

Understanding and Controlling LPARs

Draft White Paper V3

Relative Weight

Every partition must have a relative weight, [the relative weight is also known as the initial processing weight](#). This relative weight is used to determine which LPAR should get resources. However, the relative weight is only considered there is contention between the LPARs for resources.

If only one LPAR needs resources and the other two LPARs are idle, then the weights will not be factored. However, if one of the idle LPARs becomes active and starts competing for resources, then PR/SM will divide the processing power between the two LPARs based upon the relative weights.

Relative weights are unit-less. That is, they have no unit of measure. Instead, the term 'relative' means that the weights are compared against each other. They need not sum to any particular value, either. Below, in Example 1, is a sample configuration for a mainframe with three configured LPARs:

Example 1:

LPAR	Weight	Sum of Weights	Relative %
LPAR1	614	614+91+295 = 1000	614/1000 = ~61%
LPAR2	91	614+91+295 = 1000	91/1000 = ~9%
LPAR3	295	614+91+295 = 1000	295/1000 = ~30%

**note: the sum of the relative weights in this example is 1000. However, the weights need not sum to any specific value.*

By changing any other the weights on these LPARs will affect the resources available to EACH LPAR... because the weights are all relative to each other. In Example 2, if you change LPAR1's weight from 614 to 200, here is what occurs:

Example:

LPAR	Weight	Sum of Weights	Relative %
LPAR1	614 200	614 200+91+295 = 586	200/586 = ~34%
LPAR2	91	614 200+91+295 = 586	91/586 = ~16%
LPAR3	295	614 200+91+295 = 586	295/586 = ~50%

Notice, in Example 2, that the relative weight of LPAR2 increased from the original 9% to 16% even though you didn't directly change LPAR2. Similarly the relative weight of LPAR3 increased from the original 30% to 50% even though you didn't directly change LPAR3.

Hardcapping

You have the right to enforce the relative weight, also known as HARDCAP or PR/SM hardcap. This means you **never** allow the LPAR to use more than its share of resources, as specified in the relative weight. Looking at the information in Example 2, if you chose to enforce the relative weight of LPAR2 at 91 (16% of the machine) then LPAR2 could never use more than 16%, even if the other LPARs were not using any resources. Outsourcers who charge customers for a specific amount of computing power often like to use hardcaps to ensure that those customers never use more than they've paid for. Occasionally hardcaps are used to limit the resources available in test environments.

Defined Capacity

Beginning with zSeries running z/OS in 64-bit mode in October 2000, IBM added another setting for customers to use when setting up their LPARs, called "Defined Capacity." Instead of using hardcaps, many Sub-Capacity customers choose to use the more flexible (and more appropriate) defined capacity. Defined capacity allows you to set a size of a partition, in MSUs. This setting is not relative to any other LPAR. Once you set a defined capacity, you are telling WLM to monitor the rolling 4-hour average utilization of that LPAR.

Defined Capacity is an LPAR setting available to LPARs that meet these criteria:

- zSeries hardware
- running z/OS in 64-bit mode
- shared general purpose engines (i.e., no dedicated engines)
- relative weight NOT enforced (i.e., no PR/SM hardcap)

PLEASE KEEP IN MIND: Defined Capacity is only available when there is no hardcap (that is, the relative weight is NOT enforced). *Please note that the HMC will allow you to set both a defined capacity and enforce the relative weight... but WLM will ignore the defined capacity.*

If the rolling 4-hour average utilization of the LPAR is less than the defined capacity, then nothing happens. If the rolling 4-hour average utilization of the LPAR exceeds the defined capacity, then WLM signals to PR/SM and requests that something called 'softcapping' be initiated. Softcapping constrains the workload of the LPAR to the level of the defined capacity.

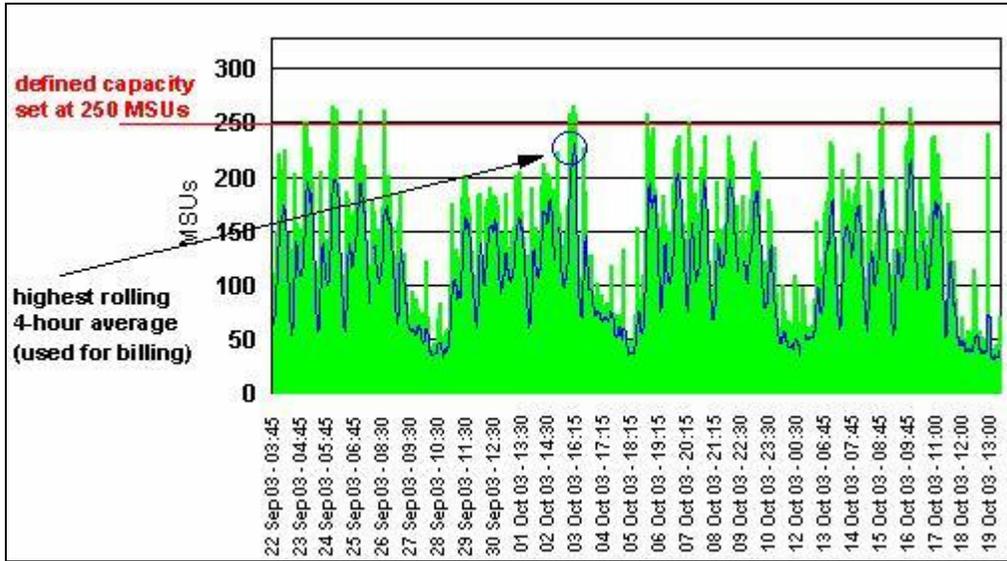
Softcapping remains in place until the rolling 4-hour average utilization of the LPAR goes back down below the defined capacity. During this period when softcapping is in place, the rolling 4-hour average may exceed the defined capacity but the Sub-Capacity Reporting Tool will report the defined capacity instead of the actual rolling 4-hour average utilization for any intervals when the defined capacity is in place... Remember, with Defined Capacity, LPAR utilization may exceed the Defined Capacity... it is only when the rolling 4-hour average utilization rises above the defined capacity that softcapping is initiated.

Establishing Defined Capacity

Defined Capacity is optional setting. It may be changed dynamically. Use the HMC to establish/change a Defined Capacity. You can also establish a Defined Capacity in the Image Profile, which will ensure the desired Defined Capacity is used if the system is re-IPLed. Defined Capacity setting is in terms in MSUs. In cases where there are different hardware MSUs and software pricing MSUs for a given zSeries server, the Defined Capacity should be established in terms of software pricing MSUs. If you establish a defined capacity for an LPAR which has a PR/SM hardcap or dedicated engines, the defined capacity will be ignored.

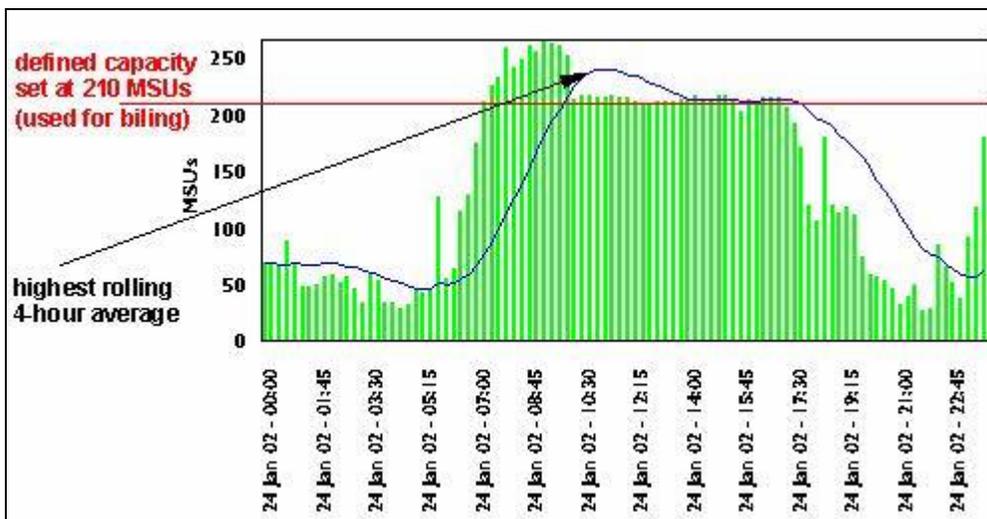
Defined Capacity and Sub-Capacity Software Charges

Defined Capacity is the most sophisticated means of controlling an LPAR's rolling 4-hour average in a Sub-Capacity software charging environment. In cases when a customer establishes a defined capacity but the rolling 4-hour average utilization of the LPAR does not reach the defined capacity, then the software charges will be based upon the highest observed rolling 4-hour average utilization (see CASE1 illustration, below).



CASE1, Rolling 4-Hour average utilization is below the defined capacity

However, in cases where the rolling 4-hour average does exceed the defined capacity and softcapping is applied, it is possible (while being softcapped) for the rolling 4-hour average to continue to rise and exceed the defined capacity. In these cases, IBM will charge customers at the defined capacity level rather than based upon the rolling 4-hour average utilization (see CASE2 illustration, below).



CASE2, Rolling 4-Hour average utilization exceeds the defined capacity

Softcapping Levels

In almost all cases, the softcapping level applied by PR/SM and WLM is equal to the defined capacity set by the customer. However, there is a special case where the softcapping level may be less than the defined capacity. Let's look at both cases using an example...

Let's say you have a 100 MSU machine with two LPARs, LPARA and LPARB. Let's say you give LPARA a relative weight of 5 and LPARB a relative weight of 45. If you convert these weights into MSUs, you get

$$\text{LPARA} = 5/(45+5) = 10\% = .10 * 100 \text{ MSUs} = 10 \text{ MSUs}$$

$$\text{LPARB} = 45/(45+5) = 90\% = .90 * 100 \text{ MSUs} = 90 \text{ MSUs}$$

Typical Case, Softcapping = Defined Capacity:

Relative Weight > Defined Capacity

If I set a defined capacity of LPARB at 75 MSUs, then my LPAR's relative weight (90 MSUs) is greater than the defined capacity. Therefore, softcapping will be at the defined capacity level of 75 MSUs.

Special Case, Softcapping < Defined Capacity for Periods of Time:

Relative Weight < Defined Capacity

If I set a defined capacity of LPARA at 20 MSUs, then my LPAR's relative weight (10 MSUs) is less than the defined capacity. Therefore, softcapping will alternate between enforcing the relative weight (10 MSUs) and no cap at all. This cap on / cap off pattern will apply to try to effectively achieve the 20 MSU defined capacity. It would be wise to avoid this case and try to establish a defined capacity which is less than or equal to the relative weight.

RMF Monitors

There are three RMF monitors available for monitoring the rolling 4-hour average:

1. Partition Data Report includes a field called "% Capping", which shows the % of time that softcapping was applied
2. CPC Capacity RMF Mon III Command, which shows rolling 4-hour average, % capping
3. RMF PM Java Client, which shows a number of software-pricing-related views

You can read more about RMF monitors for software pricing on the web, at:

<http://ibm.com/zseries/zos/rmf/rmfhtmls/rmfnews14.html>

Note: DO NOT READ THE PRICING INFORMATION ON THIS WEBSITE... IT IS INCORRECT...

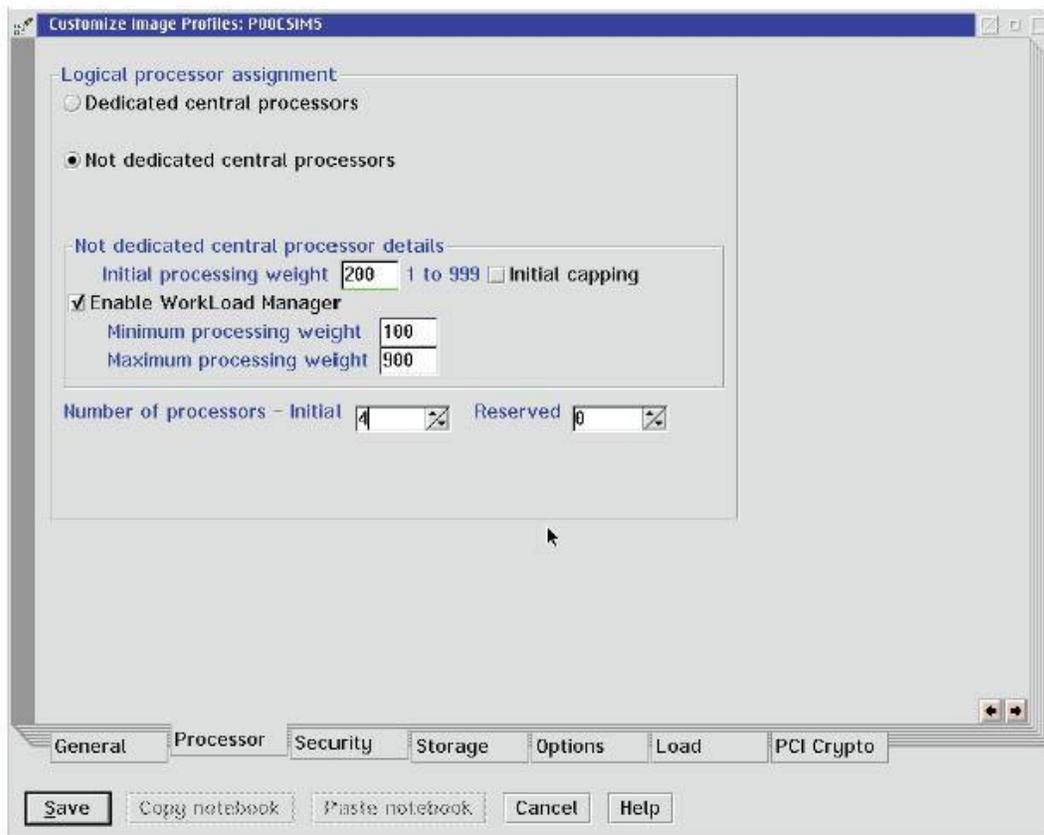
One last thing that people ask about... when configuring your LPARs on the HMC, you have an option to select "WLM Managed" for each LPAR. WLM Managed means that you want WLM to dynamically change the weights for each LPAR and dynamically increase/decrease assigned logical CPs for each LPAR. This concept of WLM management is generally known as Intelligent Resource Director (IRD). If you want IRD (WLM) dynamically managing your LPAR weights and assigned logical engines, check the 'WLM Managed' checkbox.

HMC Panels

You would use the Customize Image Profile panel on the HMC to customize a profile (establish the configuration) for one of your LPARs. Changing the profile will not affect an LPAR which is already running. However, you may choose to use this profile as the configuration when you next IPL your LPAR.

When you are updating the profile, there are some steps that you should take to ensure that Defined Capacity is successfully enabled. Here are the three steps:

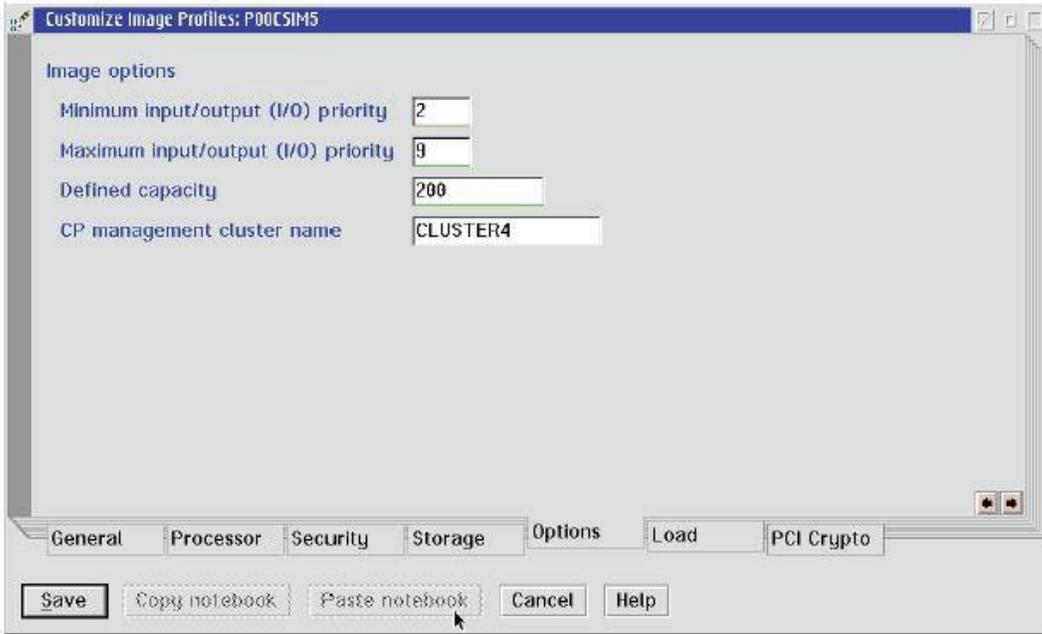
1. select 'not dedicated central processors' on the Processor tab – this enables shared processors
2. uncheck the box labeled 'Initial Capping' on the Processor tab – this disables PR/SM hardcapping
3. specify a defined capacity (in MSUs) on the Options tab – this enables Defined Capacity



Panel 1, the Processor Tab of the Customize Image Profile Panel

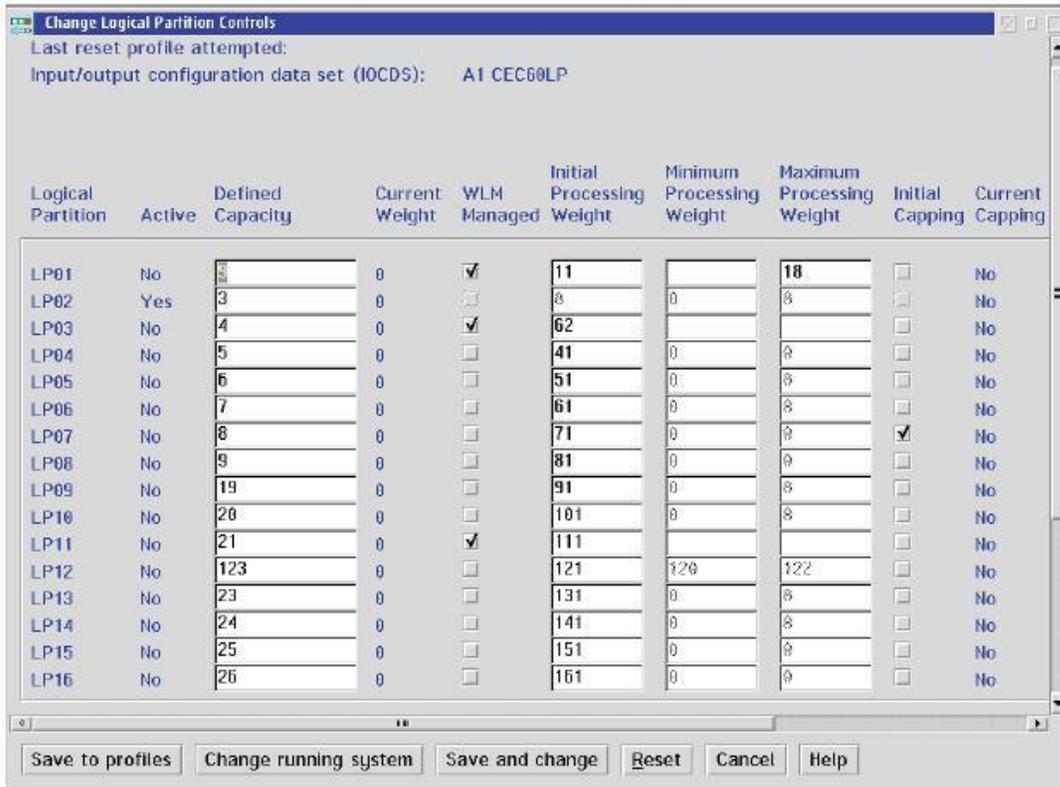
Continued on the next page...

HMC Panels, continued



Panel 2, the Options Tab of the Customize Image Profile Panel

Instead of updating the LPAR's profile, you might prefer to change the defined capacity on a running system. To do so, update the 'Defined Capacity' field on the Change Logical Partition Controls panel on the HMC. Again, be sure that you leave the 'initial capping' box unchecked.



Panel 3, the Change Logical Partition Controls Panel