

Where do my batch jobs spend all this time? – Part 1

Fabio Massimo Ottaviani – EPV Technologies

1 Introduction

In an era of smart web applications fully exploiting the virtualization concept to ultimately run in the cloud, it is astonishing to realize that batch processing is still such an important component of the workload of any z/OS environment.

As batch processing is not usually triggered by a user waiting for an immediate response, the importance of batch in performance tuning and capacity planning activities is often underestimated.

Batch jobs usually run overnight, during the so called “batch window”. During this time, online activity is restricted or even completely forbidden.

So a good first reason to tune batch jobs is to be sure that the batch window will end before the agreed online start time. This is becoming more and more important because batch activity is still growing at many sites while the batch window is normally shrinking due to 24X7 online demand.

Other reasons to tune batch jobs are:

- batch activity may encroach on the online shift; even when not delaying the online start, it may degrade overall system performance and increase costs, especially for customers using the WLC software cost policy;
- many batch jobs were designed years ago so it may be possible to achieve substantial performance improvements by using up to date technologies.

The bottom line is that tuning batch jobs is an important task.

In order to do that, the first mandatory step is to be able to answer to a very important question: “Where do my batch jobs spend all this time?”.

In this paper we will try to provide that answer by analyzing in detail the SMF 30 metrics.