installation exit IWMAREX1 must exist in the current concatenation order for the WLM application to run.</li></ul>
16	Unexpected keyword(parameter) from WLM exit IWMAREX1. Please check coding in WLM installation exit IWMAREX1 for incorrect keyword(parameter): <i>parameter</i> .  <i>parameter</i> represents the keyword(parameter) that is incorrect.

Table 16. Return Codes from IWMARIN0 (continued)

Return Code	Explanation
20	Unexpected RC="rc" from TSO ALTLIB ISPF LIBDEF for <i>data-set-name</i>  The WLM application cannot be started due to ALTLIB or LIBDEF failures for <i>data-set-name</i> . See specific REXX messages for the names of the data sets which failed.
24	ALTLIB failed during attempt to find installation exit data set. WLM application can not be started. TSO ALTLIB RC=xx.  xx represents return code from TSO ALTLIB service.

## Coding the WLM Exits

The exit stubs are shipped in SYS1.SBLSCLI0. If you have a need to customize the application with the exits, you should create a data set for your exits, copy IWMAREX1 or IWMAREX2 into the data set, and modify the exits with your options. When you create the data set for your exits, you should specify the same data set characteristics as SYS1.SBLSCLI0.

You specify the exit data set on the TSO/E EXEC statement when you start the WLM application. Specify the fully qualified data set name containing the exit. For example, suppose you created data set IPCS.EXIT, and coded IWMAREX1. To start the WLM application with the exit, you specify:

```
EX 'SYS1.SBLSCLI0(IWMARIN0)' 'EXIT(IPCS.EXIT)'
```

## IWMARIN1

If you have previously allocated the required data sets, either in a logon procedure, or in a CLIST, you can use IWMARIN1 to start the WLM application. To use IWMARIN1, specify:

```
EX 'SYS1.SBLSCLI0(IWMARIN1)'
```

---

## Customizing the WLM Application Libraries - IWMAREX1

WLM provides the IWMAREX1 exit to specify the IPCS/WLM libraries names. If you have re-named or customized the following IPCS/WLM libraries, use IWMAREX1 to specify your names.

SYS1.SBLSCLI0	Application REXX code data set.
SYS1.SBLSKEL0	Application skeleton data set.
SYS1.SBLSPNL0	Application panel data set.
SYS1.SBLSTBL0	Application keylists and commands data set.
SYS1.SBLSMSG0	Application messages data set

If you have renamed them in your installation, use IWMAREX1 to set up the allocations.

IWMAREX1 is a REXX routine for specifying installation-customized data sets required for starting the WLM ISPF application.

If you have renamed or changed the WLM/IPCS data sets in your installation, use IWMAREX1 to set up the allocations.

## Processing

IWMAREX1 is called from the IWMARIN0 REXX exec.

## Parameters

IWMAREX1 has the following parameters:

### REXXDS

The application REXX code data set.

### SKELDS

The application skeleton data set.

### PANELDS

The application panel data set.

### TABLEDS

The application tables (keylist and commands) data set.

### MESSAGEDS

The application messages data set.

## Example

### Example: Using IWMAREX1

Suppose the IPCS/WLM application resides in SYS1.IPCS.SBLSCLI0, and you have re-named your IPCS/WLM libraries to:

```
SYS1.IPCS.SBLSMSG0  
SYS1.IPCS.SBLSPNL0  
SYS1.IPCS.SBLSKEL0  
SYS1.IPCS.SBLSTBL0
```

Suppose you have created your exit in a data set called WLM.EXITS. You code IWMAREX1 in the following way:

```
/* REXX */  
queue 'REXDS(SYS1.IPCS.SBLSCLI0)'  
queue 'SKELDS(SYS1.IPCS.SBLSKEL0)'  
queue 'PANELDS(SYS1.IPCS.SBLSPNL0)'  
queue 'TABLEDS(SYS1.IPCS.SBLSTBL0)'  
queue 'MESSAGEDS(SYS1.IPCS.SBLSMSG0)'  
Exit 0
```

To start the WLM application with the exit, you specify:

```
EX 'SYS1.IPCS.SBLSCLI0(IWMARIN0)' 'EXIT(WLM.EXITS)'
```

---

## Customizing the WLM Application Data Sets - IWMAREX2

WLM provides the IWMAREX2 exit to specify:

- The application recovery data set and data set characteristics.
- The service definition data set and data set characteristics.

You should use this exit to allocate data sets according to your installation's storage management policies. You can not change any of the following TSO/E ALLOCATE parameters in IWMAREX2:

- Data set name (DA)
- Record format (RECFM)
- Logical record length (LRECL)
- Data set organization (DSORG)

Unless you code otherwise in this exit, the WLM application uses the following data set information as the defaults:

```
UNIT(SYSDA)  
TRACKS SPACE(5,5)
```

## Processing

This exit is called from the IWMARZAL REXX exec.

## Parameters

IWMAREX2 has the following parameters:

### ARDSDIR

Specifies the number of 256 byte records to be allocated for the application recovery data set.

### ARDSOPTS

Specifies the options for the application recovery data sets, with the following sub-parameters:

**UNIT(xxxx)**

Specifies the unit type for the TSO ALLOCATE command.

**STORCLAS**

The SMS storage class.

**MGMTCLAS**

The SMS management class.

**DATACLAS**

The SMS data class.

To determine what to specify for unit, check which UNIT type is coded for the TSO ALLOCATE command in your installation.

**ARDSSPACE**

Specifies the options to allocate a data set. The options are:

**SPACE(quantity,(increment))****TRACKS****CYLINDERS****BLOCKS(value)****PDDSDIR**

Specifies the number of 256 byte records to be allocated for the print service definition data set.

**PDDSOPTS**

Specifies the options for the print service definition data set, with the following sub-parameters:

**UNIT(xxxx)**

Specifies the unit type for the TSO ALLOCATE command.

**STORCLAS**

The SMS storage class.

**MGMTCLAS**

The SMS management class.

**DATACLAS**

The SMS data class.

To determine what to specify for unit, check which UNIT type is coded for the TSO ALLOCATE command in your installation.

**PDDSSPACE**

Specifies the options to allocate a print definition data set. The options are:

**SPACE(quantity,(increment))****TRACKS****CYLINDERS****BLOCKS(value)****SDDSDIR**

Specifies the number of 256 byte records to be allocated for the service definition data set.

**SDDSOPTS**

Specifies the options for the service definition data sets, with the following sub-parameters:

**UNIT(xxxx)**

Specifies the unit type for the TSO ALLOCATE command.

**STORCLAS**

The SMS storage class.

**MGMTCLAS**

The SMS management class.

## DATACLAS

The SMS data class.

## SDDSSPACE

Specifies the options to allocate a data set. The options are:

**SPACE(quantity,(increment))**

**TRACKS**

**CYLINDERS**

**BLOCKS(value)**

## Examples

### Examples: Using IWMAREX2

- Suppose you want to specify that the UNIT type for your TSO allocate commands in your installation is type SYSDS. In exit IWMAREX2, you specify the following:

```
/* REXX */
'ARDSOPTS(UNIT(SYSDS))'
Exit 0
```

- Suppose you want to specify the service definition data sets as SMS managed data sets in the standard storage class, and in the NOMIG management class. In exit IWMAREX2, you specify the following:

```
/* REXX */
queue 'ARDSOPTS(STORCLAS(STANDARD) MGMTCLAS(NOMIG))'
queue 'ARDSSPACE(SPACE(10,10) TRACKS)'
Exit 0
```

---

## Adding WLM as an ISPF Menu Option

To add the WLM application as an option on your ISPF primary panel, you should make a copy of the ISPF primary option menu - ISR@PRIM. You need to add some information to the processing section of the panel. What you add depends on whether you use IWMARIN0 or IWMARIN1 to start the application.

- If you use IWMARIN0, specify:

```
WLM,'CMD(EXEC "SYS1.SBLSCLI0(IWMARIN0)") NEWAPPL(IWMP) PASSLIB'
```

- If you use the WLM exits together with IWMARIN0, specify :

```
WLM,'CMD(EXEC "SYS1.SBLSCLI0(IWMARIN0)" "EXIT(ISPF.EXITS)") NEWAPPL(IWMP) PASSLIB'
```

- If you use IWMARIN1, specify:

```
WLM,'CMD(EXEC "SYS1.SBLSCLI0(IWMARIN1)") NEWAPPL(IWMP) PASSLIB'
```

Make sure you concatenate the library containing your customized primary panel before any others in your logon procedure or CLIST.

The following example shows the two lines added to a copy of the ISPF primary panel. In the example, IWMARIN1 is specified. The two lines are highlighted in the figure.

### Example: Adding WLM as an Option on your ISPF Menu

```
----- ISPF/PDF PRIMARY OPTION MENU -----
%OPTION ==>_ZCMD
%
% 0 +ISPF PARS - Specify terminal and user parameters +USERID -
% 1 +BROWSE - Display source data or output listings +TIME -
% 2 +EDIT - Create or change source data +PF KEYS -
% 3 +UTILITIES - Perform utility functions
% 4 +FOREGROUND - Invoke language processors in foreground
% 5 +BATCH - Submit job for language processing
% 6 +COMMAND - Enter TSO/E command or CLIST
% 7 +DIALOG TEST - Perform dialog testing
% 8 +LM UTILITIES- Perform library management utility functions
% C +CHANGES - Display summary of changes for this release
% W +WLM - WLM administrative application
% T +TUTORIAL - Display information about ISPF/PDF
% X +EXIT - Terminate ISPF using log and list defaults
%
+Enter%END+command to terminate ISPF.
%
)INIT
 .HELP = ISR00003
 &ZPRIM = YES /* ALWAYS A PRIMARY OPTION MENU */
 &ZHTOP = ISR00003 /* TUTORIAL TABLE OF CONTENTS */
 &ZHINDEX = ISR91000 /* TUTORIAL INDEX - 1ST PAGE */
 VPUT (ZHTOP,ZHINDEX) PROFILE
)PROC
 &ZSEL = TRANS( TRUNC (&ZCMD, '.')
 0, 'PANEL(ISPOPTA)'
 1, 'PGM(ISRBRO) PARM(ISRBRO01)'
 2, 'PGM(ISREDIT) PARM(P,ISREDM01)'
 3, 'PANEL(ISRUTIL)'
 4, 'PANEL(ISRFPA)'
 5, 'PGM(ISRJB1) PARM(ISRJPA) NOCHECK'
 6, 'PGM(ISRPTC)'
 7, 'PGM(ISRYXDR) NOCHECK'
 8, 'PANEL(ISRLPRIM)'
 C, 'PGM(ISPTUTOR) PARM(ISR00005)'
 W, 'CMD(EXEC "SYS1.SBLSCLI0(IMMARIN0)") NEWAPPL(IWMP) PASSLIB'
 T, 'PGM(ISPTUTOR) PARM(ISR00000)'
 ' '
 X, 'EXIT'
 *, '?' )
 &ZTRAIL = .TRAIL
)END
```

## Moving Pop-up Windows

If you would like to customize the placement of the pop-up windows in the WLM application, you can use a manual ISPF function. You must place the cursor anywhere on the active window frame and press ENTER. ISPF acknowledges the window move request by displaying WINDOW MOVE pending message. Then place the cursor where you want the upper left corner of the pop-up placed. press Enter a second time, and the pop-up is moved to the new location.

**Note:** The placement lasts only for the duration of the session. If you exit the application, your changes are lost.

There are some other options for moving pop-up windows. For more information about them, see *ISPF Dialog Management Guide and Reference*.

## Customizing the Keylists

The WLM application uses a set of keylists which you can customize to your purposes. To edit a keylist, type:

```
KEYLIST
```

on the command line from any panel in the WLM application. The application displays the keylist utility panel, as shown in Figure 82. From this panel, select the keylist you would like to work with, and choose an action from **OPTIONS** on the menu bar.

```

Options  Change Keylists
-----
ISPKLUP      Keylist Utility for IWMP      ROW 1 TO 9 OF 9
Command ==> _____

Enter keylist name _____ OR

Select one keylist name from the list below:
Select  Keylist  T  -
-       KEYSBRP  S
-       KEYSWRK  S
-       KEYS001  S
-       KEYS002  S
-       KEYS01H  S
-       KEYS01P  S
-       KEYS01S  S *** CURRENTLY ACTIVE KEYLIST ***
-       KEYS02A  S
-       KEYS02B  S
-       KEYS02P  S

***** BOTTOM OF DATA *****

```

Figure 82. Keylist Utility Panel

The keylists and the type of panels on which they are used are:

Keylist	Function
KEYS001	Non-scrollable, non pop-up
KEYS002	Scrollable, non pop-up
KEYS01P	Non-scrollable, pop-up
KEYS02P	Scrollable, pop-up
KEYS01H	All help panels
KEYS01S	Non-scrollable, non pop-up, no PF4=RETURN
KEYS02A	Scrollable, non pop-up, PF1=HELPD, PF10=ACTIONS.
KEYS02B	Scrollable, non pop-up, PF1=HELPD, PF10=LEFT, PF11=RIGHT.
KEYSBRP	Browse panel
KEYSWRK	"Working..." panel

## Appendix B. CPU Capacity Table

The following table shows the unweighted CPU service units per second by CPU model. You use this information to define your minimum and maximum capacity for a resource group.

Table 17. Processors: zSeries 900

Processor Model	Service Units Per Second of Task or SRB Execution Time	Seconds of Task or SRB Execution Time Per Service Unit
<b>zSeries 900</b>		
zSeries 900, Model 101	11585.8	0.000086
zSeries 900, Model 102	10891.8	0.000092
zSeries 900, Model 103	10430.2	0.000096
zSeries 900, Model 104	10081.9	0.000099
zSeries 900, Model 105	9732.4	0.000103
zSeries 900, Model 106	9384.2	0.000107
zSeries 900, Model 107	9153.3	0.000109
zSeries 900, Model 108	8805.7	0.000114
zSeries 900, Model 109	8456.7	0.000118
zSeries 900, Model 110	9334.9	0.000107
zSeries 900, Model 111	9211.3	0.000109
zSeries 900, Model 112	8968.6	0.000112
zSeries 900, Model 113	8724.1	0.000115
zSeries 900, Model 114	8602.2	0.000116
zSeries 900, Model 115	8359.5	0.00012
zSeries 900, Model 116	8117.7	0.000123
zSeries 900, Model 1C1	12112	0.000083
zSeries 900, Model 1C2	11502.5	0.000087
zSeries 900, Model 1C3	11142.1	0.00009
zSeries 900, Model 1C4	10781.7	0.000093
zSeries 900, Model 1C5	10540.2	0.000095
zSeries 900, Model 1C6	10296	0.000097
zSeries 900, Model 1C7	10056.6	0.000099
zSeries 900, Model 1C8	9816	0.000102
zSeries 900, Model 1C9	9575.1	0.000104

Table 18. Processors: S/390 9672, 2003, 3000, and 7060

Processor Model	Service Units Per Second of Task or SRB Execution Time	Seconds of Task or SRB Execution Time Per Service Unit
<b>S/390 9672 Models R1, E, and P</b>		
S/390 9672, Model R11 (1 CP per CPC)	696.0	0.001437
S/390 9672, Model R21 (2 CPs per CPC)	661.2	0.001512
S/390 9672, Model R31 (3 CPs per CPC)	619.4	0.001614
S/390 9672, Model R41 (4 CPs per CPC)	584.6	0.001710
S/390 9672, Model R51 (5 CPs per CPC)	542.9	0.001842
S/390 9672, Model R61 (6 CPs per CPC)	508.0	0.001968
<b>S/390 9672 R2 Models</b>		
S/390 9672, Model RA2	718.0	0.001393

Table 18. Processors: S/390 9672, 2003, 3000, and 7060 (continued)

Processor Model	Service Units Per Second of Task or SRB Execution Time	Seconds of Task or SRB Execution Time Per Service Unit
S/390 9672, Model RB2	488.7	0.002046
S/390 9672, Model RC2	287.3	0.003481
S/390 9672, Model R12	1029.1	0.000972
S/390 9672, Model R22	967.4	0.001034
S/390 9672, Model R32	926.2	0.001080
S/390 9672, Model R42	885.1	0.001130
S/390 9672, Model R52	843.9	0.001185
S/390 9672 (logical partition of 6 CPs)	802.8	0.001246
S/390 9672, Model R72	761.6	0.001313
<b>S/390 9672 R3 Models</b>		
S/390 9672 (logical partition of 1 CP)	1090.0	0.000917
S/390 9672 (logical partition of 2 CPs)	1035.5	0.000966
S/390 9672 (logical partition of 3 CPs)	1013.7	0.000986
S/390 9672 (logical partition of 4 CPs)	991.9	0.001008
S/390 9672, Model R53	981.0	0.001019
S/390 9672, Model R63	959.2	0.001043
S/390 9672, Model R73	926.5	0.001079
S/390 9672, Model R83	904.7	0.001105
S/390 9672 (logical partition of 9 CPs)	872.0	0.001147
S/390 9672, Model RX3	839.3	0.001191
<b>S/390 9672 R4 Models</b>		
S/390 9672, Model RA4	1587.8	0.000630
S/390 9672, Model RB4	1524.2	0.000656
S/390 9672 (logical partition of 1 CP)	2062.7	0.000485
S/390 9672 (logical partition of 2 CPs)	1980.2	0.000505
S/390 9672, Model RC4 (unique uni-base)	1897.8	0.000527
S/390 9672, Model R14	2264.7	0.000442
S/390 9672, Model R24	2174.2	0.000460
S/390 9672, Model R34	2083.6	0.000480
S/390 9672, Model R44	2038.2	0.000491
S/390 9672, Model R54	1970.2	0.000508
S/390 9672, Model R64	1902.3	0.000526
S/390 9672, Model R74	1857.0	0.000539
S/390 9672, Model R84	1789.1	0.000559
S/390 9672, Model R94	1721.2	0.000581
S/390 9672, Model RX4	1630.7	0.000613
<b>S/390 9672 RY4 Models</b>		
S/390 9672 (logical partition of 1 CP)	2486.8	0.000402
S/390 9672 (logical partition of 2 CPs)	2362.3	0.000423
S/390 9672 (logical partition of 3 CPs)	2287.7	0.000437
S/390 9672 (logical partition of 4 CPs)	2213.0	0.000452
S/390 9672 (logical partition of 5 CPs)	2163.3	0.000462
S/390 9672 (logical partition of 6 CPs)	2088.8	0.000479
S/390 9672 (logical partition of 7 CPs)	2014.1	0.000496
S/390 9672 (logical partition of 8 CPs)	1939.4	0.000516
S/390 9672 (logical partition of 9 CPs)	1864.8	0.000536
S/390 9672, Model RY4	1790.1	0.000559
<b>S/390 9672 G4 Models</b>		

Table 18. Processors: S/390 9672, 2003, 3000, and 7060 (continued)

Processor Model	Service Units Per Second of Task or SRB Execution Time	Seconds of Task or SRB Execution Time Per Service Unit
S/390 9672, Model RA5	2312.5	0.000432
S/390 9672, Model RB5	2058.1	0.000486
S/390 9672 (Partitioned 1-way off a 9672-RC5)	2622.5	0.000381
S/390 9672 (Partitioned 2-way off a 9672-RC5)	2439.0	0.000410
S/390 9672, Model RC5	2255.4	0.000443
S/390 9672, Model R15	2937.9	0.000340
S/390 9672, Model R25	2732.2	0.000366
S/390 9672, Model R35	2555.9	0.000391
S/390 9672, Model R45	2438.3	0.000410
S/390 9672 (Partitioned 1-way off a 9672-R55 through RX5)	3129.3	0.000320
S/390 9672 (Partitioned 2-way off a 9672-R55 through RX5)	2910.1	0.000344
S/390 9672 (Partitioned 3-way off a 9672-R55 through RX5)	2722.5	0.000367
S/390 9672 (Partitioned 4-way off a 9672-R55 through RX5)	2597.4	0.000385
S/390 9672, Model R55	2472.2	0.000404
S/390 9672, Model R65	2378.5	0.000420
S/390 9672, Model R75	2253.2	0.000444
S/390 9672, Model R85	2127.9	0.000470
S/390 9672, Model R95	2002.8	0.000499
S/390 9672, Model RX5	1908.9	0.000524
S/390 9672 (Partitioned 1-way off a 9672-RY5)	3496.5	0.000286
S/390 9672 (Partitioned 2-way off a 9672-RY5)	3252.0	0.000308
S/390 9672 (Partitioned 3-way off a 9672-RY5)	3042.4	0.000329
S/390 9672 (Partitioned 4-way off a 9672-RY5)	2902.8	0.000344
S/390 9672 (Partitioned 5-way off a 9672-RY5)	2762.9	0.000362
S/390 9672 (Partitioned 6-way off a 9672-RY5)	2657.8	0.000376
S/390 9672 (Partitioned 7-way off a 9672-RY5)	2518.1	0.000397
S/390 9672 (Partitioned 8-way off a 9672-RY5)	2378.1	0.000421
S/390 9672 (Partitioned 9-way off a 9672-RY5)	2238.1	0.000447
S/390 9672, Model RY5	2133.0	0.000469
<b>S/390 9672 G5 Models</b>		
S/390 9672, Model T16	6067.5	0.000165
S/390 9672, Model T26	5643.7	0.000177
S/390 9672, Model R36	5460.8	0.000183
S/390 9672, Model R46	5278.8	0.000189
S/390 9672, Model R56	5158.0	0.000194
S/390 9672, Model R66	5036.2	0.000199
S/390 9672, Model R76	4915.5	0.000203
S/390 9672, Model R86	4733.7	0.000211
S/390 9672, Model R96	4490.6	0.000223
S/390 9672, Model RX6	4247.4	0.000235
S/390 9672, Model Y16	7246.4	0.000138
S/390 9672, Model Y26	6739.7	0.000148
S/390 9672, Model Y36	6449.0	0.000155
S/390 9672, Model Y46	6230.5	0.000161
S/390 9672, Model Y56	6086.0	0.000164
S/390 9672, Model Y66	5941.3	0.000168
S/390 9672, Model Y76	5797.1	0.000173
S/390 9672, Model Y86	5578.8	0.000179
S/390 9672, Model Y96	5361.9	0.000187
S/390 9672, Model YX6	5071.3	0.000197
S/390 9672, Model RA6	4212.7	0.000237
S/390 9672, Model RB6	3960.4	0.000253
S/390 9672, Model R16	5588.5	0.000179

Table 18. Processors: S/390 9672, 2003, 3000, and 7060 (continued)

Processor Model	Service Units Per Second of Task or SRB Execution Time	Seconds of Task or SRB Execution Time Per Service Unit
S/390 9672, Model R26	5141.4	0.000194
S/390 9672, Model RC6	5029.9	0.000199
S/390 9672, Model RD6	4918.5	0.000203
<b>S/390 9672 G6 Models</b>		
S/390 9672, Model X17	8346.4	0.000120
S/390 9672, Model X27	7928.6	0.000126
S/390 9672, Model X37	7677.5	0.000130
S/390 9672, Model X47	7511.7	0.000133
S/390 9672, Model X57	7262.8	0.000138
S/390 9672, Model X67	7095.3	0.000141
S/390 9672, Model X77	6762.5	0.000148
S/390 9672, Model X87	6512.0	0.000154
S/390 9672, Model X97	6177.6	0.000162
S/390 9672, Model XX7	6010.5	0.000166
S/390 9672, Model XY7	5759.5	0.000174
S/390 9672, Model XZ7	5592.5	0.000179
S/390 9672, Model Z17	9667.7	0.000103
S/390 9672, Model Z27	9184.8	0.000109
S/390 9672, Model Z37	8893.8	0.000112
S/390 9672, Model Z47	8700.4	0.000115
S/390 9672, Model Z57	8412.2	0.000119
S/390 9672, Model Z67	8217.8	0.000122
S/390 9672, Model Z77	7831.6	0.000128
S/390 9672, Model Z87	7540.1	0.000133
S/390 9672, Model Z97	7249.7	0.000138
S/390 9672, Model ZX7	6959.5	0.000144
S/390 9672, Model ZY7	6765.3	0.000148
S/390 9672, Model ZZ7	6475.1	0.000154
<b>S/390 2003 Models</b>		
S/390 2003, Model 102	177.1	0.005646
S/390 2003, Model 103	265.9	0.003761
S/390 2003, Model 104	443.9	0.002253
S/390 2003, Model 105	620.9	0.001610
S/390 2003, Model 106	798.6	0.001252
S/390 2003, Model 107	1197.9	0.000835
S/390 2003, Model 115	1433.7	0.000697
S/390 2003, Model 116	1769.7	0.000565
S/390 2003 (Partitioned 1-way off a 2003-124)	1307.4	0.000765
S/390 2003, Model 124	1202.8	0.000831
S/390 2003, Model 125	1304.6	0.000767
S/390 2003, Model 126	1610.5	0.000621
S/390 2003, Model 135	1233.0	0.000811
S/390 2003, Model 136	1521.9	0.000657
S/390 2003, Model 146	1433.4	0.000698
S/390 2003, Model 156	1362.6	0.000734
S/390 2003 (Partitioned 1-way off a 2003-1C5)	1618.3	0.000618
S/390 2003 (Partitioned 2-way off a 2003-1C5)	1488.8	0.000672
S/390 2003, Model 1C5	1407.8	0.000710
S/390 2003, Model 203	265.9	0.003761
S/390 2003, Model 204	443.9	0.002253
S/390 2003, Model 205	620.9	0.001611

Table 18. Processors: S/390 9672, 2003, 3000, and 7060 (continued)

Processor Model	Service Units Per Second of Task or SRB Execution Time	Seconds of Task or SRB Execution Time Per Service Unit
S/390 2003, Model 206	798.6	0.001252
S/390 2003, Model 207	1197.9	0.000835
S/390 2003, Model 215	1433.7	0.000697
S/390 2003, Model 216	1792.1	0.000558
S/390 2003 (Partitioned 1-way off a 2003-224)	1313.2	0.000761
S/390 2003, Model 224	1208.2	0.000828
S/390 2003 (Partitioned 1-way off a 2003-225)	1719.9	0.000581
S/390 2003, Model 225	1582.3	0.000632
S/390 2003 (Partitioned 1-way off a 2003-227)	1982.4	0.000504
S/390 2003, Model 227	1863.5	0.000537
S/390 2003, Model 237	1724.7	0.000580
S/390 2003 (Partitioned 1-way off a 2003-246)	1769.7	0.000565
S/390 2003 (Partitioned 2-way off a 2003-246)	1610.5	0.000621
S/390 2003 (Partitioned 3-way off a 2003-246)	1521.9	0.000657
S/390 2003, Model 246	1433.4	0.000698
S/390 2003, Model 247	1645.4	0.000608
S/390 2003, Model 257	1526.4	0.000655
S/390 2003 (Partitioned 1-way off a 2003-2C5)	1618.3	0.000618
S/390 2003 (Partitioned 2-way off a 2003-2C5)	1488.8	0.000672
S/390 2003, Model 2C5	1407.8	0.000710
<b>S/390 3000 Models</b>		
S/390 3000, Model A10	1792.1	0.000558
S/390 3000, Model A20	1582.3	0.000632
<b>S/390 7060 Models</b>		
S/390 7060, Model H30	2881.3	0.000347
S/390 7060, Model H50	5519.1	0.000181
S/390 7060, Model H70	5077.8	0.000197

Table 19. Processors: ES/9000 9021, 9121, and 9221

Processor Model	Service Units Per Second of Task or SRB Execution Time	Seconds of Task or SRB Execution Time Per Service Unit
<b>ES/9000 9021 Processors</b>		
ES/9000 Model 9021-330	1050.0	0.000952
ES/9000 Model 9021-340	1162.3	0.000860
ES/9000 Model 9021-500	1104.2	0.000906
ES/9000 Model 9021-520	2324.9	0.000430
ES/9000 Model 9021-580	1069.4	0.000935
ES/9000 Model 9021-620 single image	1022.9	0.000978
ES/9000 Model 9021-620 partitioned 2-way	1104.2	0.000906
ES/9000 Model 9021-640	2208.7	0.000453
ES/9000 Model 9021-660	2208.7	0.000453
ES/9000 Model 9021-660 partitioned 1-way	2324.9	0.000430
ES/9000 Model 9021-711	3018.1	0.000331
ES/9000 Model 9021-720 single image	941.5	0.001062
ES/9000 Model 9021-720 partitioned 3-way	1069.3	0.000935
ES/9000 Model 9021-740	2139.0	0.000468
ES/9000 Model 9021-820 single image	2046.0	0.000489

Table 19. Processors: ES/9000 9021, 9121, and 9221 (continued)

Processor Model	Service Units Per Second of Task or SRB Execution Time	Seconds of Task or SRB Execution Time Per Service Unit
ES/9000 Model 9021-820 partitioned 2-way	2208.7	0.000453
ES/9000 Model 9021-821	2867.2	0.000349
ES/9000 Model 9021-822 single image	2867.2	0.000349
ES/9000 Model 9021-822 partitioned 1-way	3018.1	0.000331
ES/9000 Model 9021-831	2776.6	0.000360
ES/9000 Model 9021-832 single image	2776.6	0.000360
ES/9000 Model 9021-832 partitioned 2-way	2867.2	0.000349
ES/9000 Model 9021-832 partitioned 1-way	3018.1	0.000331
ES/9000 Model 9021-860	1976.2	0.000506
ES/9000 Model 9021-860 partitioned 2-way	2208.7	0.000453
ES/9000 Model 9021-860 partitioned 3-way	2139.0	0.000468
ES/9000 Model 9021-900 single image	1883.2	0.000531
ES/9000 Model 9021-900 partitioned 3-way	2139.0	0.000468
ES/9000 Model 9021-941	2655.9	0.000377
ES/9000 Model 9021-942 single image	2655.9	0.000377
ES/9000 Model 9021-942 partitioned 2-way	2867.2	0.000349
ES/9000 Model 9021-952 single image	2565.3	0.000390
ES/9000 Model 9021-952 partitioned 3-way	2776.6	0.000360
ES/9000 Model 9021-952 partitioned 2-way	2867.2	0.000349
ES/9000 Model 9021-962 single image	2444.6	0.000409
ES/9000 Model 9021-962 partitioned 3-way	2776.6	0.000360
ES/9000 Model 9021-972 single image	2384.3	0.000419
ES/9000 Model 9021-972 partitioned 4-way	2655.9	0.000377
ES/9000 Model 9021-972 partitioned 3-way	2776.6	0.000360
ES/9000 Model 9021-982 single image	2323.9	0.000430
ES/9000 Model 9021-982 partitioned 4-way	2655.9	0.000377
ES/9000 Model 9021-9X2	2172.7	0.000460
<b>ES/9000 9121 Processors</b>		
ES/9000 Model 9121-180	299.7	0.003337
ES/9000 Model 9121-190	399.0	0.002506
ES/9000 Model 9121-210	591.1	0.001691
ES/9000 Model 9121-260	782.2	0.001278
ES/9000 Model 9121-311	1004.8	0.000995
ES/9000 Model 9121-320	996.5	0.001004
ES/9000 Model 9121-411	1201.3	0.000832
ES/9000 Model 9121-440	743.1	0.001346
ES/9000 Model 9121-480	946.6	0.001056
ES/9000 Model 9121-490	946.6	0.001056
ES/9000 Model 9121-490 partitioned 1-way	996.5	0.001004
ES/9000 Model 9121-511	1536.0	0.000651
ES/9000 Model 9121-521	1141.2	0.000876
ES/9000 Model 9121-522 single image	1141.2	0.000876
ES/9000 Model 9121-522 partitioned 1-way	1201.3	0.000832
ES/9000 Model 9121-570	916.8	0.001091
ES/9000 Model 9121-570 partitioned 1-way	996.5	0.001004
ES/9000 Model 9121-570 partitioned 2-way	946.6	0.001056
ES/9000 Model 9121-610	876.9	0.001140
ES/9000 Model 9121-610 partitioned 2-way	946.6	0.001056
ES/9000 Model 9121-621	1459.2	0.000685
ES/9000 Model 9121-622 single image	1459.2	0.000685
ES/9000 Model 9121-622 partitioned 1-way	1536.0	0.000651
ES/9000 Model 9121-732 single image	1413.1	0.000708
ES/9000 Model 9121-732 partitioned 2-way	1459.2	0.000685

Table 19. Processors: ES/9000 9021, 9121, and 9221 (continued)

Processor Model	Service Units Per Second of Task or SRB Execution Time	Seconds of Task or SRB Execution Time Per Service Unit
ES/9000 Model 9121-732 partitioned 1-way	1536.0	0.000651
ES/9000 Model 9121-742 single image	1351.6	0.000740
ES/9000 Model 9121-742 partitioned 2-way	1459.2	0.000685
<b>ES/9000 9221 Processors</b>		
ES/9000 Model 9221-120	83.5	0.011976
ES/9000 Model 9221-130	145.8	0.006848
ES/9000 Model 9221-150	208.9	0.004794
ES/9000 Model 9221-170	278.1	0.003595
ES/9000 Model 9221-191	481.6	0.002076
ES/9000 Model 9221-200	215.6	0.004637
ES/9000 Model 9221-201	597.3	0.001674
ES/9000 Model 9221-211	719.1	0.001390
ES/9000 Model 9221-221	392.8	0.002545
ES/9000 Model 9221-421	638.4	0.001566

If you plan to use these constants for purposes other than those suggested in this manual, please observe the following limitations:

- Actual customer workloads and performance may vary. For a more exact comparison of processors, see the internal throughput rate (ITR) numbers in *Large Systems Performance Reference (LSPR)*.
- CPU time can vary for different runs of the same job step. One or more of the following factors might cause variations in CPU time: CPU architecture (such as storage buffering), cycle stealing with integrated channels, and the amount of queue searching (see *z/OS MVS System Management Facilities (SMF)*).
- The constants do not account for the effects of PR/SM LPAR mode. For example, a logical 1-way partition in an S/390 9672, Model RX3, has 1090 service units per second, while a 10-way partition on the same machine has 839.3 service units per second.

## Using SMF Task Time

For installations with no prior service data, the task time reported in SMF record Type 4, 5, 30, 34, and 35 records can be converted to service units using the preceding tables.

### Examples of Resource Groups

- To give a department a “dedicated 4381” amount of capacity, you can specify the following resource group:

**Name** 4381-91E

**Description**

Capacity of a 4381

**Capacity Minimum**

309

**Capacity Maximum**

309

- To assure half the capacity of a 9672 Model Y96 system, with 9 physical CPs, to a service class, specify the following:

**Name** HALFSYS3

**Description**

Preserve half of SYS3 (9672 Mod Y96)

**Capacity Maximum**

24129

Since the 9672 Model Y96 has 5362 raw CPU service units per second, with 9 physical CPs,  $0.5 \times 9 \times 5362$  is 24129. The 24129 is captured service units allocated to the address spaces in the resource group.

---

## Appendix C. Notices

This information was developed for products and services offered in the USA.

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  
USA

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.

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

## Notices

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, NY 12601-5400  
USA

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

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

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

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

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

### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

---

## Trademarks

The following terms are trademarks for the IBM Corporation in the United States and/or other countries:

- AD/CYCLE
- APPC
- C/370
- CBIPO
- CICS
- CICS/ESA
- DB2
- DFSMS
- DFSMS/MVS
- Distributed Relational Database Architecture
- DRDA
- ES/9000
- ESCON
- Extended Services
- Hiperbatch
- IBM
- IBMLink
- IMS/ESA
- MVS/ESA
- MVS/SP
- OS/2
- OS/390
- PR/SM
- RACF
- Resource Link
- Resource Measurement Facility
- RMF
- S/390
- SAA
- SOM
- SOMobjects
- Sysplex
- System Object Model
- System/390
- TSO/E
- VTAM
- z/Architecture
- z/OS
- z/OS.e
- zSeries
- z/VM
- 3090

LINUX is a registered trademark of Linus Torvalds and others.

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



---

## Appendix D. Accessibility

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully. The major accessibility features in z/OS enable users to:

- Use assistive technologies such as screen-readers and screen magnifier software
- Operate specific or equivalent features using only the keyboard
- Customize display attributes such as color, contrast, and font size

---

### Using assistive technologies

Assistive technology products, such as screen-readers, function with the user interfaces found in z/OS. Consult the assistive technology documentation for specific information when using it to access z/OS interfaces.

---

### Keyboard navigation of the user interface

Users can access z/OS user interfaces using TSO/E or ISPF. Refer to *z/OS TSO/E Primer*, *z/OS TSO/E User's Guide*, and *z/OS ISPF User's Guide Volume I* for information about accessing TSO/E and ISPF interfaces. These guides describe how to use TSO/E and ISPF, including the use of keyboard shortcuts or function keys (PF keys). Each guide includes the default settings for the PF keys and explains how to modify their functions.



---

## Workload Management Terms

### A

**active service policy.** The service policy that determines workload management processing.

**application environment.** A group of application functions requested by a client that execute in server address spaces.

**application owning region (AOR).** In a CICSplex configuration, a CICS region devoted to running applications.

**automatic control.** One of two distinct methods of managing application environments. Under automatic control, the name of the startup JCL procedure has been defined for an application environment, giving workload management the ability to automatically start server address spaces. Contrast with *manual control*.

### C

**CICSplex.** A configuration of interconnected CICS systems in which each system is dedicated to one of the main elements of the overall workload. See also *application owning region*, *file owning region*, and *terminal owning region*.

**classification rules.** The rules workload management and subsystems use to assign a service class and, optionally, a report class to a work request. A classification rule consists of one or more of work qualifiers such as subsystem type, subsystem instance, userid, accounting information, transaction name, transaction class, source LU, netid, and LU name.

**compatibility mode.** Prior to z/OS V1R3 a mode of processing, in which the IEAIPSxx and IEAICSxx parmlib members determine system resource management. See also *goal mode*.

**control region.** The main storage region that contains the subsystem work manager or subsystem resource manager control program.

**couple data set.** A data set created through the XCF couple data set format utility. The data set is shared by MVS systems in a sysplex. There are several types of couple data sets for different purposes. See also *WLM couple data set*.

**CPU service units.** A measure of the task control block (TCB) execution time multiplied by an SRM constant which is CPU model dependent. See also *unweighted CPU service units per second*, and *service unit*.

### D

**delay monitoring services.** The workload management services that monitor the delays encountered by a work request.

**distributed data facility (DDF).** An optional feature that allows a DB2 application to access data at other DB2s and at remote relational database systems that support IBM's Distributed Relational Database Architecture (DRDA).

**duration.** The length of a service class performance period in service units.

**dynamic alias management.** A service definition option — when enabled, workload management will dynamically reassign parallel access volume aliases to help work meet its goals and to minimize IOS queuing.

### E

**enclave.** A transaction that can span multiple dispatchable units (SRBs and tasks) in one or more address spaces and is reported on and managed as a unit.

**execution velocity.** A service goal naming the rate at which you expect work to be processed for a given service class or a measure of the acceptable processor and storage delays while work is running.

### F

**fold qualifier names.** When defining classification rules, the "Fold qualifier names" option, when set to the default Y, means that the qualifier names will be folded to upper case as soon as you type them in and then press Enter. If you set this option to N, then the qualifier names will remain in the case they are typed in.

### G

**goal mode.** A mode of processing where the active service policy determines system resource management.

### I

**importance level.** the degree of importance of a service goal relative to other service class goals, in five levels: lowest, low, medium, high, highest.

**I/O priority management.** A service definition option — when enabled, I/O priorities will be managed separately from dispatching priorities, according to the goals of the work.

**I/O service units.** A measure of individual data set I/O activity and JES spool reads and writes for all data sets associated with an address space.

**installation.** A particular computing system, including the work it does and the people who manage it, operate it, apply it to problems, service it, and use the results it produces.

**installed service definition.** The service definition residing in the WLM couple data set for WLM.

## L

**logical unit (LU).** In *VTAM*, the source and recipient of data transmissions. Data is transmitted from one logical unit (LU) to another LU. For example, a terminal can be an LU, or a CICS or IMS system can be an LU.

**LU.** Logical unit.

**LU name.** The second level of the source LU name after the “.” for fully qualified names.

**LU 6.2 session.** A session that is initiated by VTAM\* programs on behalf of a logical unit (LU) 6.2 application program, or a session initiated by a remote LU in which the application program specifies that the VTAM programs are to control the session by using the APPCCMD macro instruction.

## M

**manual control.** One of two distinct methods of managing application environments in goal mode. Under manual control, the name of the startup JCL procedure has not been defined for an application environment. The installation must therefore manually start server address spaces when needed. Contrast with *automatic control*.

**masking.** Using a % for a single character replacement in classification rules. See also *wild carding*.

**MVS image.** The system-id of any MVS system included in the sysplex as it appears in SYS1.PARMLIB member.

## P

**performance administration.** The process of defining and adjusting workload management goals and resource groups based on installation business objectives.

**performance block.** A piece of storage containing workload management’s record of execution delay information about work requests.

**performance management.** The process workload management uses to decide how to match resources to work according to performance goals and processing capacity.

**performance period.** A service goal and importance level assigned to a service class for a specific duration. You define performance periods for work that has changing performance requirements as work consumes resources.

**policy.** See *service policy*.

## R

**relational database management system (RDBMS).** A relational database manager that supports SAA.

**report class.** A group of work for which reporting information is collected separately. For example, you can have a report class for information combining two different service classes, or a report class for information on a single transaction.

**resource.** When used as part of a scheduling environment, a resource is an abstract element that can represent an actual physical entity (such as a peripheral device), or an intangible quality (such as a certain time of day). A resource is listed in a scheduling environment along with a required state of ON or OFF. If the corresponding resource state on a given system matches the required state, then the requirement is satisfied for that resource.

**resource group.** An amount of processing capacity across one or more MVS images, assigned to one or more service classes.

## S

**scheduling environment.** A list of resource names along with their required states. If an MVS image satisfies all of the requirements in the scheduling environment associated with a given unit of work, then that unit of work can be assigned to that MVS image. If any of the requirements are *not* satisfied, then that unit of work *cannot* be assigned to that MVS image.

**server address space.** Any address space that does work on behalf of a transaction manager or a resource manager. For example, a server address space could be a CICS AOR, or an IMS control region.

**service administration application.** The online ISPF application used by the service administrator to specify the workload management service definition.

**service class.** A group of work which has the same performance goals, resource requirements, or business importance. For workload management, you assign a service goal and optionally a resource group to a service class.

**service coefficient.** A value that specifies which type of resource consumption should be emphasized in the calculation of service rate. The types of resource consumption are CPU, IOC, MSO, and SRB.

**service definition.** A definition of the workloads and classification rules in an installation. The definition includes workloads, service classes, systems, resource groups, service policies, and classification rules.

**service level administrator.** The user role introduced by workload management whose main task is to make sure overall installation operation is consistent with performance goals and objectives.

**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 policy.** A named set of performance goals workload management uses as a guideline to match resources to work. See also *active service policy*.

**service request block (SRB) service units.** A measure of the SRB execution time for both local and global SRBs, multiplied by an SRM constant which is CPU model dependent.

**service unit.** The amount of service consumed by a work request as calculated by service definition coefficients and CPU, SRB, I/O, and storage service units.

**single-system sysplex.** A sysplex in which only one MVS system is initialized as part of the sysplex. In a single-system sysplex, XCF provides XCF services on the system, but does not provide signalling services between MVS systems. See also *multi-system sysplex*, *XCF-local mode*, and *monoplex*.

**source LU.** A fully qualified two level name separated by a ".", where the first level is the network id and the second is the LU name, OR merely a single LU name. See also *LU name*.

**storage service units.** A measure of the central storage page frames multiplied by 1/50 of the CPU service units. The 1/50 is a scaling factor designed to bring the storage service component in line with the CPU component.

**subsystem instance.** 1) For application environments, a unique combination of subsystem type (as specified in the service definition for an application environment) and subsystem name (as specified by the work manager subsystem when it connects to workload

management). 2) For classification, a work qualifier used to distinguish multiple instances of a subsystem.

**subsystem work manager.** An address space defined in the SYS1.PARMLIB member as SUBSYS=nnn.

## T

**terminal owning region (TOR).** A CICS region devoted to managing the terminal network.

## U

**unweighted CPU service units per second.** The unweighted service units per second of task or SRB execution time. This measure is CPU-model dependent, but is independent of the values of the service coefficients.

## V

**velocity.** A service goal naming the rate at which you expect work to be processed for a given service class or a measure of the acceptable processor and storage delays while work is running.

## W

**wild carding.** The use of an asterisk (\*) as a multiple character replacement in classification rules. See also *masking*.

**WLM couple data set.** A type of couple data set that is created through the XCF couple data set format utility for the WLM function. The data set contains the service definition information.

**workload.** A group of work to be tracked, managed and reported as a unit. Also, a group of service classes.

**workload management mode.** The mode in which workload management manages system resources on an MVS image. Prior to z/OS V1R3 this mode could be either *compatibility mode*, or *goal mode*. Starting with z/OS V1R3, compatibility mode has been removed.

**work qualifier.** An attribute of incoming work. Work qualifiers include: subsystem type, subsystem instance, userid, accounting information, transaction name, transaction class, source LU, netid, and LU name.

**work request.** A piece of work, such as a request for service, a batch job, an APPC, CICS, or IMS transaction, a TSO LOGON, or a TSO command.



---

# Index

## A

- access
  - restricting 138
- accessibility 237
- accounting information
  - nesting 58
  - qualifier 58
- action field 174
- active service policy
  - definition 239
- adjusting velocity goals 45
- administration application
  - definition 240
- alias management, dynamic 85
- APPC/MVS scheduler (ASCH) 56
- application
  - action field 174
  - classification groups 193
  - command line 174
  - commands 174
  - create a group panel 193
  - create workload panel 182
  - customizing keylists 224
  - definition menu 176
  - function keys 175
  - scrollable area 173
  - starting 220, 222, 223
  - using the menu bar 172
  - workload selection list 182
  - workloads 182
- application environment
  - definition 239
- application environments 106
  - authorizing 109
  - CB (component broker) 101
  - changing the definition of 108
  - DB2 101
  - defining 101
  - getting started with 101
  - handling error conditions in 109
  - IWEB 101
  - making changes to servers 108
  - managing 107
  - overview 11
  - selecting server limits for 104
  - SOM 101
  - specifying to workload management 102
  - using operator commands for 107
- application owning region
  - definition 239
- ASCH (APPC/MVS scheduler)
  - overview of work 52
  - work qualifiers supported by 56
- automatic control
  - definition 239
- average response time
  - limit 40
  - limits 40

## B

- batch initiators 21
- business importance
  - definition 8

## C

- capacity, minimum and maximum 37
- CB (component broker)
  - application environments 101
  - overview of work 52
  - work qualifiers supported by 56
- changing goal types in performance periods 47
- channel subsystem priority queuing 162
- CICS (customer information control system)
  - overview of work 52
  - work qualifiers supported by 56
- CICSplex
  - definition 239
- classification
  - defining rules 49
  - defining the order 68
  - inheritance 68
  - nesting 51, 58
  - qualifiers supported 55
  - supporting subsystems 51
- classification rules
  - creating for a subsystem type 188
  - definition 239
- collection name
  - qualifier 59
- compatibility mode
  - definition 239
- component broker (CB) 56
- connection type
  - qualifier 59
- control region
  - definition 239
- correlation information
  - qualifier 59
- couple data set
  - allocating 145
  - calculating the size of 145
  - definition 239
  - increasing the size of 148
  - installing service definition 207
  - restricting access to 138
  - SEXTXCF command 149
  - updating COUPLExx 149
- COUPLExx member
  - DATA keyword 150
- coupling facility
  - defining a structure 155, 159
- CPSM environment 20
- CPU protection 40, 90
- CPU service units
  - definition 239, 241

- customer information control system (CICS) 56

## D

- DB2
  - application environments 101
  - overview of work 52
  - work qualifiers supported by 56
- DB2 distributed data facility environment 20
- DDF (distributed data facility)
  - overview of work 52
  - work qualifiers supported by 56
- DDF environment 20
- defining application environments 101
- defining scheduling environments 113
- defining velocity goals 44
- disability 237
- discretionary goal 40
- discretionary goal management
  - migration considerations 152
- discretionary goals
  - using 46
- distributed data facility (DDF) 56
- documents, licensed xiii
- duration
  - definition 41
  - in a performance period 41
- dynamic alias management 85
  - definition 239
  - migration considerations 152
- dynamic channel path management 161
  - coupling facility structure 159

## E

- enclave
  - definition 239
- Enterprise Storage Server 85
- execution velocity
  - calculation 84
  - definition 239
  - with I/O priority management 84
- exempting
  - from transaction server management 91

## F

- fold qualifier names
  - definition 239
  - option explained 189
- functionality level 140

## G

- goal
  - definition 39

- goal mode
  - definition 239
- goal types 18
- goals, performance
  - defining 42
  - definition 7

## H

- heterogeneous report class 81
- homogeneous report class 81

## I

- I/O priority
  - management 84
- I/O priority management
  - definition 240
- I/O priority queueing
  - defining 27
  - dynamic 27
- I/O service units
  - definition 240
- importance
  - definition 239
- importance levels
  - in performance periods 47
- IMS (information management system)
  - overview of work 52
  - work qualifiers supported by 56
- information management system (IMS) 56
- inheritance
  - in classification rules 68
- initiators, batch 21
- Install Definition Utility
  - member IWMINSTL 133
  - member IWMSSDEF 133
- installation exit
  - IWMAREX1 218
  - IWMAREX2 220
- installed service definition
  - definition 240
- intelligent resource director
  - channel subsystem priority queuing 162
  - coupling facility structure 159
  - dynamic channel path management 161
  - LPAR CPUManagement 161
  - making it work 164
- IRD
  - coupling facility structure 159
- IWEB
  - application environments 101
  - overview of work 52
  - work qualifiers supported by 56
- IWMAM040 message 213
- IWMAM041 message 213
- IWMAM042 message 213
- IWMAM043 message 213
- IWMAM044 message 213
- IWMAM046 message 213
- IWMAM050 message 214
- IWMAM051 message 214
- IWMAM052 message 214

- IWMAM054 message 214
- IWMAM055 message 214
- IWMAM058 message 214
- IWMAM072 message 214
- IWMAM077 message 215
- IWMAREX1 installation exit 218
- IWMAREX2 installation exit 220
- IWMINSTL 133
- IWMSSDEF 133

## J

- JES (job entry subsystem)
  - overview of work 52
  - work qualifiers supported by 56
- JES2 (job entry subsystem 2)
  - overview of work 52
  - work qualifiers supported by 56
- JES2 batch initiators 21
- JES3 (job entry subsystem 3)
  - overview of work 52
  - work qualifiers supported by 56
- JES3 batch initiators 21
- job entry subsystem (JES) 56

## K

- keyboard 237

## L

- licensed documents xiii
- LookAt message retrieval tool xiii
- LPAR clustering
  - coupling facility structure 159
- LPAR CPU management 161
- LPAR weight management 161
- LSFM
  - overview of work 52
  - work qualifiers supported by 56
- LU 6.2 session
  - definition 240
- LU name
  - definition 240
  - qualifier 59

## M

- managing resource states 117
- manual control
  - definition 240
- masking
  - definition 240
- masking notation
  - examples 74
- maximum capacity 37
- menu bar
  - definition menu 178
  - on selection list 173
- message retrieval tool, LookAt xiii
- migration
  - creating service definition for the first time 131
  - discretionary goal management considerations 152

- migration (*continued*)
  - dynamic alias management
    - considerations 152
  - from pre-version 5 131
  - multisystem enclaves
    - considerations 152
    - overview 131
    - velocity considerations 151
    - with an existing service
      - definition 135
  - minimum capacity 37
  - mixed-release sysplex 140
  - mode
    - definition 241
  - MQ (MQSeries Workflow)
    - overview of work 52
    - work qualifiers supported by 56
  - MQSeries Workflow (MQ) 56
  - multiple periods 18
  - multisystem enclaves
    - coupling facility structure 155
    - migration considerations 152

## N

- nesting
  - accounting information 69
  - subsystem parameter 69
- netid
  - qualifier 59
- non z/OS partition CPU
  - management 22, 160, 167
- notation
  - start position 73
- notepad
  - updating 27
  - using 27
- Notices 233

## O

- OMVS (OS/390 UNIX System Services)
  - overview of work 52
  - work qualifiers supported by 56
- OS/390 UNIX System Services (OMVS) 56

## P

- package name
  - qualifier 60
- parallel access volume 85
- partitioned data set
  - restricting access to 138
- PERFORM
  - qualifier 60
- performance administration
  - definition 4
- performance goal
  - definition 39
- performance goals
  - defining 42
  - definition 7
- performance group
  - qualifier 60

- performance management
  - definition 5
- performance period
  - definition 240
  - maximum number 41
- performance periods
  - using 47
  - using importance levels in 47
- periods, multiple 18
- plan name
  - qualifier 60
- policy overrides
  - defining 32, 186
  - examples 32
  - resource group association 32
  - resource group capacity 32
- printing
  - service definition 211
  - service policy 211
- priority
  - qualifier 61
- procedure name
  - qualifier 61
- process name
  - qualifier 61
- processor model
  - related to service units 225
  - related to task/SRB execution time 225
- protection
  - CPU 90
  - storage 89
- protection options for critical work
  - defining 89

## Q

- qualifier
  - accounting information 58
  - collection name 59
  - connection type 59
  - correlation information 59
  - grouping 75
  - LU name 59
  - nesting 69
  - netid 59
  - package name 60
  - PERFORM 60
  - performance group 60
  - plan name 60
  - priority 61
  - procedure name 61
  - process name 61
  - scheduling environment name 62
  - subsystem collection name 62
  - subsystem instance 62
  - subsystem parameter 63
  - sysplex name 64
  - system name 64
  - transaction class 65
  - transaction name 66
  - userid 67
- qualifier names
  - folding 189, 239
- qualifiers
  - definition 49

## R

- RACF
  - restricting access to WLM service
    - definition 138
- report class
  - assigning 81
  - definition 81, 240
  - heterogeneous 81
  - inhomogeneous 81
- resource
  - definition 240
- resource group
  - defining 37
  - definition 240
  - limitations 38
  - maximum capacity 37
  - minimum capacity 37
  - overview 9
  - removing a service class from 187
- resource requirements 113
  - definition 7
- resource states, managing 117
- response time goals, system
  - determining 43
- response time with percentile
  - limit 40
- restricting access to WLM service
  - definition 138

## S

- SAF
  - restricting access to WLM service
    - definition 138
- samples included in velocity goals 45
- scheduling environment
  - definition 240
- scheduling environment name
  - qualifier 62
- scheduling environments
  - associating with incoming work 121
  - defining 113
  - getting started with 113
  - managing resource states 117
  - overview 13
  - specifying to workload
    - management 114
  - using special characters 114
- security server
  - restricting access to WLM service
    - definition 138
- selection lists
  - in classification rules 192
  - status line 173
- server management, transaction 91
- service class
  - definition 39, 241
  - maximum number 39
  - SYSOTHER 73
  - SYSSTC 72
  - SYSTEM 71
  - system-provided 71
- service coefficient
  - defaults 84
  - definition 241

- service coefficients
  - defining 83
- service definition
  - base 14
  - contents 27
  - creating for the first time 131
  - defining 27
  - definition 6, 241
  - hierarchy 14
  - installing 133, 207
  - printing 211
  - printing as GML 211
  - restricting access to 138
  - storing in MVS PDS 171
  - storing in WLM couple data set 171
  - worksheets for defining 125
- service level agreement (SLA)
  - definition 241
- service policy 32
  - activating 210
  - create service policy panel 181
  - defining 7, 31
  - defining overrides 186
  - definition 241
  - in service definition 31
  - overview 14
  - printing 211
- service unit
  - definition 241
  - related to processor model 225
  - related to task/SRB execution time 225
- SETXCF command
  - example 149
- shark 85
- shortcut keys 237
- SOM
  - application environments 101
  - overview of work 52
  - work qualifiers supported by 56
- source LU
  - definition 241
- SRB (service request block)
  - execution time
    - related to processor model 225
    - related to service units 225
- SRB service units
  - definition 241
- start position
  - using 73
- started task control (STC) 56
- started tasks
  - defining goals for 46
  - servers 46
- started tasks (STC)
  - defining classification rules 76
  - defining service classes 76
- STC (started task control)
  - overview of work 52
  - work qualifiers supported by 56
- STC (started tasks)
  - defining classification rules 76
  - defining service classes 76
- storage protection 89
- storage service units
  - definition 241

- substring notation
  - examples 74
- subsystem instance
  - definition 241
- subsystem collection name
  - qualifier 62
- subsystem instance
  - qualifier 62
- subsystem parameter
  - qualifier 63
- subsystem type
  - creating 192
  - deleting 193
  - modify rules for 188
- subsystem types
  - in classification 55
- SYS\_ 114
- SYS1.SAMPLIB
  - member IWMINSTL 133
  - member IWMSSDEF 133
- SYSH
  - overview of work 52
  - work qualifiers supported by 56
- SYSH subsystem 23
- SYSOTHER
  - service class 73
- sysplex couple data set
  - upgrading 133
- sysplex name
  - qualifier 64
- SYSSTC
  - service class 72
- SYSTEM
  - service class 71
- system name
  - qualifier 64
- system response time goals
  - determining 43

## T

- terminal owning region (TOR)
  - definition 241
- transaction class
  - qualifier 65
- transaction name
  - qualifier 66
- transaction server management,
  - exemption from 91
- TSO
  - overview of work 52
  - work qualifiers supported by 56

## U

- userid
  - qualifier 67

## V

- velocity
  - definition 241
  - formula 42
  - limit 40
  - migration considerations 151

- velocity goals
  - adjusting 45
  - defining 44
  - samples included in 45

## W

- wild carding
  - definition 241
- wildcard notation 74
  - examples 75
- WLM couple data set
  - allocating 145
  - calculating the size of 145
  - definition 241
  - increasing the size of 148
  - installing service definition 207
  - restricting access to 138
  - SEXTXCF command 149
  - updating COUPLExx 149
- work environments 17
- work qualifier
  - definition 241
- work request
  - definition 241
- workload
  - definition 241
  - in service definition 35
- workload balancing 19
  - definition 5
- worksheets for defining a service
  - definition 125

---

# Readers' Comments — We'd Like to Hear from You

z/OS

MVS Planning: Workload Management

Publication No. SA22-7602-04

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>				

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>				
Complete	<input type="checkbox"/>				
Easy to find	<input type="checkbox"/>				
Easy to understand	<input type="checkbox"/>				
Well organized	<input type="checkbox"/>				
Applicable to your tasks	<input type="checkbox"/>				

Please tell us how we can improve this book:

Thank you for your responses. May we contact you?  Yes  No

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

\_\_\_\_\_  
Name

\_\_\_\_\_  
Address

\_\_\_\_\_  
Company or Organization

\_\_\_\_\_  
Phone No.



Fold and Tape

Please do not staple

Fold and Tape



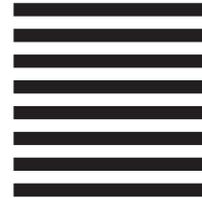
NO POSTAGE  
NECESSARY  
IF MAILED IN THE  
UNITED STATES

# BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Deutschland Entwicklung GmbH  
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, 5655-G52

Printed in U.S.A.

SA22-7602-04



Spine information:



z/OS

**z/OS V1R4.0 MVS Planning: Workload  
Management**

SA22-7602-04