

21st International Conference on Computing in High Energy and Nuclear Physics (CHEP2015)



Contribution ID: 537

Type: oral presentation

Achieving production-level use of HEP software at the Argonne Leadership Computing Facility

Thursday, April 16, 2015 11:45 AM (15 minutes)

HEP's demand for computing resources has grown beyond the capacity of the Grid, and these demands will accelerate with the higher energy and luminosity planned for Run II. Mira, the ten petaflops supercomputer at the Argonne Leadership Computing Facility, is a potentially significant compute resource for HEP research. Through an award of fifty million hours on Mira, we have delivered millions of events to LHC Experiments by establishing the means of marshaling jobs through serial stages on local clusters, and parallel stages on Mira. We are running several HEP applications, including Alpgen, Pythia, Sherpa, and Geant4. Event generators, such as Sherpa, typically have a split workload: a small scale integration phase, and a second, more scalable, event-generation phase. To accommodate this workload on Mira we have developed two Python-based Django applications, Balsam and ARGO. Balsam is a generalized scheduler interface which uses a plugin system for interacting with scheduler software such as Condor, Cobalt, and Torque. ARGO is a workflow manager that submits jobs to instances of Balsam. Through these mechanisms, the serial and parallel tasks within jobs are executed on the appropriate resources. This approach and its integration with the PanDA production system will be discussed.

Primary authors: BENJAMIN, Doug (Duke University (US)); CHILDERS, Taylor (Argonne National Laboratory (US)); LE COMPTE, Thomas (Argonne National Laboratory (US)); URAM, Tom (urn:Google)

Presenter: URAM, Tom (urn:Google)

Session Classification: Track 4 Session

Track Classification: Track4: Middleware, software development and tools, experiment frameworks, tools for distributed computing