

Understanding WLM General Delays

Fabio Massimo Ottaviani – EPV Technologies

White paper – WLM series

The idea of z/OS Workload Manager is to make a contract between the installation (end user) and the operating system. The installation classifies the work running on the z/OS operating system in distinct Service Classes and defines goals for them that express the expectation of how the work should perform. WLM uses these goal definitions to manage the work across all systems of a sysplex environment. (System Programmer's Guide to Workload Manager, IBM Redbooks, SG24-6472-00, June 2005)

In order to accomplish this complex task, WLM needs to know the amount of time spent by each workload at the different resources (using and waiting for them).

This is the reason why WLM periodically samples the state of each dispatchable unit (TCB and SRB) associated both to an Address Space or an Enclave.

These samples are then aggregated at the service class period level and used to:

- determine what resource is the primary bottleneck for the work in the service class period,
- evaluate the impact of some possible action WLM might take
- make policy adjustment decisions.

This paper will shortly describe the WLM using and delay states.