

ATLAS Global Shares Implementation in the PanDA Workload Management System

Tuesday 10 July 2018 16:40 (20 minutes)

PanDA (Production and Distributed Analysis) is the workload management system for ATLAS across the Worldwide LHC Computing Grid. While analysis tasks are submitted to PanDA by over a thousand users following personal schedules (e.g. PhD or conference deadlines), production campaigns are scheduled by a central Physics Coordination group based on the organization's calendar. The Physics Coordination group needs to allocate the amount of Grid resources dedicated to each activity, in order to manage sharing of CPU resources among various parallel campaigns and to make sure that results can be achieved in time for important deadlines.

While dynamic and static shares on batch systems have been around for a long time, we are trying to move away from local resource partitioning and manage shares at a global level in the PanDA system. The global solution is not straightforward, given different requirements of the activities (number of cores, memory, I/O and CPU intensity), the heterogeneity of Grid resources (site/HW capabilities, batch configuration and queue setup) and constraints on data locality. We have therefore started the Global Shares project that follows a requirements-driven multi-step execution plan, starting from definition of nestable shares, implementing share-aware job dispatch, aligning internal processes with global shares and finally implementing a pilot stream control for controlling the batch slots while keeping late binding.

This contribution will explain the development work and architectural changes in PanDA to implement Global Shares, and provide an operational point of view with the difficulties we found along the way.

Authors: BARREIRO MEGINO, Fernando Harald (University of Texas at Arlington); DE, Kaushik (University of Texas at Arlington (US)); WALKER, Rodney (Ludwig Maximilians Universitat (DE)); MAENO, Tadashi (Brookhaven National Laboratory (US)); DI GIROLAMO, Alessandro (CERN)

Presenter: BARREIRO MEGINO, Fernando Harald (University of Texas at Arlington)

Session Classification: Posters

Track Classification: Track 3 –Distributed computing