

Dedicated Central Processors

- An LPAR can have (non-shared) dedicated CPs.
- When an LPAR that uses dedicated CPs is activated, a physical CP is assigned to each defined logical CP.
 - The LPAR then has exclusive use of its physical CPs. The physical CPs that belong to an LPAR that uses dedicated CPs are always available for its use, but the capacity that is not used cannot be used by other LPARs.

- **Workloads best suited** for logical partitions that use dedicated processors are those that maintain a fairly even external throughput rate (ETR) while using most of the capacity of the logical partition.

- LPARs with timing dependencies might require dedicated processors.
- NOTE: ETR** = is the number of transactions completed per wall clock second.
- The installation goals for a logical partition that uses dedicated processors should be similar to the goals of the processor complex.

- For example, if the goal is an average processor utilization rate of 70%–85%, then the same goal should be used for a logical partition that uses dedicated processors.*

Shared Central Processors

- LPARs can share CPs with other LPARs.
- A sharing LPAR does not have exclusive use of the physical CPs.
 - There is no correspondence between the logical CPs in a sharing LPAR and the physical CPs on which the logical CPs are dispatched (except on a one-way CPC).
 - A logical CP can be dispatched on any physical CP and, on successive dispatches, the logical CP can be dispatched on different physical CPs.
 - The number of CPs available for use by sharing LPARs is determined by adding the number of CPs already assigned to active, dedicated LPARs and subtracting that sum from the total number of physical CPs available.
 - The total of all logical CPs for all sharing LPARs can be larger than the number of physical CPs serving the sharing LPARs.

- For example, if four LPARs are active on a six-way CPC and each LPAR is defined to have four CPs, the total number of online logical CPs is 16.*

- **Workloads best suited** for sharing logical partitions are those that have a widely fluctuating ETR or would not fit well into the capacity of a dedicated logical partition.

- If a workload can use only a small portion of the capacity of a dedicated logical partition, customers should redefine the logical partition to use shared processors to free the available capacity for use by other logical partitions.
- A workload with a widely fluctuating ETR will experience peaks and valleys in its processor-utilization curve and have fluctuations over extremely short periods of time (minutes or seconds).
- This type of workload could take advantage of the **time-** and **event-** driven dispatching available
 - > With event-driven dispatching, a sharing logical partition receives the resources required as needed and leaves the capacity free for other logical partitions when not needed.
 - > Time driven LPAR dispatching uses a time-slice interval (i.e. 0.25 ms)

- When combining workloads on a processor complex by means of logical partitions, customers should examine their average and peak requirements carefully. If the workloads fluctuate over very short intervals, the total capacity of the system must meet the sum of the average requirements for each workload. If processor utilization fluctuates over longer periods, and the peak utilization periods for these workloads occurs simultaneously, then the total capacity of the logical partitions must meet the sum of the peak requirements for each workload.

- Sharing logical partitions that use event-driven dispatching are better able to maintain high transaction rates with fluctuating demand while being responsive.
 - However, the ITR for a sharing logical partition is lower than the ITR for a dedicated logical partition.

NOTE: ITR = The amount of transactions completed per SRM (CPU) second.

- The capability to limit CPU usage for any or all logical partitions with shared processors is provided by the PR/SM capping function (see later in this issue).
 - The capping function enhances LPAR workload balancing controls for environments with a requirement that the CPU resources for a logical partition be limited.
 - Capped logical partitions are recommended for use when CPU resources must be limited for business reasons (in accordance with a contract), or when the impact that one logical partition can have on other logical partitions needs to be limited.

Processing Weights

- Logical partition processor weights are a **priority policy** for logical partitions based on a user's dispatching priorities.
- They are specified when allocating the resources of the machine to its workloads. Using processor weights with shared processors, the customer can use subprocessor granularity for allocating machine resources to the workload.

- Processing weights are used to specify the portion of the shared CP resources allocated to an LPAR.
- Each LPAR is treated as a separate workload, managed against the processor weight policy.
- Processor weights define relative priorities of LPARs for determining which image receives the resource when there is contention for that resource.
- When there is no contention for the resource (i.e., all LPARs are not, at the moment, trying to use all of the CP resource they are entitled to), the other logical partitions automatically fill that "white space" and use that excess capacity.

- NOTE:** Even the redistribution of the "white space" is done in accordance with the processor weight policy.
- When the logical partition that was under-utilizing its entitled resource requires it again, the resource moves back to it in accordance with the processor weights. The policy can be dynamically updated, with changes taking effect immediately.

- An LPAR with dedicated CPs is not affected by processing weights.

- NOTE:** Yet a lower level resource manager is when z/OS is executing in an LPAR uses WLM. Multiple workloads within a z/OS image are given an even finer granularity of workload management via the service class component of z/OS. Therefore, the priority of the LPAR with respect to other logical partitions is managed by PR/SM. The priority of the individual workloads and applications within the z/OS logical partition is managed by WLM.

CheatSheet

#46 zTidBits

z10 Processors and Weights

Processing weights are not z/OS specific.

PR/SM (hypervisor)	
zVSE	300
zOSTEST	100
zVM	900

- Although PR/SM always manages sharing LPARs according to the specified weights, there are times when an LPAR will receive either more or less than its processing share:

1. An LPAR will receive more than its processing share when there is excess CP capacity provided it has work to do and other LPARs are not using their share.
2. An LPAR will receive less than its processing share when its workload demand drops below the capacity specified by its weight.
3. An LPAR will not receive more than its processing share when the CP resources for that LPAR are capped.

NOTE: The recommendation is to specify processing weights to satisfy peak requirements of all the LPARs.

Example – consider a z10 with six CPs and three LPARs as follows:

LPAR NAME	Logical CPs	Weight
zVSE	1	300
zOSTEST	6	100
zVM(Linux)	2	900

1300

The processing weights for all active, sharing LPARs are added together and are considered to be 100% of the processing resources available to the shared CPs.

- The share of processing resources for each LP is calculated by dividing the processing weight for each sharing LPAR by the total processing weight.

Example, at peak CP utilization levels, the dispatcher allocates shared processing resources to each of the LPARs as follows:

zVSE	300/1300 = 23.1%
zOSTEST	100/1300 = 7.7%
zVM	900/1300 = 69.2%
	100.0%

Processing weights can range from 1 to 999 (weights of less than 2% difference are not significant).

See #36 zTidBits

- The share of processing resource for each online logical CP with **HiperDispatch** 'disabled' in the logical partition is calculated by dividing the share for each LP by the number of online logical CPs. The share for each logical CP is as follows:

zVSE	23.1 / 1 CP = 23.1%
zOSTEST	7.7 / 6 CPs = 1.3%
zVM	69.2 / 2 CPs = 34.6%

These % s are used to determine preemption priority for I/O interrupts.

- As long as there is excess CP capacity, processing weights have no effect on the CP resources consumed and serves as a symbolic reference to LPAR importance. Weights affect processing only when the number of logical CPs that need processing time is greater than the number of physical CPs available.

Processing weights and shared CP, ICF, zAAP, IFL and zIIP processors:

- Shared general purpose, ICF, zAAP, IFL and zIIP processors are each managed as separate "pools" of physical resources.

- The processing weights assigned to logical partitions using shared ICF, zAAP, IFL, zIIP or general purpose processors are totaled and managed together only with the total weights from all of the logical partitions using the same processor type.

- The calculations shown in above examples are done independently for general purpose, ICF, zAAP, IFL, or zIIP that are actively installed.

Processing weights for logical partitions with multiple shared CP types:

- When a logical partition is defined to use one or more shared general purpose CPs and one or more shared zAAPs / zIIPs, each of the types of logical processors is managed independently.

- The shared general purpose CPs compete with all other shared general purpose CPs defined in other logical partitions in the configuration. The zAAPs compete with all other shared zAAPs defined in other logical partitions.

- General purpose, zAAP, and zIIP processors have a **separate** processing weight specified

- Note:** If WLM weight management (IRD) is being used for such an LPAR, only the weight of the shared general purpose processor portion of the LPAR will be altered by WLM. The specified weight for an LPAR's zAAP / zIIP processors is unaltered by WLM.

- **Effects of processing weights:** First, every time a sharing LPAR is activated, the share of all other active LPARs, using the same processor types, changes because the total of the processing weights has changed. Since the processing share of LPARs can vary, the actual utilization reported by monitors such as RMF can be different from the weights.

- In systems at less than 100% utilization, some LPARs could receive greater than their share if other LPARs are not using their share.

- The number of CPs defined also affects the maximum resource allocated to that LPAR.
- For example,** an LPAR defined to have two CPs on a three-way CPC can never be allocated more than 67% of the CP resources no matter what its processing weight.

- **Capping processing weights:** The PR/SM capping function provides the capability of limiting CPU resource usage for one or more LPARs. The relative processing weight of an LPAR is its capping value.

- A capped LP running at its cap does not have access to the CP resources that are not utilized by other LPs. However, CP resources that are not used by a capped LPAR can be used by other LPARs. **NOTE:** Equitable distribution of CP resources is always maintained.

- Capping values can be dynamically adjusted.
- The capping function can be turned on and off, and provides the capability of specifying capping for individual LPs without a re-IPL of the LPAR.