

21st International Conference on Computing in High Energy and Nuclear Physics (CHEP2015)



Contribution ID: 6

Type: poster presentation

How much higher can HTCondor fly

The HTCondor batch system is heavily used in the HEP community as the batch system for several WLCG resources. Moreover it is the backbone of the GlideInWMS, the main pilot system used by CMS. To prepare for LHC Run 2, we are probing the scalability limits of new versions and configurations of HTCondor with the goal of reaching at least 200,000 simultaneous running jobs in a single pool.

A sleeper pool of this size was achieved without a major impact in real jobs by using only 10,000 real slots distributed at several WLCG sites. We will report on how this was made and the impact it has in future scalability tests, not only of HTCondor but of any web faced service.

High configurability is one of the main capabilities of HTCondor. In addition to the test conditions and the tested topology, we include the suggested configuration options used to obtain the scaling results. Finally, we will list the features present in newer versions of HTCondor that allow for sustained operations at scales well beyond what was previously possible.

Primary authors: FAJARDO HERNANDEZ, Edgar (Univ. of California San Diego (US)); DOST, Jeffrey Michael (Univ. of California San Diego (US))

Co-authors: TIRADANI, Anthony (Fermilab); BOCKELMAN, Brian Paul (University of Nebraska (US)); Dr HOLZMAN, Burt (Fermi National Accelerator Lab. (US)); Dr MASON, David Alexander (Fermi National Accelerator Lab. (US)); Mr FREY, Jaime (University of Wisconsin); LETTS, James (Univ. of California San Diego (US)); TANNENBAUM, Todd (Univ of Wisconsin-Madison, Wisconsin, USA)

Presenter: FAJARDO HERNANDEZ, Edgar (Univ. of California San Diego (US))

Track Classification: Track4: Middleware, software development and tools, experiment frameworks, tools for distributed computing