

Reaching new scales in the CMS Global pool

Thursday, September 21, 2023 11:25 AM (20 minutes)

The computing resource needs of LHC experiments, such as CMS, are expected to continue growing significantly over the next decade, during the Run 3 and especially the HL-LHC era. The SI team manages a set of federated HTCondor pools, currently aggregating around 400k CPU cores distributed worldwide, supporting the simultaneous execution of over 200k CMS computing tasks. In order to detect and overcome performance degradation driven by scalability barriers, the SI team regularly runs tests to explore the scalability reach of our infrastructure. In this contribution, we will report on the test results for potential scalability limitations of our infrastructure.

Desired slot length

Speaker release

Yes

Primary authors: PEREZ-CALERO YZQUIERDO, Antonio (Centro de Investigaciones Energéticas Medioambientales y Tecnológicas); KIZINEVIC, Edita (CERN); KHAN, Farrukh Aftab (Fermi National Accelerator Lab. (US)); KIM, Hyunwoo (Fermi National Accelerator Lab. (US)); MASCHERONI, Marco (Univ. of California San Diego (US)); TSIPINAKIS, Nikos (CERN)

Presenter: MASCHERONI, Marco (Univ. of California San Diego (US))

Session Classification: HTCondor User Presentations

Track Classification: HTCondor user presentations