



International Technical Support Organization

## IBM System z9 - LPAR Management

[ibm.com/redbooks](http://ibm.com/redbooks)



# Redbooks Workshop

IBM ITSO - International Technical Support Organization

IBM System z9 Workshop

© 2005 IBM Corporation

[ibm.com/redbooks](http://ibm.com/redbooks)

International Technical Support Organization



## Agenda



**Redbooks Workshop**  
IBM ITSO - International Technical Support Organization

IBM System z9 Workshop

© 2005 IBM Corporation

2

## LPAR Management



## z9-109 PR/SM Support

- PR/SM provides support for :
  - 64 logical processors per partition
    - Number of processors supported passed as input to LPAR during IML.
  - 512 GB partition storage size
  - MIDAWs (modified indirect data addressing word)
    - New facility reported in 'Store Channel Subsystem Characteristics' CHSC
    - Allows partition 'Start Sub-channel' to specify MIDAW being used in 'operation request block'
  - Coupling facility "isolate" partition
  - Partition link incident enablement facility
  - Link recovery threshold facility

## Z9-109 PR/SM Support....

- Partition Multiprocessing improvements
  - Reduce MP z/OS spin lock resolution effects
- Dynamic HSA management support
- Program-directed re-IPL and dump/re-IPL
- Separate management of CP, ICF, IFL and zAAP processors
- Concurrent driver upgrade
  - PR/SM driver build and patch management
  - Flexible PR/SM-SE interface
- Server Time Protocol (STP)
- Enhanced Book Availability (EBA)
- Partition multiple sub-channel set enablement
  - PR/SM usage of multiple sub-channel sets
  - Minor update to 'Store Channel Subsystem Characteristics' to make the facility visible to a partition

## Couple facility "isolate" fencing

- z/OS partition initiates request to 'isolate' (reset) a target partition to coupling facility.
- Target partition receives notification to perform 'isolation' operation and communicates to coupling facility when complete
- PR/SM posts new sub-channel CRW indicating 'partition isolation' of target partition to PR/SM on completion of target partition isolation
- PR/SM resets target partition
  - The PSW values for the logical processors in the configuration are not changed
  - All logical processors in the configuration enter the stopped state

## Partition link incident enablement facility

- Allows routing a link-incident-record event CRW to a particular partition, i.e. z/OS
- New facility in 'Store Channel Subsystem Characteristics' CHSC
- Enabled by new operation code for CHSC 'Set Domain Attributes'.
- If no partition enabled via 'Set Domain Attributes' current processing performed (CRW goes to first active non-idle partition found)
- If multiple partitions have been enabled via 'Set Domain Attributes', PR/SM gives to first active non-idle partition found

## Link Threshold recovery facility

- New facility in 'Store Channel Subsystem Characteristics' CHSC
- Enabled/reset by new operation codes for 'Set Channel Subsystem Characteristics' CHSC (no PR/SM involvement)
- PR/SM code to allow new CHSC 'Store Channel Subsystem Information' to be issued directly by partition

## Dynamic HSA Management

- New PCCALL ('D9'X) to query current state of HSA storage allocation including absolute increment number and book location
- PR/SM will issue this every time storage allocation request is made or storage may be relocated. PR/SM internal control blocks updated as needed. Includes:
  - Partition activation
  - Partition Attach Storage Element SCLP
  - Partition Assign Storage Increment SCLP
  - Partition Configure Expanded Storage Element SCLP
  - Partition storage reassignment during optimization
  - Concurrent Book Replacement storage evacuation

## Program Directed IPL

- New subcodes of Diag'308'x (IPL Functions) for IPL and dump/re-IPL. Both traditional CCW and SCSI IPL formats are supported.
  - subcode 3 -- IPL Clear
  - subcode 4 -- IPL Normal/Dump
  - subcode 5 -- Set IPL Parameters
  - subcode 6 -- Store IPL parameter
- The IPL parameters used when a program-directed re-IPL is issued are those which were provided by Diagnose '308'x 'set IPL parameters' or the last IPL, whichever is most recent.
- CCW-type re-IPL LPAR actions are identical to those that take place through the existing Support Element / LPAR interface.

## Separate management of GP/ICF/IFL/zAAPs

- Four separate and distinct physical processor pools (GPs, ICFs, IFLs and zAAPs) supported
- Main internal change is that logical processor weights for each CPU type are individually managed

## Separate management of GP/ICF/IFL/zAAPs

### PR/SM PU Dispatching “Pools”

- **PU Pool – Physical PUs to dispatch to online logical PUs**
- **Z9-109 with 10 CPs, 1 ICF, 2 IFLs, and 3 zAAPs**
  - CP pool contains 10 CP engines
  - ICF pool contains 1 ICF
  - IFL pool contains 2 IFLs
  - zAAP pool contains 3 zAAPs
- **z990 with 10 CPs, 1 ICF, 2 IFLs, and 3 zAAPs**
  - CP pool contains 10 CP engines
  - Specialty pool contains 6 engines – ICFs, IFLs, zAAPs

## PR/SM PU Pool Rules

- **Logical PUs dispatched from supporting pool only**
  - Logical CPs from CP pool only, for example
- **Pool “width”**
  - Width equals the number of PUs in the pool
  - Limits LPAR maximum number of shared logical PUs brought ON
- **PUs placed in pools by**
  - Activate (POR)
  - Concurrent Upgrade – On/Off CoD, CBU, CIU, CUoD MES
  - Dedicated LPAR deactivation
  - Dedicated LPAR configure logical PU OFF
- **PUs removed from pools by**
  - Concurrent Downgrade - On/Off CoD, CBU, PU Conversion MES
  - Dedicated LPAR activation (“width” permitting)
  - Dedicated LPAR configure logical PU ON (“width” permitting)

## Z9-109 PR/SM Weight/Share Calculations

- **LPAR Share = Pool PUs x (LPAR Pool Weight)/(Total Pool Weight)**
  - Can't exceed number of Online Logical Processors dispatched from the pool
- **Pool PUs (Physical) – CP =10, zAAP = 3, IFL = 2, ICF =1**
- **Total Pool Weights – CP = 1000, zAAP = 1000, IFL= 400, ICF= 100**

| LPAR Name              | LPAR Weight | Shared Logical PUs On |      |     |     | PU Share |      |     |     |
|------------------------|-------------|-----------------------|------|-----|-----|----------|------|-----|-----|
|                        |             | CP                    | zAAP | IFL | ICF | CP       | zAAP | IFL | ICF |
| MVS1                   | 250         | 10                    | 2    | NA  | NA  | 2.5      | .75  | NA  | NA  |
| MVS2                   | 750         | 10                    | 3    | NA  | NA  | 7.5      | 2.25 | NA  | NA  |
| CF1                    | 50          | 0                     | NA   | NA  | 1   | 0        | NA   | NA  | .5  |
| CF2                    | 50          | 0                     | NA   | NA  | 1   | 0        | NA   | NA  | .5  |
| VM1                    | 100         | 0                     | NA   | 2   | NA  | 0        | NA   | .5  | NA  |
| LINUX1                 | 300         | 0                     | NA   | 2   | NA  | 0        | NA   | 1.5 | NA  |
| Pool Weight >          |             | 1000                  | 1000 | 400 | 100 |          |      |     |     |
| Total PUs (Physical) > |             | 10                    | 3    | 2   | 1   |          |      |     |     |



CPs



zAAPs



IFLs



ICFs

- **Probably the intended result**

## Z900 PR/SM Weight/Share Calculations

- LPAR Share = Pool PUs x (LPAR Pool Weight)/(Total Pool Weight)
  - Can't exceed number of Online Logical Processors dispatched from the pool
- Pool PUs (Physical) – CP =10, Specialty = 6 (ICF, IFL, zAAP)
- Total Pool Weights – CP = 1000, **Specialty = 1500**

| LPAR Name              | LPAR Weight | Shared Logical PUs ON |      |     |     | PU Share |           |
|------------------------|-------------|-----------------------|------|-----|-----|----------|-----------|
|                        |             | CP                    | zAAP | IFL | ICF | CP       | Specialty |
| MVS1                   | 250         | 10                    | 2    | NA  | NA  | 2.5      | 1         |
| MVS2                   | 750         | 10                    | 3    | NA  | NA  | 7.5      | 3         |
| CF1                    | 50          | 0                     | NA   | NA  | 1   | 0        | .2        |
| CF2                    | 50          | 0                     | NA   | NA  | 1   | 0        | .2        |
| VM1                    | 100         | 0                     | NA   | 2   | NA  | 0        | .4        |
| LINUX1                 | 300         | 0                     | NA   | 2   | NA  | 0        | 1.2       |
| Pool Weight >          | 1000        | 1500                  |      |     |     |          |           |
| Total PUs (Physical) > |             |                       |      |     |     | 10       | 6         |



CPs



Specialty

- Probably **NOT** the intended result!

## Z990... RMF Mon III CPC report

```

RMF V1R7  CPC Capacity
Samples: 120  System: Z03A  Date: 02/08/05  Time: 14.51.00  Range: 120
Partition:  Z03          2084 Model 301
CPC Capacity: 352  Weight % of Max: 100      4h MSU Average: 2
Image Capacity: 50  WLM Capping %: 0.0      4h MSU Maximum: 8

Partition  --- MSU ---  Cap Proc  Logical Util %  - Physical Util % -
           Def  Act  Def  Num   Effect  Total  LPAR  Effect  Total
*CP
Z01         0    0  NO   1.0    0.0    0.1    0.0    0.0    0.0
Z02         50   3  NO   3.0    1.4    1.5    0.1    0.7    0.8
Z03         50   2  NO   3.0    1.1    1.3    0.1    0.6    0.6
PHYSICAL
           0.4
*ICF
CF01         1.0    3.5    3.6    0.0    7.1    7.4
CF02         1.0    0.0    0.3    0.0    0.0    0.2
LIG29        1.0    0.0    0.0    0.0    0.0    0.0
LIG30        1.0    0.0    0.0    0.0    0.0    0.0
Z03         1.0    2.2    2.4    0.1    2.2    2.4
PHYSICAL
           0.1
    
```



## Z9-109... RMF Mon III CPC report

```

RMF V1R7 CPC Capacity

Samples: 120 System: Z03A Date: 02/08/05 Time: 14.51.00 Range: 120

Partition: Z03 2094 Model 706
CPC Capacity: 500 Weight % of Max: 100 4h MSU Average: 2
Image Capacity: 50 WLM Capping %: 0.0 4h MSU Maximum: 8

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

*CP
Z01      0 0 NO 1.0 0.0 0.1 0.0 0.0 0.0
Z02     50 3 NO 3.0 1.4 1.5 0.1 0.7 0.8
Z03     50 2 NO 3.0 1.1 1.3 0.1 0.6 0.6
PHYSICAL
          0.4 0.4

*IFA
Z03      1.0 2.2 2.4 0.2 2.2 2.4
PHYSICAL
          0.2 0.2

*ICF
CF01     1.0 3.5 3.6 0.3 7.1 7.4
CF02     1.0 0.0 0.3 0.1 0.1 0.2
PHYSICAL
          0.5 0.5

*IFL
LIG29     1.0 0.0 0.0 0.0 0.0 0.0
LIG30     1.0 0.0 0.0 0.0 0.0 0.0
PHYSICAL
          0.0 0.0
    
```

## Z990... Partition Data report

```

PARTITION DATA REPORT

z/OS V1R7 SYSTEM ID Z03A DATE 02/08/2005 INTERVAL 19.59.999
RPT VERSION V1R7 RMF TIME 12.20.00 CYCLE 1.000 SECONDS

MVS PARTITION NAME Z03 NUMBER OF PHYSICAL PROCESSORS 11
IMAGE CAPACITY 50 CP 6
NUMBER OF CONFIGURED PARTITIONS 26 ICF 5
WAIT COMPLETION NO
DISPATCH INTERVAL DYNAMIC

----- PARTITION DATA ----- LOGICAL PARTITION PROCESSOR DATA --- AVERAGE PROCESSOR UTILIZATION PERCENTAGES ---
NAME S WGT DEF ACT DEF WLM% NUM TYPE EFFECTIVE TOTAL EFFECTIVE TOTAL LPAR MGMT EFFECTIVE TOTAL
Z01 A 100 0 0 NO 0.0 1.0 CP 00.00.50.024 00.00.52.001 1.38 1.41 0.07 0.69 0.76
Z02 A 100 50 2 NO 0.0 3.0 CP 00.00.40.200 00.00.41.636 1.11 1.32 0.05 0.61 0.66
Z03 A 100 50 2 NO 0.0 3.0 CP 00.00.40.400 00.00.42.353 1.12 1.26 0.06 0.57 0.63
*PHYSICAL* 00.00.18.728 0.40 0.40
TOTAL 00.02.10.624 00.02.33.684 0.58 1.87 2.45

CF01 A 150 1.0 ICF 00.00.07.665 00.00.08.169 0.64 0.68 0.02 0.32 0.34
CF02 A 150 1.0 ICF 00.00.03.202 00.00.03.669 0.27 0.31 0.02 0.13 0.15
LIG29 A 40 1.0 ICF 00.00.00.000 00.00.00.000 0.00 0.00 0.00 0.00 0.00
LIG30 A 40 1.0 ICF 00.00.00.000 00.00.00.000 0.00 0.00 0.00 0.00 0.00
Z03 A 150 1.0 ICF 00.00.07.021 00.00.08.000 0.62 0.66 0.02 0.31 0.32
*PHYSICAL* 00.00.12.541 0.52 0.52
TOTAL 00.00.18.420 00.00.31.949 0.58 0.79 1.36
    
```

## Z9-109... Partition Data report

```

PARTITION DATA REPORT

z/OS V1R7          SYSTEM ID Z03A          DATE 02/08/2005          INTERVAL
19.59.999          RPT VERSION V1R7 RMF          TIME 12.20.00          CYCLE 1.000

SECONDS

MVS PARTITION NAME          203          NUMBER OF PHYSICAL PROCESSORS          11
IMAGE CAPACITY              50          CP          6
NUMBER OF CONFIGURED PARTITIONS          26          IFA          1
WAIT COMPLETION              NO          IFL          2
DISPATCH INTERVAL          DYNAMIC          ICF          2

----- PARTITION DATA ----- -- LOGICAL PARTITION PROCESSOR DATA -- -- AVERAGE PROCESSOR
UTILIZATION PERCENTAGES --
-- PHYSICAL PROCESSORS --
NAME S WGT DEF ACT DEF WLM% NUM TYPE EFFECTIVE TOTAL EFFECTIVE TOTAL
LPAR MGMT EFFECTIVE TOTAL

201 A 100 0 0 NO 0.0 1.0 CP 00.00.50.024 00.00.52.001 1.38 1.41
0.07 0.69 0.76
202 A 100 50 2 NO 0.0 3.0 CP 00.00.40.200 00.00.41.636 1.11 1.32
0.05 0.61 0.66
203 A 100 50 2 NO 0.0 3.0 CP 00.00.40.400 00.00.42.353 1.12 1.26
0.06 0.57 0.63
*PHYSICAL*
0.40 0.40
-----
TOTAL
0.58 1.87 2.45
-----
203 A 150 1.0 IFA 00.00.07.021 00.00.08.000 0.62 0.66
0.02 0.31 0.32
*PHYSICAL*
0.52 0.52
-----
TOTAL
0.54 0.31 0.84
    
```

## Z9-109... Partition Data report

```

LIG29 A 40 1.0 IFL 00.00.00.000 00.00.00.000 0.00
0.00 0.00 0.00 0.00
LIG30 A 40 1.0 IFL 00.00.00.000 00.00.00.000 0.00
0.00 0.00 0.00 0.00
*PHYSICAL*
0.00 0.00
-----
TOTAL
0.00 0.00 0.00
-----
CF01 A 150 1.0 ICF 00.00.07.665 00.00.08.169 0.64
0.68 0.02 0.32 0.34
CF02 A 150 1.0 ICF 00.00.03.202 00.00.03.669 0.27
0.31 0.02 0.13 0.15
*PHYSICAL*
0.52 0.52
-----
TOTAL
0.56 0.45 1.01
    
```