

Memory handling in the ATLAS submission system from job definition to sites limits

Monday, October 10, 2016 11:15 AM (15 minutes)

The ATLAS workload management system is a pilot system based on a late binding philosophy that avoided for many years to pass fine grained job requirements to the batch system. In particular for memory most of the requirements were set to request 4GB vmem as defined in the EGI portal VO card, i.e. 2GB RAM + 2GB swap. However in the past few years several changes have happened in the operating system kernel and in the applications that make such a definition of memory to use for requesting slots obsolete and ATLAS has introduced the new PRODSYS2 workload management which has a more flexible system to evaluate the memory requirements and to submit to appropriate queues. The work stemmed in particular from the introduction of 64bit multicore workloads and the increased memory requirements of some of the single core applications. This paper describes the overall review and changes of memory handling starting from the definition of tasks, the way tasks memory requirements are set using scout jobs and the new memory tool produced in the process to do that, how the factories set these values and finally how the jobs are treated by the sites through the CEs, batch systems and ultimately the kernel.

Tertiary Keyword (Optional)

Computing facilities

Secondary Keyword (Optional)

Computing middleware

Primary Keyword (Mandatory)

Distributed workload management

Primary author: FORTI, Alessandra (University of Manchester (GB))

Session Classification: Track 3: Distributed Computing

Track Classification: Track 3: Distributed Computing