

# Group Capacity and the mystery of the unenforced limit

*Fabio Massimo Ottaviani – EPV Technologies*

April 2012

## 1 Introduction

Most sites pay IBM, and other ISV, software bills based on the WLC software pricing policy; in this policy the license fees depend on the CPU usage (in MSUs) rather than the machine capacity.

CPU usage is calculated based on a 4-hour rolling average<sup>1</sup>; depending on the workload characteristics this value can be much lower than the power of the machine, which is normally sized to guarantee the service levels during a few peak hours.

The “bad news” is that the WLC software license fee is a monthly fee based on the maximum value of the measured 4-hour rolling average. The complexity of today’s systems and workloads, together with human errors can make it very probable that a company would pay for the full capacity of the machine most of the time.

To guarantee the expected savings IBM introduced the possibility to set limits to the MSU which can be used in the 4-hour rolling average:

- by a single LPAR, “defined capacity limit”;
- by a group of LPARs, “group capacity limit”.

The defined capacity limit can be very useful to avoid certain LPARs, normally running non business critical workloads, can increase the software costs.

The group capacity limit is much more important; it can guarantee that you don’t pay more than the limit value (or more than the sum of the limit values if more than one LPAR group has been created). This is the reason why the majority of the z/OS sites use group capacity limits, for protection against the risk of unplanned software costs.

Unfortunately it can happen that the group capacity limit is not enforced as expected leading to undesired results.

After a short introduction to Group Capacity concepts, we will discuss this issue based on the experience of one of our customers.

---

<sup>1</sup> The sum of the measured 4-hour rolling MSU averages for all the LPARs in the CPC.