Measuring the HiperDispatch activity

Fabio Massimo Ottaviani – EPV Technologies

March 2017

1 Introduction

The capacity of recent mainframe machines is largely determined by the capability of the z/OS system and workload components to exploit the processor cache architecture.

The optimal situation, the lowest number of processor cycles used per instruction, occurs when data and instructions are available in the Level 1 cache.
Performance becomes progressively worse as data and instructions are found in the higher levels of cache, including the memory, and eventually in a different book or node/drawer.

The likelyhood of finding the necessary data and instructions, in the levels of cache closest to the processor, depends on the probability of the same workload running as much as possible on the same logical and physical processor.
To maximize this probability a close cooperation between the z/OS operating system and the PR/SM hypervisor is required. This cooperation has been realized through a highly complex component called HiperDispatch.

In this paper, we will focus on the metrics available in SMF 70 records that have to be used to understand the HiperDispatch activity and measure its effectiveness.
Some real life examples, based on EPV for z/OS views, will also be discussed.
All formulas and examples applies both to standard CPU and to zIIP except those referring to Group Capacity, which applies to CPU only.