

Pyglidein - a simple HTCondor glidein service

Thursday 13 October 2016 16:30 (15 minutes)

A major challenge for data production at the IceCube Neutrino Observatory presents itself in connecting a large set of small clusters together to form a larger computing grid. Most of these clusters do not provide a Grid interface. Using a local account on each submit machine, HTCondor glideins can be submitted to virtually any type of scheduler. The glideins then connect back to a main HTCondor pool, where jobs can run normally with no special syntax. To respond to dynamic load, a simple server advertises the number of idle jobs in the queue and the resources they request. The submit script can query this server to optimize glideins to what is needed, or not submit if there is no demand. Configuring HTCondor dynamic slots in the glideins allows us to efficiently handle varying memory requirements as well as whole-node jobs.

One step of the IceCube simulation chain, photon propagation in the ice, heavily relies on GPUs for faster execution. Therefore, one important requirement for any workload management system in IceCube is that it can handle GPU resources. Within the pyglidein system, we have successfully configured HTCondor glideins to use any GPU allocated to it, with jobs using the standard HTCondor GPU syntax to request and use a GPU. This mechanism allows us to seamlessly integrate our local GPU cluster with remote non-Grid GPU clusters, including specially allocated resources at XSEDE Supercomputers.

Tertiary Keyword (Optional)

Secondary Keyword (Optional)

Computing middleware

Primary Keyword (Mandatory)

Data processing workflows and frameworks/pipelines

Author: SCHULTZ, David (University of Wisconsin-Madison)

Co-authors: Dr RIEDEL, Benedikt (University of Alberta); MERINO, Gonzalo (IceCube)

Presenter: SCHULTZ, David (University of Wisconsin-Madison)

Session Classification: Posters B / Break

Track Classification: Track 7: Middleware, Monitoring and Accounting