

Managing CPU Utilization with WLM Resource Groups - Part 1

Fabio Massimo Ottaviani – EPV Technologies

February 2013

1 Introduction

At any z/OS site hardware costs are mostly driven by installed CPU capacity while software costs are mostly driven by installed CPU capacity (OTC) and used CPU (MLC). CPU utilization measurement and control is therefore the key factor required to reduce costs.

This is the reason why every company collects a lot of different measurements related to CPU utilization which are the mandatory input to tuning and capacity planning activities.

In addition (and sometimes as an alternative) to tuning and capacity planning, different capping techniques can be exploited to manage CPU utilization in order to reduce costs.

If you want to control the amount of CPU used at the LPAR level you can choose between hardware capping and defined capacity. By using group capacity you can also extend the defined capacity mechanism in order to potentially control an entire CEC. However if you want to implement a more granular control on the workloads running inside an LPAR, the only available solution are the WLM resource groups.

While WLM Resource groups have existed for many years they have been scarcely used and are not properly understood by many customers. In order to make their utilization easier they have been enhanced in z/OS 1.8 by introducing two new resource group types.

In this paper we will discuss all the WLM resource group types showing how to use them to manage CPU utilization better.