

EGI Clinics: HTCondor-CE

Tuesday 3 November 15:30 CET

Experiences / Tips from sites that recently switched from CREAM to HTCondor-CE:

- KIT and DESY: <https://indico.cern.ch/event/936993/contributions/4020907/>
- INFN : <https://indico.cern.ch/event/936993/contributions/4020900/>

Questions and Answers

- (Name, affiliation): (Comment or question)
- (Jeff Templon, Nikhef): What commands / APIs does the HTCondor APEL thing use? We don't use APEL now and don't plan to in the future, we do need accounting (also for our own HTCondor batch systems). How does one do this?

A: (Brian's answer): HTCondor has a lot of accounting tools already, so that is most likely where you want to go for internal accounting. This is done by either parsing the job history files or reading `condor_history` directly. For running jobs, you will want to run `condor_q` instead, which takes the same arguments as `condor_history`. Each job contains the full job ad which contains a lot of information -- you can pass arguments to `condor_q` to get any of these -- run "`condor_q -help`" to see all the options.

(Todd Tannenbaum): all `condor_q` does is query the schedd; without arguments, it just displays a small subset of attributes, but you can also ask for additional information or pass constraints (to select a subset of jobs). `condor_q` also has formatting options if you don't like the default output format. HTCondor also includes the `condor_gangliad` which can push information to Ganglia.

- (Evgeny Stambulchik, WIS): Any tips, pointers etc to start going with APEL and BDII with non-HTCondor batch systems - manually.

A: (Brian Lin) I believe that modifying the APEL accounting would be a fairly simple task as there aren't very many dependencies between the CE and local batch. However, BDII would be a much larger project since the LDIFs between CE and local batch are tightly integrated. Let's continue this technical discussion on htcondor-users@cs.wisc.edu or <https://github.com/htcondor/htcondor-ce/issues/>

- (Victor Mendoza) Can an HTCondor-CE be configured to route jobs to several HTCondor Central Managers?

A: Do you mean multiple at a time, i.e. duplicating jobs? It can't duplicate jobs, but you can route a job to any schedd or other batch system submit host. If you have administrative domain over each pool, though, I (Brian Lin) would recommend having a single HTCondor-CE on a submit host that is configured to "flock" to each of those other pools

- (Victor Mendoza) Is fifemon still the best approach for monitoring (analytics) an HTCondor cluster? what about analytics for HTCondor-CE?

A: (Todd Tannenbaum, paraphrased by Brian Lin) Yes, fifemon is an excellent solution for monitoring your local HTCondor cluster!

A: (Brian Lin) The htcondor-ce-view package has a web service that has some monitoring information available. This has mainly been used in the OSG so we'd love feedback! (Victor mentioned that they initially tried the HTCondor-CE but it used up 100% of a CPU so they abandoned it, Brian suggested that if they can find the configuration that caused this, they should open a GitHub issue <https://github.com/htcondor/htcondor-ce/issues/>)

- (Ghita, CC-IN2P3): What is the best way to allow local users to submit jobs to our local grid cluster and account for their usage? (paraphrased)

A: (Brian Lin) HTCondor-CE is designed specifically for pilot jobs so you should avoid having users directly submit jobs to an HTCondor-CE. Instead, you should set up a submit host for that cluster specifically for your local users. You can configure this submit host to enforce GSI authentication if that's how you account for VO usage. Alternatively, you could assign each user to a VO-based HTCondor accounting group and forego the GSI authentication requirement.