Logical processor topology in HiperDispatch mode

Fabio Massimo Ottaviani – EPV Technologies

January 2017

1 Introduction

HiperDispatch's main goal is to maximize the probability for a workload to be re-dispatched to the same group of logical and physical processors, in order to allow it to reuse instructions and data previously loaded in the cache levels closer to the processor.

In order to decide the logical processor topology (mapping to the physical processors) HiperDispatch implements vertical polarization to split the LPAR logical processors in three pools:

- High polarity (high processor share); they will have a target share corresponding to 100% of a physical processor which will be pseudo-dedicated to each of them;

- Medium polarity (medium processor share); they will normally have a target share greater than 0% and less than 100% of a physical processor; these medium logical processors have the remainder of the LPAR's shares after the allocation of the logical processors with the high share; they will compete for the physical processors as before HiperDispatch;

- Low polarity (low processor share); they will receive a target share corresponding to 0% of a physical processor; these are considered discretionary logical processors which are not needed to allow the LPAR to fully utilize the physical processor resource associated with its weight; if there is not unused capacity left from other LPARs in the CEC they will be parked and not used.

By using the information provided in SMF 70 and 113 many consequences of the HiperDispatch activity can be analyzed. However, no information is provided in these records about the logical to physical processors topology.

This information is available in the SMF 99 subtype 14 (SMF 99-14) records and it can be easily collected in an SQL DB by EPV zParser customers.
However the IBM WLM team decided to make processor topology accessible to everyone and provided a free tool: the WLM Topology Report.

The tool provides an Excel spreadsheet that displays lots of interesting information such as:
- the association of logical processors to books, chips, drawers, and nodes,
- the vertical polarization of the processors (high, medium, low),
- the processor type (CPU and zIIP),
- the association to WLM affinity nodes,
- any topology changes.

In this paper, we will show what you have to do in order to install and use the WLM Topology Report. We will also show and discuss two real life examples.