

ATLAS World-cloud and networking in PanDA

Monday 10 October 2016 11:45 (15 minutes)

The ATLAS computing model was originally designed as static clouds (usually national or geographical groupings of sites) around the Tier 1 centers, which confined tasks and most of the data traffic. Since those early days, the sites' network bandwidth has increased at O(1000) and the difference in functionalities between Tier 1s and Tier 2s has reduced. After years of manual, intermediate solutions, we have now ramped up to full usage of World-cloud, the latest step in the PanDA Workload Management System to increase resource utilization on the ATLAS Grid, for all workflows (MC production, data (re)processing, etc.).

We have based the development on two new site concepts. Nuclei sites are the Tier 1s and large Tier 2s, where tasks will be assigned and the output aggregated, and satellites are the sites that will execute the jobs and send the output to their nucleus.

Nuclei and satellite sites are dynamically paired by PanDA for each task based on the input data availability, capability matching, site load and network connectivity. This contribution will introduce the conceptual changes for World-cloud, the development necessary in PanDA, an insight into the network model and the first half year of operational experience.

Secondary Keyword (Optional)

Computing models

Primary Keyword (Mandatory)

Distributed workload management

Tertiary Keyword (Optional)

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

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

Session Classification: Track 3: Distributed Computing

Track Classification: Track 3: Distributed Computing