

Feature Updates In Pyglidein, An HTCondor Glidein Generator

Tuesday, July 10, 2018 4:40 PM (20 minutes)

IceCube is a cubic kilometer neutrino detector located at the south pole. IceCube's simulation and production processing requirements far exceed the number of available CPUs and GPUs in house. Collaboration members commit resources in the form of cluster time at institutions around the world. IceCube also signs up for allocations from large clusters in the United States like XSEDE. All of these disparate cluster resources are homogenized with IceCube's own Glidein software, Pyglidein. Pyglidein uses a pull model to launch HTCondor glideins around the world, which reduces the need for complicated firewall configuration changes at glidein sites. We present the most recent Pyglidein feature enhancements, including improvements to monitoring and integration testing. Glidein jobs now ship logs directly to S3 using presigned PUT URLs, monitor the state of resources via HTCondor startd cron jobs, and send metrics back to IceCube. The increased visibility reduces debugging time and frustration from administrators and data processors. The Pyglidein client has been containerized alongside HTCondor, Torque, and SLURM to automate integration testing across all three schedulers. New features can be confidently released quickly without disrupting processing jobs.

Primary author: SKARLUPKA, Heath (University of Wisconsin)

Co-authors: MERINO, Gonzalo (IceCube); BRIK, Vladimir (University of Wisconsin at Madison); SCHULTZ, David (University of Wisconsin-Madison)

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

Session Classification: Posters

Track Classification: Track 3 –Distributed computing