

# Minutes of the ABP Computing Working Group meeting

31st August 2017

B. Jones reported about the first month of production with HTCondor. 100k cores are currently available through HTCondor. Errors when submitting jobs, when copying files and of jobs being killed before completion were reported. Often the errors occurring when submitting a job are linked to the AFS token that is needed by the scheduler to write files. The latency when submitting jobs has been addressed by introducing sub-collectors to reduce the load on the main job collector. To further improve, a second negotiator will be introduced, as a result each pool will have its own fair-share mechanism. Jobs may stay on hold when transferring files if the AFS or kerberos tokens are expired / lost. In particular a bug related to the AFS token was solved. As a side effect, this bug was overloading the scheduler, resulting in further issue of latency when submitting the job.

The jobs are killed before completion either because of wall time or memory. The proper queue has to be set for the wall time, and the memory can be adjusted with a dedicated command. It occurred that some MAD-X jobs did not get the default 2000Mb and were killed, this bug is rare and understood, but it can be avoided by specifying *request\_memory=2000* in the job description. It is possible to check the requirement of a given job with *condor\_q* (running jobs) or *condor\_hist* (stopped jobs).

I/O issues are usually linked to the file system (EOS/AFS). Some jobs used to fail during I/O operations due to credential renewal issues, this issue was fixed.

The option *-spool* when submitting a job allows to leave the output files on the scheduler, preventing the transfer of large amount of files.

B. Jones asked whether a solution based on Docker would be interesting for people using HTCondor directly from the desktop. R. De Maria said answered that for now it is easy to setup HTCondor on a desktop and allows to bypass issues with other files systems. B. Jones said that it is good strategy, yet IT will not provide support. R. De Maria mentioned that if the overhead is small, a Docker could be useful. A. Mergetthi said that he is very interested in the Docker solution for BOINC.

R. De Maria mentioned that he experienced jobs disappearing from the queue but were never completed or killed. B. Jones answered that it is rare but could occur in specific failure of the worker node. In that case, IT should be capable of finding the jobs in their own log.

L. Methers experienced a similar issue, where jobs seem to restart from the start when approaching the end of the procedure. B. Jones said that in some specific conditions the job can complete and be restarted, on purpose either because it is requested by the user or because a failure is detected internally and therefore the jobs is restarted. The log file of those files should contain the information on what happened.