


HTCondor Accounting Update

John Gordon
APEL Team
December 2016

WLCG Operations Coordination



HTCondor Accounting

- Disclaimer: the following is my current understanding. I haven't always approached the definitive sources.

HTCondor

- **HTCondor** is an open-source high-throughput computing software framework for coarse-grained distributed parallelization of computationally intensive tasks. It can be used to manage workload on a dedicated cluster of computers, and/or to farm out work to idle desktop computers – so-called cycle scavenging.
- [HTCondor - Wikipedia](#)
- <https://en.wikipedia.org/wiki/HTCondor>

What is HTCondor?

- HTCondor is a specialized workload management system for compute-intensive jobs. Like other full-featured batch systems, HTCondor provides a job queueing mechanism, scheduling policy, priority scheme, resource monitoring, and resource management. Users submit their serial or parallel jobs to HTCondor, HTCondor places them into a queue, chooses when and where to run the jobs based upon a policy, carefully monitors their progress, and ultimately informs the user upon completion.
- <https://research.cs.wisc.edu/htcondor/description.html>

Why HTCondor?

- Opensource – LSF and GE are commercial
- Torque/Maui was reaching the limits of scalability with cluster size
- Alternatives? SLURM, UGE
- HTCondor seems to work in US LHC sites

Support for CEs

- ARC
 - Recent versions of ARC CE support HTCondor allowing publishing accounting to APEL via JURA.
 - Evidence that it loses track of flocked jobs
 - The usual issues around re-publishing
- Gratia
 - Gratia supports HTCondor which is used by US/OSG sites.
 - Gratia publishes LHC VO Summaries to APEL(*).

(*) Gratia also include a small number of other VOs

Support for CEs

- CREAM

- An APEL batch parser for Condor exists and was released but we don't think anyone has deployed it yet.
- Does CREAM still support Condor as a backend?
- Andrew Lahiff RAL/CMS wrote a plug-in which produces PBS-like job records from HTCondor which can then be parsed by the APEL PBS parser.
- In theory this could be much simpler than other batch systems as the identity information is kept in the job history files and not separately in blah logs like CREAM.
- A simpler APEL client workflow and independent of the BDII

Support for CEs

- (HT)CondorCE
 - Currently no APEL Support
 - Ian Steers at CERN instrumented accounting in the CondorCE which writes to a local accounting system. They parse the job history files which also contain the X.509 identity.
 - It also produced APEL Job Records which it successfully sent by SSM to APEL
 - This was developed separately from LSF accounting at CERN but has now been integrated into the new CERN accounting which is just starting to send data to APEL.
 - A parser based on Ian's could be integrated with the apel client or developed as a standalone publisher like ARC.

Summary

- The ARC CE already supports HTCondor
- Work is being planned for an APEL Parser for HTCondor
- A parser written to get all info from the job history files could be CE-neutral and so used with CREAM or HTCondorCE