


ibm.com



e-business



z/OS V1R9 WLM and RMF Enhancements



Redbooks

International Technical Support Organization

© Copyright IBM Corp. 2007. All rights reserved.

Trademarks



eNetwork	DFSMS/MVS	IMS	RMF
geoManager	DFSMSdfp	IMS/ESA	RS/6000
AD/Cycle	DFSMSdss	IP PrintWay	S/390
ADSTAR	DFSMSshm	IPDS	S/390 Parallel Enterprise Server
AFP	DFSMSrmm	Language Environment	SecureWay
APL2	DFSORT	Multiprise	StorWatch
APPN	Enterprise System 3090	MQSeries	Sysplex Timer
BookManger	Enterprise System 4381	MVS/ESA	System/390
BookMaster	Enterprise System 9000	Network Station	System REXX
C/370	ES/3090	NetSpool	SystemView
CallPath	ES/4381	OfficeVision/MVS	SOM
CICS	ES/9000	Open Class	SOMobjects
CICS/ESA	ESA/390	OpenEdition	SP
CICS/MVS	ESCON	OS/2	VisualAge
CICSPlex	First Failure Support Technology	OS/390	VisualGen
COBOL/370	FlowMark	Parallel Sysplex	VisualLift
DataPropagator	FFST	Print Services Facility	VTAM
DisplayWrite	GDDM	PrintWay	WebSphere
DB2	ImagePlus	ProductPac	3090
DB2 Universal Database	Intelligent Miner	PR/SM	3890/XP
DFSMS	IBM	QMFr	z/OS
	IBM System z	RACF	z/OS.e

Domino (Lotus Development Corporation)

DFS (Transarc Corporation)

Java (Sun Microsystems, Inc.)

Lotus (Lotus Development Corporation)

Tivoli (Tivoli Systems Inc.)

Tivoli Management Framework

(Tivoli Systems Inc.)

Tivoli Manger (Tivoli Systems Inc.)

UNIX (X/Open Company Limited)

Windows (Microsoft Corporation)

Windows NT (Microsoft Corporation)



© Copyright IBM Corp. 2007. All rights reserved.

z/OS V1R9 WLM Enhancements



- ❑ Promote jobs that are cancelled
 - Increase dispatching priority when cancelled
- ❑ Start a minimum number of servers
- ❑ Enhancements for blocked workloads
- ❑ zAAP support for routing



© Copyright IBM Corp. 2007. All rights reserved.

Promote Jobs that are Cancelled



- ❑ Cancelling an address space sometimes takes a long time to complete because the job is being pre-empted by higher priority work
 - Job does not get dispatched
 - Majority of cancel processing runs in the address space being cancelled
 - Runs at the dispatch priority of the address space
- ❑ With z/OS V1R9:
 - Increase dispatch priority of cancelled address space
 - SYSEVENT BRINGIN service deleted
 - If used, abend code x'15F', reason code 4
- ❑ New service used SYSEVENT CANCEL
 - Can cause a swap in and creates a dispatch change



© Copyright IBM Corp. 2007. All rights reserved.

Start a Minimum Number of Servers



- ❑ If a WLM managed server does a lot of processing before it connects to WLM, the startup of the minimum amount of servers can last a long time
- ❑ A new parameter is specified, WLM starts the minimum amount of servers without waiting on the connect status of the previous started servers.
 - In such environments the startup of the minimum servers can be accelerated
- ❑ A new parameter on IWM4SLI macro service



© Copyright IBM Corp. 2007. All rights reserved.

IWM4SLI Macro Service



- ❑ **AE_SERVERMAX** parameter - establish maximum number
 - Useful for applications such as MQSeries and workflow that connects to backend applications supporting a limited number of parallel connections
- ❑ Use the **AE_SERVERMIN** parameter to establish a minimum number as this allows an application to keep a number of servers active, even during low utilization periods
- ❑ Specify **AE_SPREADADMIN=YES** to ensure that the defined minimum number of servers are distributed evenly across all of the service classes used to execute work requests in the application environment



© Copyright IBM Corp. 2007. All rights reserved.

New Parameters with z/OS V1R9

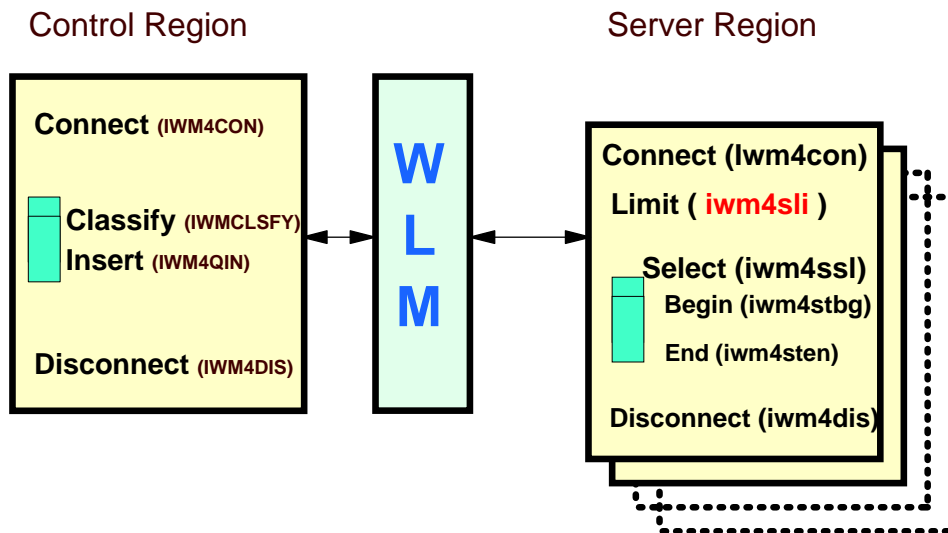


- ❑ When `AE_SERVERMIN=ae_servermin` is specified
 - Optional parameter, which indicates whether WLM starts the minimum number of servers one by one or in parallel
 - The default is `START_MINIMUM=SERIAL`
- ❑ `START_MINIMUM=SERIAL` - The server tasks specified in `AE_SERVERMIN` are started one by one
 - This means the next server will only be started if the previous server has connected to WLM
- ❑ `START_MINIMUM=PARALLEL` - The server tasks specified in `AE_SERVERMIN` are started in parallel
 - This means WLM starts additional servers even when the previous servers have not connected to WLM



© Copyright IBM Corp. 2007. All rights reserved.

Server Regions Connecting to WLM

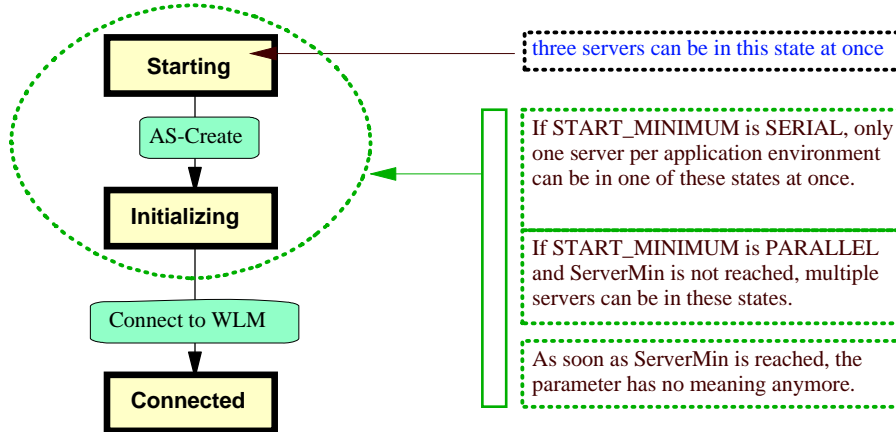


© Copyright IBM Corp. 2007. All rights reserved.

Server Regions Establishing ServerMIN



Server States



© Copyright IBM Corp. 2007. All rights reserved.

ibm.com



WLM Blocked Workloads



Redbooks

International Technical Support Organization

© Copyright IBM Corp. 2007. All rights reserved.

Blocked Workload Analysis



Low priority work

- Requests a resource
- Gets suspended in favour of more important work
- Remains undispached over long periods
- Cannot free up the resource

High priority work

- Requests the same resource
- Gets blocked due to the resources unavailability

z/OS V1R9

Provide limited access to the CPU for low priority work
Promote dispatching priority for short periods

© Copyright IBM Corp. 2007. All rights reserved.

WLM Blocked Workloads Enhancement



- During periods of 100% CP utilization, it is possible that discretionary workloads are not dispatched for execution - (low dispatch priorities)
- These discretionary workloads may obtain and hold serially reusable resources required by other workloads
 - May block progress of higher dispatch priority workloads
- In z/OS V1R9, now specify that any address space and any enclaves that have work that is ready to run but does not get CP service within a certain time interval can be temporarily promoted to a higher dispatch priority
 - RMF supports this function by reporting relevant measurements

Blocked Workload Support



- ❑ When low priority work obtains a resource that high importance work is waiting for
 - But is blocked by a large CP consumer of medium importance (priority inversion)
 - The high priority work is now in effect blocked by the medium priority work
- ❑ z/OS V1R9 - SRM can promote low priority work holding a resource to a higher dispatch priority to allow sufficient access to a CP in order to release the held resource
 - Eliminates need to manually increase the priority of low priority work holding a resource to accelerate resource contention issue



© Copyright IBM Corp. 2007. All rights reserved.

IEAOPTxx CPU Management Constants



- ❑ **BLWLINTHD**
 - Specifies the threshold time interval for which a blocking address space or enclave must wait before being considered for promotion
 - Minimum is 5 seconds - Maximum is 65535 seconds
 - Default is 60 seconds
- ❑ **BLWLTRPCT**
 - Specifies how much of the CP capacity is to be used to promote blocking workloads
 - Minimum is 0 units - Maximum is 200 units (=20%) (0% implies trickle should not occur)
 - Default is 5 (=0.5%) - max is 200 (=20%)



© Copyright IBM Corp. 2007. All rights reserved.

RMF Support for Blocked Workloads



- ❑ RMF enhances the Workload Activity report and the CPU Activity report to provide information about the handling of blocked workload
- ❑ RMF provides new overview conditions for the Postprocessor based on SMF record 72-3
- ❑ This new functionality is available as SPE and needs to be installed as APAR OA18244 and PTFs will be provided for z/OS V1.8 and V1.7 RMF
- ❑ RMF reporting for WLM's blocked workload support
 - By collecting additional data in SMF records
 - By extending the RMF Postprocessor reports
 - By providing new metrics for RMF PM and the Data Portal



© Copyright IBM Corp. 2007. All rights reserved.

RMF CPU Activity Report



- ❑ Use to adjust the BLWLTRPCT parameter
- ❑ Waiters greater than 0, the system has blocked work
 - Indicates a need to increase BLWLTRPCT
- ❑ Lists the number of blocked dispatchable work units that are eligible for dispatching priority promotion
- ❑ Lists the defined average promotion rate and the percentage used during the measurement interval
 - A value below 100% indicates that not all blocking workloads could be promoted
- ❑ List average and peak number of address spaces and enclaves found blocked

```
                                C P U   A C T I V I T Y
                                z/OS V1R9
CPU  2094  MODEL  714  H/W MODEL  S18

BLOCKED WORKLOAD ANALYSIS

OPT PARAMETERS: BLWLTRPCT (%)  0.5  PROMOTE RATE:  DEFINED  50000  WAITERS FOR PROMOTE:  AVG  0.001
                  BLWLINTHD    60                USED (%)   95                PEAK    15
```



© Copyright IBM Corp. 2007. All rights reserved.

RMF Blocked Workload Analysis



- ❑ **DEFINED** - Average number of blocked dispatchable work units which may get promoted in their dispatching priority per second. This value is derived from OPT parameter BLWLTRPCT.
- ❑ **USED (%)** - The utilization percentage of the defined promote rate during the reporting interval.
- ❑ **AVG** - Average number of address spaces and enclaves found blocked according to BLWLINTHD during the report interval.
- ❑ **PEAK** - Highest number of address spaces and enclaves found blocked during the report interval.



© Copyright IBM Corp. 2007. All rights reserved.

SMF Record Types - Blocked Workloads



- ❑ **SMF Record 99 Subtype 1** shows the number of address spaces or enclaves waiting longer than the threshold
 - The field name is SMF99_CCTRCWTR and shows the number of waiters

System state information section for SMF record type 99

Offsets	Name	Length	Format	Description
28 120	SMF99_CCTINTHD	2	Binary	OPT parameter BLWLINTHD starvation threshold
28 122	SMF99_CCTTRPCT	2	Binary	OPT parameter BLWLTRPCT for percentage of CP trickling
31 138	SMF99_CCTRCWTR	4	Binary	Number of address spaces or enclaves waiting longer than the threshold



© Copyright IBM Corp. 2007. All rights reserved.

RMF Workload Activity Report



RMF's WLMGL report, the following can be analyzed:

- Identify service classes running work units at a promoted dispatching priority and amount of CPU time
- PROMOTED shows the CPU time in seconds that transactions in this group were running at a promoted dispatching priority

```
                                W O R K L O A D   A C T I V I T Y
-TRANSACTIONS-  TRANS-TIME HHH.MM.SS.TTT  --DASD I/O--  ---SERVICE---  --SERVICE TIMES--  ----APPL%----  -----STORAGE-----
AVG    10.07    ACTUAL          1.00.895  SSCHRT  20.8  IOC    96199  CPU    1.400  CP     0.20  AVG    980.42
MPL    10.07    EXECUTION        59.616  RESP   2.8  CPU   266554  SRB    0.200  AAPCP  0.00  TOTAL  9868.59
ENDED   1      QUEUED          1.278  CONN   0.2  MSO    9131  RCT    0.000  IIPCP  0.00  SHARED 127.67
END/S   0.00    R/S AFFIN          0      DISC   0.0  SRB   39610  IIT    0.200
#SWAPS  30      INELIGIBLE         0      Q+PEND 0.5  TOT   411494  HST    0.000  AAP    0.00  --PAGE-IN RATES--
EXCTD   0      CONVERSION         0      IOSQ   2.0  /SEC   457  AAP    0.000  IIP    N/A  SINGLE  0.0
AVG ENC  0.00    STD DEV           0
REM ENC  0.00
MS ENC  0.00
                                ABSRPTN  45
                                TRX SERV 45  PROMOTED 0.333
                                HSP      0.0
```



© Copyright IBM Corp. 2007. All rights reserved.

RMF CPU Activity Report



Field WAITERS FOR PROMOTE shows the "number of waiters"

- This report can be used to adjust the BLWLTRPCT parameter
- As long as the number of waiters is greater than 0, the system has blocked work, indicating a need to increase BLWLTRPCT

If you get a problem with blocked work holding resources for too long although you see no waiters in the RMF data

- Consider a decrease of BLWLINTHD



© Copyright IBM Corp. 2007. All rights reserved.

SMF record type 70 subtype 1 (CPU Activity) CPU control section



Offset	Name	Length	Format	Description
5 x5	SMF70STF	1	Binary	Flag. BIT 0 to 4 meaning not changed. BIT 5 = OPT parameter BLWLTRPCT changed. BIT 6 = OPT parameter BLWLINTHD changed BIT 7 = reserved
94 x5E		2		Reserved
96 x60	SMF70PMI	4	Binary	Accumulated number of blocked dispatchable units per second that may get promoted in their dispatch priority. To get the average promote event rate, divide SMF70PMI by SMF70SAM.
100 x64	SMF70PMU	4	Binary	Number of blocked dispatchable units being promoted during the interval
104 x68	SMF70PMW	4	Binary	Accumulated number of address spaces and enclaves being blocked during the interval. To get the average number of waiters for promote, divide SMF70PMW by SMF70SAM.
108 x6C	SMF70PMP	4	Binary	Maximum number of address spaces and enclaves found being blocked during the interval
112 x70	SMF70PMT	2	Binary	1/1000s of a CP for promote slices (OPT parameter BLWLTRPCT)
114 x72	SMF70PML	2	Binary	Swapped-in starvation threshold When an address space or enclave has not received CPU service within this time interval, it is considered blocked (OPT parameter BLWLINTHD)



© Copyright IBM Corp. 2007. All rights reserved.

Assist Processor Routing Changes



- ❑ For zAAPs and zLIPS - new information returned
 - IWMSRSRS FUNCTION=SELECT or IWMSRSRS FUNCTION=SPECIFIC
 - Returns a list of registered servers in a sysplex along with 4 weight values for each server (3 individual, 1 combined)
 - tells caller relative number of requests to send to server
 - IWM4SRSC
 - Returns the weight(s) for one specific server identified by its STOKEN
 - IWMWSYSQ EXTENDED_DATA=YES
 - Returns capacity information for all processor types of all systems in the sysplex



© Copyright IBM Corp. 2007. All rights reserved.

z/OS V1R9 IWMSRSRS Returns



- ❑ In previous releases:
 - Only one weight (SYSR_WEIGHT) is returned and it is based only on regular CP capacity
 - A system was NOT returned if it had less than 5% of displaceable capacity at the selected importance level
- ❑ With z/OS V1R9
 - A system is NOT returned if it has NO displaceable capacity for regular CPs at the selected importance level
 - As a result potentially more systems may be returned
 - If a system has no displaceable capacity for an assist processor it is returned and missing capacity is reflected in the weight



© Copyright IBM Corp. 2007. All rights reserved.

Calculating the System Weights



Step 1

- ❑ Finding the importance level where at least 5% of displaceable capacity is available
 - Result: This is importance level 5
 - System B and C meet the criteria
 - System A does not meet the criteria

Selected Importance Level	=	5
Total capacity at level 5	40+390+700 =	1130
CPU System Weight for System A	40*64/1130 =	2
CPU System Weight for System B	390*64/1130 =	22
CPU System Weight for System C	700*64/1130 =	40

Level	System A		System B		System C	
	SUs	%	SUs	%	SUs	%
0 (System)	2000	100	3000	100	2000	100
1	1800	90	2700	90	1700	85
2	1400	70	2100	70	1500	75
3	600	30	1800	60	1400	70
4	160	8	900	30	900	45
5	40	2	390	22	700	40
6 (Disc)	0		120	4	60	3
7 (Free)	0		30	1	0	



© Copyright IBM Corp. 2007. All rights reserved.

IWMSRSRS Details



- ❑ Prior to z/OS V1R9, system A would have been removed from the list of eligible systems. Now with the new changes, system A is still considered because the consumption is NOT zero
- ❑ **SYSR_WEIGHT** now is the combined weight of all available processor resources
 - The individual weights have been added:
 - **SYSR_CPU_WEIGHT** (= to **SYSR_WEIGHT** of previous releases)
 - **SYSR_ZAAP_WEIGHT**
 - **SYSR_ZIIP_WEIGHT**



© Copyright IBM Corp. 2007. All rights reserved.

IWMSRSRS Details for New Support



- ❑ **Step 2**
 - Calculate the system weights for the assist processors
 - Systems which have been excluded in step 1 are no longer considered
 - In this step no system is excluded
 - Results are shown on the right
- ❑ **Step 3**
 - Calculate the server weight
 - Use the proportion of the work using the processor types to calculate a system weight for the server
 - Scale the resulting weight to 64
- ❑ **Step 4**
 - Include other weight factors (not depicted)

Processor Weights - Step 2

ProcType	Importance	System A	System B	System C
Regular CPs	5	2	22	40
zAAPs	4	15	51	0
zIIPs	7	20	15	29

Server Processor Usage - Step 3

No zAAPs

Server X	System A			System B			System C		
	CP	zAAP	zIIP	CP	zAAP	zIIP	CP	zAAP	zIIP
Service	200	200	200	100	300	100	200	0	50
Portion	33	33	33	20	60	20	80	0	20

Server System Weights - Step 4

	System A	System B	System C
Calculation	$\frac{1}{100} \cdot \sum_i [\text{Portion}(i) \cdot \text{ProcType}(i)]$		
System weight (unscaled)	12	38	37
Scaling	$64 \cdot \frac{\text{UnscaledWeight}(i)}{\sum \text{UnscaledWeight}(i)}$		
Resulting System Weight	9	28	27



© Copyright IBM Corp. 2007. All rights reserved.

Step 4 Server Weight Calculation



- Calculate the server weights
- Divide system weights by number of servers on system
- (only for function=SPECIFIC): include other server performance factors
- In case one of the servers has a mixed weight = 0: return only 1 server (and undivided weight): the server with maximum performance

Server Weights - Step 4

System C

Server	Server C_1				Server C_2				Server C_3			
	W	CPU W	zAAP W	zIIP W	W	CPU W	zAAP W	zIIP W	W	CPU W	zAAP W	zIIP W
System Weight	27	40	0	29	27	40	0	29	27	40	0	29
# of Servers	3											
Divided weights	9	13	0	10	9	13	0	10	9	13	0	10
Server performance factors (PI, Queue)	100%				25%				66%			
Server Weights	9	13	0	10	4	6	0	5	6	8	0	6



© Copyright IBM Corp. 2007. All rights reserved.

IWMSRSRS Additional Enhancements



- Restriction that no more than 96 servers per system could be returned is now relieved
 - Now up to 300 servers per system will be returned
 - Also valid for IWMSRSRS FUNCTION=QUERY
- This support is available for z/OS V1R6 and above via APAR OA18531
- The C interface to IWMSRSRS IWMDNSRV has also been extended to return the new weights



© Copyright IBM Corp. 2007. All rights reserved.

IWM4SRSC Service



- ❑ IWM4SRSC service provides information of how well a server is suitable to receive work
- ❑ The IWM4SRSC service allows to check a specific server before routing work to it from WLM
 - Information obtained can be used for making balanced routing decisions
 - For example, with programs like Sysplex Distributor, and its exploiters, for example, TCP/IP
- ❑ A weight is calculated based on six factors:
 - A combination of the three processor weights (CPU weight, ZAAP weight and ZIIP weight) and their respective consumed service units



© Copyright IBM Corp. 2007. All rights reserved.



IWM4SRSC Service






IWM4SRSC

Service for non-registered Servers
with LPAR Scope

Work Requests



Routing Recommendation



© Copyright IBM Corp. 2007. All rights reserved.

IWM4SRSC Service New Support



- ❑ Calculation of processor weight already exists today for regular CPUs - now extended to assist processors
- ❑ Resulting server weight is calculated based on the proportion of usage by the server (service class)
- ❑ New server weight, and individual weights are returned
- ❑ New CPUWEIGHT corresponds to the “old” WEIGHT from previous releases and other new parameters:
 - ZAAPWEIGHT and ZIIPWEIGHT which return the assist processor weights
 - CPUPROPORTION, ZAAPPROPORTION, and ZIIPPROPORTION are calculated based on the consumed service on each processor type

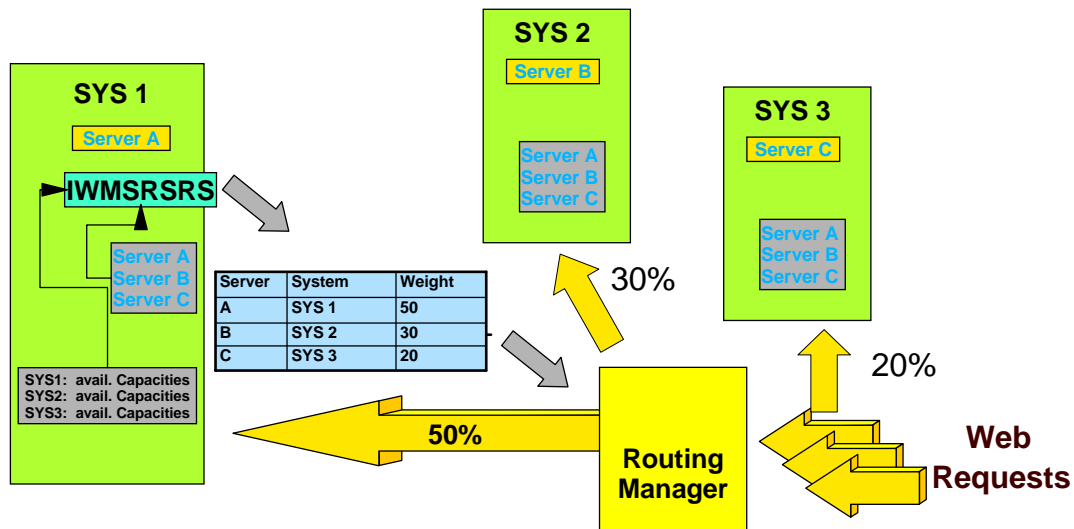


© Copyright IBM Corp. 2007. All rights reserved.

Obtain Routing Recommendation



Routing manager wants to distribute the incoming requests among the web servers in a balanced way




© Copyright IBM Corp. 2007. All rights reserved.

ibm.com



e-business



RMF Enhancements with z/OS V1R9



Redbooks

International Technical Support Organization

© Copyright IBM Corp. 2007. All rights reserved.

z/OS V1R9 RMF Enhancements




- ❑ With z/OS V1.9 SMF 74.2 data can be visualized by the Monitor III Data Portal as well
 - In the future, the OMEGAMON XE data collector will also exploit the XCF data provided by the Portal
- ❑ A new RMF Spreadsheet Reporter macro allows now to analyze SMF 74.2 data graphically
- ❑ WLM implemented new algorithms to help low priority workloads to get sufficient access to the CP resource
 - RMF reports now promotion statistics for low priority workloads as part of the Postprocessor CPU Activity Report

ibm.com



e-business



RMF Monitor III Data Portal for z/OS



Redbooks

International Technical Support Organization

© Copyright IBM Corp. 2007. All rights reserved.

RMF Monitor III Data Portal for z/OS



- ❑ Instead of using the RMF PM, you can access RMF Monitor III data using DDS via any Web browser
 - DDS behaves like an HTTP Server
 - Therefore, if RMF is currently running and you have not started DDS
- ❑ You must start DDS as follows as an example:
 - F RMF,DDS
 - From a Web browser
 - <http://<hostname>:8803>

RMF Monitor III Data Portal for z/OS



RMF Monitor III Data Portal for z/OS

Explore
Overview
My View
Home



© IBM Corporation, 1998-2007

Important notes:

- When using this application you will be prompted to login to the Sysplex with a valid userid and password.
- This application requires Javascript to be enabled
- For some functions (such as "My View") you must allow your browser to store cookies.
- This application has been successfully tested with:
 - Microsoft Internet Explorer Version 6.0 or higher for Windows (Recommended: 7.0 or higher)
 - The Mozilla Suite or Firefox Browser Version 1.0 or higher from www.mozilla.org for various platforms. (Recommended: 2.0 or higher)
 - Netscape Browser Version 7.0 or higher for various platforms (Recommended: 8.0 or higher)

FAQ

RMF



© Copyright IBM Corp. 2007. All rights reserved.

New Full RMF Reports with z/OS V1R9



RMF Monitor III Data Portal for z/OS

Explore
Overview
My View
Home



New

Full RMF Reports:

CACHDET	CACHSUM	CFACT	CFOVER	CFSYS	SYSSUM
XCFIGROUP	XCFOVW	XCFCPATH	XCFSYS		

Available metrics for: ,SANDBOX,SYSPLEX

Metric description	Help	Id
% delay	Explanation	8D0160
% delay for enqueue	Explanation	8D1A20
% delay for i/o	Explanation	8D1A80
% delay for operator	Explanation	8D1AE0
% delay for processor	Explanation	8D1B40
% delay for storage	Explanation	8D1BA0
% delay for swsub	Explanation	8D1C00
% using	Explanation	8D04A0
% using for i/o	Explanation	8D1D40
% using for processor	Explanation	8D1DB0
% workflow	Explanation	8D0550
% workflow for i/o	Explanation	8D1ED0
% workflow for processor	Explanation	8D1F30

FAQ

RMF



© Copyright IBM Corp. 2007. All rights reserved.

Coupling Facility Overview Report



RMF Monitor III Data Portal for z/OS

RMF Report [,SANDBOX,SYSPLEX] : CFOVER (Coupling Facility Overview)

Time Range: 05/17/2007 11:02:00 - 05/17/2007 11:03:00

Policy Name: CFRM37 Policy Activation Date: 05/15/07 Policy Activation Time: 18.19.06 Policy Reformat Required: NO

CF Name	Model	Version	CF Level	CF Dynamic Dispatching	Status of CF	CF Storage Volatile	Proc Util %	Processors defined	Number of Dedicated Processors	Number of Shared Processors	Average Weighting of Shared Processors	Processors effective	Request Rate	Storage Size (in Bytes)
CF1	2094	S18	15	ON	Okay	YES	0.0	1	0	1	10	1.0	30.3	980418560
CF2	2094	S18	15	ON	Okay	YES	0.0	1	0	1	10	1.0	5.5	980418560
CF3	2094	S18	15	ON	Okay	YES	0.0	1	0	1	10	1.0	21.1	443547648



© Copyright IBM Corp. 2007. All rights reserved.

Coupling Facility Activity Report



'N/A' is shown in this field if the CF level is lower than 15

RMF Monitor III Data Portal for z/OS

RMF Report [,SANDBOX,SYSPLEX] : CFACT (Coupling Facility Activity)

Time Range: 05/16/2007 14:16:00 - 05/16/2007 14:17:00

Structure Name	Structure Type	Structure Status	Extended Structure Status	System Name	CF Utilization %	Structure Execution %	Sync Rate	Sync Avg Service Time	Sync Request Count	Async Rate	Async Avg Service Time	Asyn Requ Count
DB8FU_LOCK1	LOCK	A	ActivePersistent	*ALL	2.6	0.0	0.0	0	0	0.0	0	0
	LOCK			SC65			0.0	0	0	0.0	0	0
	LOCK			SC70			0.0	0	0	0.0	0	0
DB8FU_SCA	LIST	A	ActivePersistent	*ALL	6.8	0.0	0.0	0	0	0.0	0	0
	LIST			SC65			0.0	0	0	0.0	0	0
	LIST			SC70			0.0	0	0	0.0	0	0
EJESGDS_WTSCPLX4	LIST	A	ActivePersistent	*ALL	0.1	0.0	0.0	0	0	0.0	0	0
	LIST			SC65			0.0	0	0	0.0	0	0
	LIST			SC70			0.0	0	0	0.0	0	0
IGWLOCK00	LOCK	A	ActivePersistent	*ALL	46.4	0.0	0.0	0	0	0.0	0	0
	LOCK			SC65			0.0	0	0	0.0	0	0




© Copyright IBM Corp. 2007. All rights reserved.

ibm.com



e-business



CF Duplexing Performance CFLEVEL 15



Redbooks

International Technical Support Organization

© Copyright IBM Corp. 2007. All rights reserved.

RMF Support for CF Duplexing



- ❑ New CF values are provided and integrated in RMF:
 - RMF Postprocessor CF report is extended by:
 - Dynamic CF dispatching status
 - Number and average weight of shared processors
 - Support for granular CF processor utilization accounting
 - Monitor III CF reporting is extended by:
 - CF utilization percentage of an active structure
 - CF dynamic dispatching setting
 - Number of shared and dedicated processors
 - Subchannel busy value
 - SMF record type 74 subtype 4 is extended



© Copyright IBM Corp. 2007. All rights reserved.

RMF Postprocessor CF Activity Report



- The structure summary section of the Coupling Facility Activity Usage summary groups the structure summary data by structure type
 - The values shown for each structure are extended by ' % of CF utilization ' which shows the structure-related processor busy time for an allocated structure compared to the total CF processor busy time

$$\% \text{ OF CF UTIL} = \frac{\text{Structure execution time}}{\text{CF busy time}} * 100$$



© Copyright IBM Corp. 2007. All rights reserved.

CF Activity Report - Structure Summary



COUPLING FACILITY USAGE SUMMARY												
STRUCTURE SUMMARY												
TYPE	STRUCTURE NAME	STATUS CHG	ALLOC SIZE	% OF CF STOR	# REQ	% OF ALL REQ	% OF CF UTIL	AVG REQ/ SEC	LST/DIR ENTRIES TOT/CUR	DATA ELEMENTS TOT/CUR	LOCK ENTRIES TOT/CUR	DIR REC DIR XI'S
LIST	DB8FU_SCA	ACTIVE	9M	0.9	0	0.0	6.2	0.00	6464	13K	N/A	N/A
									112	425	N/A	N/A
	ISTMNPS	ACTIVE	13M	1.3	0	0.0	0.0	0.00	4163	8308	N/A	N/A
									1	0	N/A	N/A
	IXC_DEFAULT_1	ACTIVE	32M	3.4	19953	11.5	12.0	2.77	5767	5747	N/A	N/A
									1	16	N/A	N/A
	IXC_DEFAULT_3	ACTIVE	17M	1.8	2243	1.3	3.6	0.31	2061	2047	N/A	N/A
									1	32	N/A	N/A
	RRS_MAINUR_1	ACTIVE	13M	1.3	2156	1.2	1.6	0.30	2933	12K	N/A	N/A
									6	48	N/A	N/A
	SYSTEM_LOGREC	ACTIVE	17M	1.8	8	0.0	0.0	0.00	2027	16K	N/A	N/A
		PRIM							68	329	N/A	N/A
	SYSTEM_OPERLOG	ACTIVE	33M	3.5	817	0.5	0.8	0.11	29K	29K	N/A	N/A
		SEC							20K	22K	N/A	N/A
LOCK	DB8FU_LOCK1	ACTIVE	8M	0.9	0	0.0	2.4	0.00	11K	0	2097K	N/A
									0	0	N/A	N/A
	IGWLOCK00	ACTIVE	14M	1.5	0	0.0	39.3	0.00	33K	0	2097K	N/A
									0	0	21	N/A
	ISGLOCK	ACTIVE	9M	0.9	113494	65.5	9.3	15.76	0	0	1049K	N/A
									0	0	2805	N/A
CACHE	SYSIGGCAS_ECS	ACTIVE	5M	0.5	0	0.0	0.2	0.00	980	970	N/A	0
									0	0	N/A	0
	SYSZWLM_WORKUNIT	ACTIVE	12M	1.3	0	0.0	0.2	0.00	1052	2093	N/A	0
									0	0	N/A	0
STRUCTURE TOTALS			217M	23.2	173347	100	83.0	24.08				



© Copyright IBM Corp. 2007. All rights reserved.

D CF Command with CFLEVEL 15



D CF

IXL150I 11.43.35 DISPLAY CF 762
COUPLING FACILITY 002094.IBM.02.00000002991E
PARTITION: 0F CPCID: 00
CONTROL UNIT ID: FFF5

NAMED CF1

COUPLING FACILITY SPACE UTILIZATION

ALLOCATED SPACE		DUMP SPACE UTILIZATION	
STRUCTURES:	184832 K	STRUCTURE DUMP TABLES:	0 K
DUMP SPACE:	2048 K	TABLE COUNT:	0
FREE SPACE:	770560 K	FREE DUMP SPACE:	2048 K
TOTAL SPACE:	957440 K	TOTAL DUMP SPACE:	2048 K
		MAX REQUESTED DUMP SPACE:	0 K
VOLATILE:	YES	STORAGE INCREMENT SIZE:	512 K

CFLEVEL: 15
CFCC RELEASE 15.00, SERVICE LEVEL 00.18
BUILT ON 03/26/2007 AT 12:25:00

COUPLING FACILITY HAS 1 SHARED AND 0 DEDICATED PROCESSORS
DYNAMIC CF DISPATCHING: ON



© Copyright IBM Corp. 2007. All rights reserved.

ibm.com



WLM Group Capacity Limits



Redbooks

International Technical Support Organization

© Copyright IBM Corp. 2007. All rights reserved.

Group Capacity



- ❑ Group capacity allows soft capping across multiple LPARs defined in the same LPAR group
- ❑ Group capacity is available to LPARs in a System z9 environment that:
 - Use Workload Manager on z/OS V1R8 and above
 - Do not enforce relative weight (no PR/SM hard cap set)
 - Use the Hardware Management Console to define the capacity group
 - Use z/OS Workload Manager to monitor the rolling four-hour average utilization of all of the LPARs within the same group



© Copyright IBM Corp. 2007. All rights reserved.

Group Capacity Limit Support



- ❑ For z/OS V1R8 and z/OS V1R9, the WLM defined capacity mechanism is extended to handle LPAR groups instead of a single LPAR
 - This is called group capacity limit support and provides more flexibility when defining capacity limits for LPARs
- ❑ The group capacity limit balances the capacity between groups of partitions on the same processor
 - Requires IBM System z9 (EC or BC)
 - Software allows grouping of LPARs in the same processor
 - LPARs are then managed on a group basis using the existing WLM defined capacity mechanism



© Copyright IBM Corp. 2007. All rights reserved.

Group Capacity



- ❑ Group capacity limit is an extension of the defined capacity
 - It allows an installation to define a "soft cap" for multiple logical partitions of the same processor
- ❑ The group limit is a defined capacity (soft cap) for all partitions defined in the group
- ❑ The capacity group is defined on the Hardware Management Console
- ❑ Each capacity group has a name and a defined capacity which becomes effective to all partitions in the group



© Copyright IBM Corp. 2007. All rights reserved.

Group Capacity Limit Support



- ❑ Each partition manages itself independently from all other partitions
- ❑ Since group capacity is based on defined capacity
 - Therefore a 4 hour rolling average of the group MSU consumption is used as base for managing the partitions of the group
 - As usual, only general purpose CP service is considered
- ❑ Even when a partition receives capacity from another partition, it should not violate its defined capacity limit (if one exists)



© Copyright IBM Corp. 2007. All rights reserved.

Group Capacity Limit Support

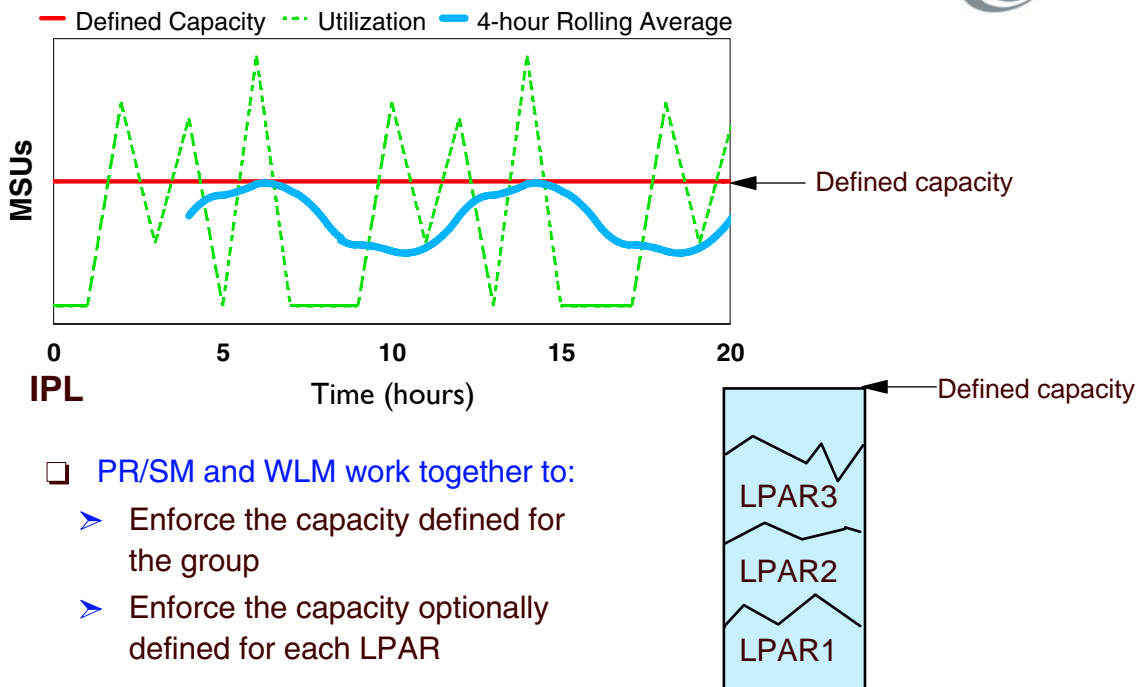


- ❑ Each partition sees the consumption of all other partitions on the processor
 - If the partition belongs to a group it identifies the other partitions of the same group
 - Calculates its defined share of the capacity group
 - Based on the partition weight
 - This share is the target for the partition if all partitions of the group want to use as much CPU resources as possible
 - If one or more LPARs do not use their share, this donated capacity is distributed over the LPARs which need additional capacity



© Copyright IBM Corp. 2007. All rights reserved.

LPAR Group Capacity



- ❑ PR/SM and WLM work together to:
 - Enforce the capacity defined for the group
 - Enforce the capacity optionally defined for each LPAR

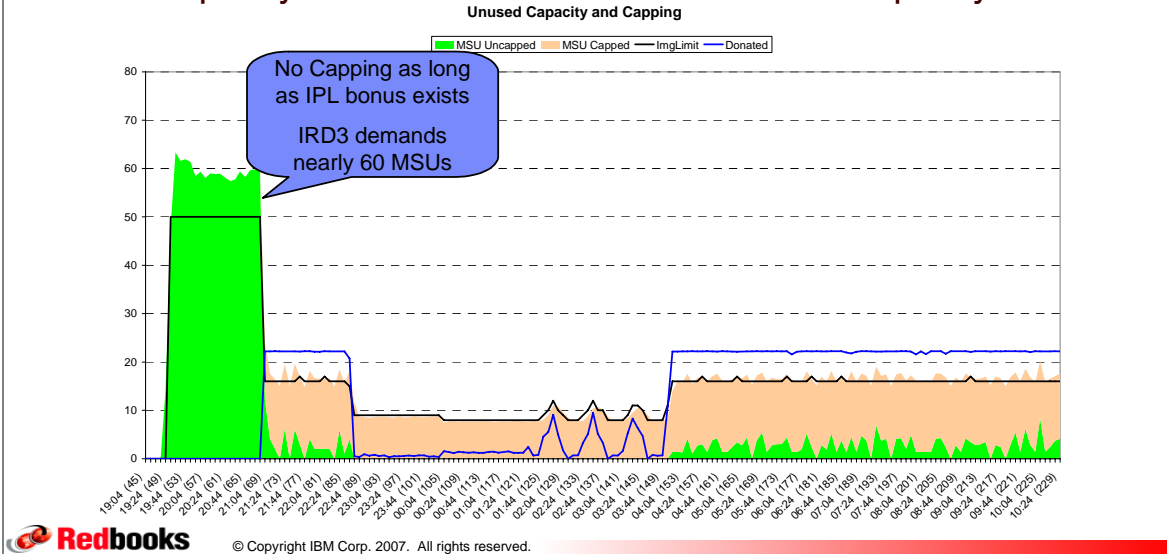


© Copyright IBM Corp. 2007. All rights reserved.

IPL Bonus



- A 4 hour bonus exists after IPL for the group
 - Meaning the group wide capping will start when the 4 hour rolling group average will reach the group capacity limit - similar to standard defined capacity



Group Capacity Limit Support



Partition	Group Name	Group Limit [MSU]	Weight	Target MSU consumption based on Weight	Defined Capacity (Softcap) [MSU]
A	Group1	200	70	93	n/a
B	Group1		50	67	80
C	Group1		30	40	30
D	n/a	n/a	100	n/a	120
E	n/a	n/a	50	n/a	n/a

Partition A: 93

- Can use up to 200 MSU based on 4 hour rolling average. If all three partitions want to use as much as possible, partition A will get $200 * 70 / (70 + 50 + 30) = 93$ MSU

Partition B: 67

- Can use up to 80 MSU because an individual softcap is defined. If all three partitions want to use as much as possible, partition B will get 67 MSU

Partition C: 40

- Can use up to 30 MSU because the defined capacity is smaller than the target based on partition weight

Select Customize/Delete Activation Profiles

Views: Groups, Exceptions, Active Tasks, Console Actions, Task List, Books, Help

Operational Customization:

- Hardware Messages
- Change LPAR Controls
- Operating System Messages
- Configure Channel Path On/Off
- Customize/Delete Activation Profiles
- Reassign Channel Path
- Customize Activity Profiles
- OSA Advanced Facilities
- Automatic Activation
- Enable I/O Priority Queuing
- Customize Scheduled Operations
- Change LPAR I/O Priority Queuing
- Customize Support Element Date/Time
- Change LPAR Group Controls

CPC Images Work Area:

- SCZP101 A01 (PLEX75:SC74)
- SCZP101 A02 (VMLINUX6)
- SCZP101 A03
- SCZP101 A04 (SANDBOX:SC63)
- SCZP101 A05 (OPPLEX:SC59)
- SCZP101 A06 (SANDBOX:SC64)
- SCZP101 A07 (WTSCVMT)
- SCZP101 A08 (WTSCPLX1:SC42)
- SCZP101 A09 (WTSCPLX1:SC04)
- SCZP101 A0A (WTSCPLX1:SC66)
- SCZP101 A0B
- SCZP101 A0C
- SCZP101 A0D
- SCZP101 A0E
- SCZP101 A0F

Redbooks © Copyright IBM Corp. 2007. All rights reserved.

Select Image Profile

Customize/Delete Activation Profiles List : SCZP101:A01

Select	Profile Name	Type	Profile Description
<input checked="" type="radio"/>	A01	Image	A01 Image profile.
<input type="radio"/>	DEFAULTLOAD	Load	This is the default Load profile.
<input type="radio"/>	IPL8403	Load	Z17RA1
<input type="radio"/>	IPL8503	Load	Z17RB1
<input type="radio"/>	IPL8603	Load	Z17RC1
<input type="radio"/>	IPL8703	Load	Z17RD1
<input type="radio"/>	IPLA803	Load	Z16RA1
<input type="radio"/>	IPLAB03	Load	Z16RB1
<input type="radio"/>	IPLAF04	Load	Z05RE1
<input type="radio"/>	IPLD102	Load	Z18RC1
<input type="radio"/>	IPLD202	Load	Z18RD1
<input type="radio"/>	PST2	Load	Parallel Sysplex Trainer 2
<input type="radio"/>	PST3	Load	Parallel Sysplex Trainer 3
<input type="radio"/>	SBOXD002	Load	Z18RB1 GA Sandbox sysplex
<input type="radio"/>	SBOXD102	Load	Z18RC1 GA Sandbox sysplex
<input type="radio"/>	SBOXD14E	Load	Z19RA1 CB01 Sandbox sysplex
<input type="radio"/>	SBOXD202	Load	Z18RC1 GA Sandbox sysplex
<input type="radio"/>	SBOXD21C	Load	Z19RB1 GA Sandbox sysplex
<input type="radio"/>	SBOXD31C	Load	Z19RC1 CB02 Sandbox sysplex
<input type="radio"/>	SBOXD41C	Load	Z19RD1 CB01 Sandbox sysplex

Buttons: Customize, Delete, Cancel, Help

Redbooks © Copyright IBM Corp. 2007. All rights reserved.

Assign Group Names



https://sczhmc6.itso.ibm.com - SCZHM6: Customize Image Profiles: SCZP101 A01 (PLEX75:SC74) : A01 : Processor - Mozilla Firefox

Customize Image Profiles: SCZP101 A01 (PLEX75:SC74) : A01 : Processor

SCZP101 A01 (PLEX75:SC74) Group Name [GROUP1]

Dedicated processors
 Select Processor Type

	Initial	Reserved
<input checked="" type="checkbox"/> Central processors (CPs)	4	2
<input type="checkbox"/> Integrated facilities for applications (IFAs)	0	0
<input checked="" type="checkbox"/> System z9 integrated information processors (zIIPs)	0	2

Not Dedicated Processor Details for :
 CPs zIIPs

CPs
 CP Details
 Initial processing weight: 10 (1 to 999) Initial capping
 Enable workload manager
 Minimum processing weight: 0
 Maximum processing weight: 0

Save Copy Notebook Paste Profile Assign Profile Cancel Help

Redbooks © Copyright IBM Corp. 2007. All rights reserved.

Select Change LPAR Group Controls



https://sczhmc6.itso.ibm.com - SCZHM6: Hardware Management Console Workplace (Version 2.9.2) - Mozilla Firefox

Views

Groups Exceptions Active Tasks Console Actions Task List Books Help

Defined CPCs Work Area

SCZP101 SCZP801 SCZP901 SCZPFK01 TC8W

Operational Customization

- Hardware Messages
- Operating System Messages
- Customize/Delete Activation Profiles
- Customize Activity Profiles
- Automatic Activation
- Customize Scheduled Operations
- Customize Support Element Date/Time
- Change LPAR Controls
- Configure Channel Path On/Off
- Reassign Channel Path
- OSA Advanced Facilities
- Enable I/O Priority Queuing
- Change LPAR I/O Priority Queuing
- Change LPAR Group Controls

Redbooks © Copyright IBM Corp. 2007. All rights reserved.

Change LPAR Group Controls



Group Name	Member Partitions	Group Capacity Value
DEFAULT	0	
GROUP1	200	



© Copyright IBM Corp. 2007. All rights reserved.

Assign New Group Name



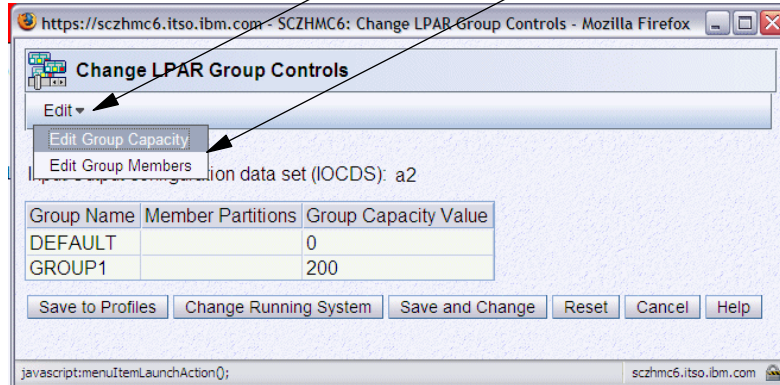
To assign a new group

	Initial	Reserved
Cent	4	2
Integ	0	0
System z9 integrated information processors (zIIPs)	0	2



© Copyright IBM Corp. 2007. All rights reserved.

To Change Members or Group Capacity



© Copyright IBM Corp. 2007. All rights reserved.

RMF Group Capacity Limit Support



- ❑ **RMF CPU activity reports information on capacity groups**
 - This allows to assess the capacity group definitions and their effect on the system
- ❑ **New Group Capacity section available in the Postprocessor CPU Activity report**
 - Informs about the share each LPAR in the group can take, the guaranteed minimum MSU share of the LPARs and the overall MSU consumption within the group
- ❑ **Monitor III CPC Capacity report, Partition Data Report section of the CPU Activity report include new header information showing group's name and capacity limit**



© Copyright IBM Corp. 2007. All rights reserved.

SMF and RMF



- ❑ **SMF Type 70 Record**
 - Contains group name, group membership (time stamp) and group MSU limit
- ❑ **Partition Data Report**
 - Shows group name and limit for the current partition plus information whether the group membership has changed during the last interval
 - Also available in the RMF Monitor III CPC Report
- ❑ **New Group Capacity Report**
 - Shows how the MSUs are used in the group
- ❑ **New Overview Conditions trend reporting also available**



© Copyright IBM Corp. 2007. All rights reserved.

Group Capacity Considerations



- ❑ **z/OS zNALC and z/OS traditional are treated the same way as any other product for group capping**
- ❑ **A partition can be dynamically removed from a group and/or added to another group**
 - The changed partition has no knowledge about unused capacity of the new group and it does not keep a history from previous activity
 - All systems must learn again about the new situation - therefore the group limit cannot be guaranteed
- ❑ **Works together with IRD Weight Management and Vary CPU Management**
 - IRD Weight Management may change the weight of partitions in a capacity group and changes the target share of the partition in the capacity group



© Copyright IBM Corp. 2007. All rights reserved.

Group Capacity Considerations



- ❑ When new members join a group...
 - If a new system is IPLed and joins a group it does not have the history of the MSU consumption of the complete group
 - Therefore it can take up to 4 hours until all systems in the group have the same view
 - During that time period the group limit cannot be guaranteed
 - If that happens the other partitions will be reduced to their capacity based on their weight in the group