

PanDA for ATLAS distributed computing in the next decade

Monday, 10 October 2016 11:00 (15 minutes)

The Production and Distributed Analysis (PanDA) system has been developed to meet ATLAS production and analysis requirements for a data-driven workload management system capable of operating at the Large Hadron Collider (LHC) data processing scale. Heterogeneous resources used by the ATLAS experiment are distributed worldwide at hundreds of sites, thousands of physicists analyse the data remotely, the volume of processed data is beyond the exabyte scale, dozens of scientific applications are supported, while data processing requires more than a few billion hours of computing usage per year. PanDA performed very well over the last decade including the LHC Run 1 data taking period. However, it was decided to upgrade the whole system concurrently with the LHC's first long shutdown in order to cope with rapidly changing computing infrastructure. After two years of reengineering efforts, PanDA has embedded capabilities for fully dynamic and flexible workload management. The static batch job paradigm was discarded in favor of a more automated and scalable model. Workloads are dynamically tailored for optimal usage of resources, with the brokerage taking network traffic and forecasts into account. Computing resources are partitioned based on dynamic knowledge of their status and characteristics. The pilot has been re-factored around a plugin structure for easier development and deployment. Bookkeeping is handled with both coarse and fine granularities for efficient utilization of pledged or opportunistic resources. Leveraging direct remote data access and federated storage relaxes the geographical coupling between processing and data. An in-house security mechanism authenticates the pilot and data management services in off-grid environments such as volunteer computing and private local clusters. The PanDA monitor has been extensively optimized for performance and extended with analytics to provide aggregated summaries of the system as well as drill-down to operational details. There are as well many other challenges planned or recently implemented, and adoption by non-LHC experiments such as bioinformatics groups successfully running Paleomix (microbial genome and metagenomes) payload on supercomputers. In this talk we will focus on the new and planned features that are most important to the next decade of distributed computing workload management.

Tertiary Keyword (Optional)

Secondary Keyword (Optional)

Primary Keyword (Mandatory)

Distributed workload management

Primary author: MAENO, Tadashi (Brookhaven National Laboratory (US))

Session Classification: Track 3: Distributed Computing

Track Classification: Track 3: Distributed Computing