

Benefits and performance of ATLAS approaches to utilizing opportunistic resources

Tuesday, October 11, 2016 2:45 PM (15 minutes)

ATLAS has been extensively exploring possibilities of using computing resources extending beyond conventional grid sites in the WLCG fabric to deliver as many computing cycles as possible and thereby enhance the significance of the Monte-Carlo samples to deliver better physics results.

The difficulties of using such opportunistic resources come from architectural differences such as unavailability of grid services, the absence of network connectivity on worker nodes or inability to use standard authorization protocols. Nevertheless, ATLAS has been extremely successful in running production payloads on a variety of sites, thanks largely to the job execution workflow design in which the job assignment, input data provisioning and execution steps are clearly separated and can be offloaded to custom services. To transparently include the opportunistic sites in the ATLAS central production system, several models with supporting services have been developed to mimic the functionality of a full WLCG site. Some are extending Computing Element services to manage job submission to non-standard local resource management systems, some are incorporating pilot functionality on edge services managing the batch systems, while the others emulate a grid site inside a fully virtualized cloud environment.

The exploitation of opportunistic resources was at an early stage throughout 2015, at the level of 10% of the total ATLAS computing power, but in the next few years it is expected to deliver much more. In addition, demonstrating the ability to use an opportunistic resource can lead to securing ATLAS allocations on the facility, hence the importance of this work goes beyond merely the initial CPU cycles gained.

In this presentation, we give an overview and compare the performance, development effort, flexibility and robustness of the various approaches. Full descriptions of each of those models are given in other contributions to this conference.

Tertiary Keyword (Optional)

Secondary Keyword (Optional)

Primary Keyword (Mandatory)

Data processing workflows and frameworks/pipelines

Primary author: FILIPCIC, Andrej (Jozef Stefan Institute (SI))

Session Classification: Track 3: Distributed Computing

Track Classification: Track 3: Distributed Computing