



Contribution ID: 215

Type: Oral

Managing the ATLAS Grid through Harvester

Tuesday, November 5, 2019 11:45 AM (15 minutes)

ATLAS Computing Management has identified the migration of all resources to Harvester, PanDA's new workload submission engine, as a critical milestone for Run 3 and 4. This contribution will focus on the Grid migration to Harvester.

We have built a redundant architecture based on CERN IT's common offerings (e.g. Openstack Virtual Machines and Database on Demand) to run the necessary Harvester and HTCondor services, capable of sustaining the load of O(1M) workers on the grid per day.

We have reviewed the ATLAS Grid region by region and moved as much possible away from blind worker submission, where multiple queues (e.g. single core, multi core, high memory) compete for resources on a site. Instead we have migrated towards more intelligent models that use information and priorities from the central PanDA workload management system and stream the right number of workers of each category to a unified queue while keeping late binding to the jobs.

We will also describe our enhanced monitoring and analytics framework. Worker and job information is synchronized with minimal delays to a CERN IT provided Elastic Search repository, where we can interact with dashboards to follow submission progress, discover site issues (e.g. broken Compute Elements) or spot empty workers.

The result is a much more efficient usage of the Grid resources with smart, built-in monitoring of resources.

Consider for promotion

No

Primary authors: BARREIRO MEGINO, Fernando Harald (University of Texas at Arlington); ALEKSEEV, Aleksandr (National Research Tomsk Polytechnic University (RU)); BERGHAUS, Frank (University of Victoria (CA)); CAMERON, David (University of Oslo (NO)); FILIPCIC, Andrej (Jozef Stefan Institute (SI)); GLUSHKOV, Ivan (University of Texas at Arlington (US)); LIN, Fa-Hui (Academia Sinica (TW)); MAENO, Tadashi (Brookhaven National Laboratory (US)); MAGINI, Nicolo (Iowa State University (US))

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

Session Classification: Track 3 – Middleware and Distributed Computing

Track Classification: Track 3 – Middleware and Distributed Computing