

Reproducible High-Throughput HPC Workloads with AiiDA

Thursday, October 22, 2020 10:15 AM (30 minutes)

The ever-growing availability of computing power and the sustained development of advanced computational methods present new challenges driven by the sheer amount of calculations and data to manage. In this talk I will present AiiDA (aiida.net), a robust open-source high-throughput infrastructure addressing the challenges arising from the needs of automated workflow management and data provenance recording. I will discuss the developments and capabilities required to reach sustained performance, with AiiDA supporting throughputs of tens of thousands processes/hour, each of them ranging from local processes to massive HPC jobs. I will also discuss how AiiDA automatically preserves and stores the full data provenance in a relational database making it queryable and traversable, thus enabling high-performance data analytics and reproducibility of the research.

Presenter: PIZZI, Giovanni (EPFL)