



Contribution ID: 13

Type: Poster

Supervised job preemption methodology for controlled memory consumption of jobs running in the ALICE Grid

Wednesday 13 March 2024 16:15 (30 minutes)

The ALICE experiment's Grid resources vary significantly in terms of memory capacity, CPU cores, and resource management. Memory allocation for scheduled jobs depends on the hardware constraints of the executing machines, system configurations, and batch queuing policies. The O2 software framework introduces multi-core tasks where deployed processes share resources. To accommodate these new use cases, most Grid sites provide ALICE with multi-core slots of a customizable amount of cores. The Grid middleware manages the resources within a slot, sub-partitioning and distributing them among allocated jobs. This allows for parallel execution of jobs with different natures and usage patterns within the same resource-sharing slot. From the scheduling system's perspective, this job set is treated as a single unit for resource usage accounting. Overconsumption by any job can lead to the entire slot being killed, terminating all co-executing ALICE jobs. To prevent this and promote job completion with reasonable resource usage, the Grid middleware should implement targeted preemption of top-consuming jobs when overall consumption approaches the system's killing threshold.

This paper analyzes site resource limiting procedures, including killing policies and memory thresholds, and the design of the ALICE Grid middleware framework's methodology for targeted preemption. Preemption decisions are made in real time, considering various factors of running payloads, weighted according to experiment priorities, to maximize efficiency and successful task completion.

Significance

The contribution presents an analysis of Grid site resource allocation limiting procedures, including killing policies and memory thresholds, and the design of the ALICE Grid middleware framework's methodology for targeted preemption of over-consuming jobs. Preemption decisions are made in real time, considering various factors of running payloads, weighted according to experiment priorities, to maximize efficiency and successful task completion.

References

Paper related to multi-core job support in the ALICE Grid - <https://dx.doi.org/10.1088/1742-6596/2438/1/012009>

Experiment context, if any

LHC ALICE Experiment

Primary authors: WIJETHUNGA, Kalana (University of Moratuwa (LK)); BERTRAN FERRER, Marta (CERN)

Presenter: WIJETHUNGA, Kalana (University of Moratuwa (LK))

Session Classification: Poster session with coffee break

Track Classification: Track 1: Computing Technology for Physics Research