



Contribution ID: 455

Type: poster presentation

Efficient provisioning for multicore applications with LSF

Tier-1 sites providing computing power for HEP experiments are usually tightly designed for high throughput performances. This is pursued by reducing the variety of supported usecases and tuning for performances those ones, the most important of which have been that of single-core jobs. Moreover, the usual workload is saturation: each available core in the farm is in use and there are queued jobs waiting for their turn to run. Enabling multicore jobs thus requires dedicating a number of hosts where to run, and waiting for them to free the needed number of cores. This drain-time introduces a loss of computing power driven by the number of unusable empty cores.

As an increasing demand for multicore capable resources have emerged, a Task Force have been constituted in WLCG, with the goal to define a simple and efficient multicore resource provisioning model.

This paper details the work done at the INFN Tier1 to enable multicore support for the LSF batch system, with the intent of reducing to the minimum the average number of unused cores.

The adopted strategy has been that of dedicating to multicore a dynamic set of nodes, whose dimension is mainly driven by the number of pending multicore requests and fairshare priority of the submitting user. The node status transition, from single to multi core et vice versa, is driven by a finite state machine which is implemented in a custom multicore director script, running in the cluster.

After describing and motivating both the implementation and the details specific to the LSF batch system, results about performance are reported. Factors having positive and negative impact on the overall efficiency are discussed and solutions to reduce at most the negative ones are proposed.

Author: DAL PRA, Stefano (INFN)

Presenter: DAL PRA, Stefano (INFN)

Track Classification: Track6: Facilities, Infrastructure, Network