The many effects of the LPAR weight – Part 1

Fabio Massimo Ottaviani – EPV Technologies

October 2016

1 Introduction

For many years z/OS customers have been used to setting a weight for each logical partition (LPAR) just to tell PR/SM how to manage shared CPs when they are in contention.
The assigned weight expresses the importance of LPARs workloads to be considered in the distribution of the CPU and zIIP capacity.

Generally speaking weight is enforced by PR/SM only when there is not enough power to satisfy all LPARs' demands. As long as the physical CPs have some spare power, all LPARs can use whatever they want. However, if initial capping is set for an LPAR its weight becomes an insurmountable limit to the usable CP capacity. No matter how much capacity is available and not being used by the other LPARs.

A not well known effect of the weight is that, if it's not coherent with the number of logical CPs assigned to the LPAR, it may negatively influence PR/SM dispatching and application performance leading to the so-called "short CP effect".

More recently, with the advent of HiperDispatch, the LPAR weight has also become a key element in the "vertical polarization" of the logical CPs which has again a very strong effect on workload performance.

Finally, LPAR weight is also used when group capacity is used to determine the LPAR minimum entitlement, which is the amount of CPU it can use in the 4-hour rolling average without being soft-capped when the group capacity limit is reached.

In this paper we will discuss the many effects of LPAR weight providing examples and suggestions.