



Contribution ID: 251

Type: Talk

Optimisation of ATLAS computing resource usage through a modern HEP Benchmark Suite via HammerCloud and PanDA

Wednesday 23 October 2024 17:09 (18 minutes)

In April 2023 HEPscore23, the new benchmark based on HEP specific applications, was adopted by WLCG, replacing HEP-SPEC06. As part of the transition to the new benchmark, the CPU core power published by the sites needed to be compared with the effective power observed while running ATLAS workloads. One aim was to verify the conversion rate between the scores of the old and the new benchmark. The other objective was to understand how the HEPscore performs when run on multi-core job slots, so exactly like the computing sites are being used in the production environment. Our study leverages the HammerCloud infrastructure and the PanDA Workload Management System to collect a large benchmark statistic across 136 computing sites using an enhanced HEP Benchmark Suite. It allows us to collect not only performance metrics, but, thanks to plugins, it also collects information such as machine load, memory usage and other user-defined metrics during the execution and stores it in an OpenSearch database. These extensive tests allow for an in-depth analysis of the actual, versus declared computing capabilities of these sites. The results provide valuable insights into the real-world performance of computing resources pledged to ATLAS, identifying areas for improvement while spotlighting sites that underperform or exceed expectations. Moreover, this helps to ensure efficient operational practices across sites. The collected metrics allowed us to detect and fix configuration issues and therefore improve the experienced performance.

Primary authors: DI GIROLAMO, Alessandro (CERN); LORY, Alexander (Ludwig Maximilians Universität (DE)); GIORDANO, Domenico (CERN); MENENDEZ BORGE, Gonzalo (CERN); VUKOTIC, Ilija (University of Chicago (US)); GLUSHKOV, Ivan (Brookhaven National Laboratory (US)); SZCZEPANEK, Natalia Diana (CERN)

Presenter: SZCZEPANEK, Natalia Diana (CERN)

Session Classification: Parallel (Track 4)

Track Classification: Track 4 - Distributed Computing