
Table des matières

Processing Capacity

(Processing Units, Entitled Capacity)

Capacity Entitlement value is the time when you can access the core. Processing capacity can be configured in fractions of 0.01 processors. The minimum amount of processing capacity that has to be assigned to a micro-partition is 0.1 processors. (Starting with POWER7+ minimum is 0.05)

On the HMC processing capacity is specified in terms of processing units. (To assign a processing capacity representing 75% of a processor, 0.75 processing units are specified on the HMC.)

Once a micro partition is activated, processing capacity is usually referred to as entitled capacity. This entitled capacity is reserved for that partition when it is started. The total entitled capacity of all the partitions configured on the system can never exceed the number of processors in that system.

Capped and uncapped mode

Micro-partitions have a specific processing mode that determines the maximum processing capacity given to them from their Shared-Processor Pool:

Uncapped mode: processing capacity can exceed the entitled capacity when resources are available in their Shared-Processor Pool (if there is competition amongst several partitions for unused processor capacity the uncapped weight will determine the sharing)

Capped mode: processing capacity can not exceed the entitled capacity

Uncapped Weight: Extra capacity is distributed between partitions in uncapped mode, by checking its uncapped weights. The higher the uncapped weight of a micro-partition, the more processing capacity the micro-partition will receive. (You must specify the uncapped weight when it is created.) The uncapped weight must be a whole number from 0 to 255. Default is 128. (If you set the uncapped weight at 0, the micro-partition will be treated as a capped partition.)

Virtual Processor

A virtual processor is a representation of a physical processor that is presented to the operating system. A virtual processor cannot have a greater processing capacity than a physical processor. The capacity of a virtual processor will be equal to or less than the processing capacity of a physical processor. In terms of capacity, a physical processor can support up to 10 virtual processors. (From Firmware 760, 20 virtual processors.)

The Hypervisor uses a 10 ms dispatch cycle; each virtual processor is guaranteed to receive its entitled capacity of processor cycles during each 10 ms dispatch window. The Entitlement Capacity will be distributed equally between the Virtual Processors by the Hypervisor.

For example, if a logical partition has 1.60 processing units and 2 virtual processors, each virtual processor has 0.80 processing units. But the same LPAR with 4 virtual processors, each virtual processor has 0.40 processing units.

For example, if you have a partition with 0.50 processing units and one virtual processor, the partition cannot exceed 1.00 processing units. However, if the same partition with 0.50 processing units was assigned two virtual processors and processing resources were available, the partition could use an additional 1.50 processing units. The number of Virtual Processors can not be extended over 10X of the Entitled Capacity. (From Firmware 760 it is 20X) (If EC=0.5 then the maximum number of Virtual Processor can not be over 5, or from Firmware 760 over 10.)

Desired, minimum and maximum values

When you set up a partition profile, you set up the desired, minimum, and maximum values you want for the profile. When a partition is started, the system chooses the partition's entitled processor capacity from this specified capacity range.

When starting a micro-partition, preference is given to the desired value, but this value cannot always be used because there may not be enough unassigned capacity in the system. In that case, a different value is chosen, which must be greater than or equal to the minimum capacity attribute. If the minimum capacity requirement cannot be met, the micro-partition will not start. The maximum value is only used as an upper limit for dynamic operations. (adding RAM/CPU as a DLPAR operation).

Processor Affinity (and processor binding):

Hypervisor always attempts to dispatch the virtual processor on to the same physical processor that it previously ran on. (To reduce context switching...) This behavior is called processor affinity. If this is not possible, the POWER Hypervisor will broaden its search out to other processors.

With AIX 6.1 TLS and AIX 7.1 binding processors is available to application running in a shared LPAR (an application process can be bound to a virtual processor). PowerVM hypervisor (from firmware level 730) maintains three levels of affinity for dispatching, such as core, chip and node level affinity. By maintaining affinity at hypervisor level as well as in AIX, applications may achieve higher level affinity through processor bindings.

Processor Folding:

Virtual processor folding puts idle virtual processors into a hibernation state so that they do not consume any resources. There are several important benefits of this feature: improved processor affinity, reduced POWER Hypervisor workload... It will reduce the Physical Core usage (physc) at LPAR level.

From:
<https://unix.ndlp.info/> - Where there is a shell, there is a way

Permanent link:
https://unix.ndlp.info/doku.php/informatique:nix:ibm:ibm_aix_powervm:definitions

Last update: 2015/11/27 09:29