

# Managing an heterogeneous scientific computing cluster with cloud-like tools: ideas and experience

*Tuesday, 10 July 2018 16:40 (20 minutes)*

Current computing paradigms often involve concepts like microservices, containerisation and, of course, Cloud Computing.

Scientific computing facilities, however, are usually conservatively managed through plain batch systems and as such can cater to a limited range of use cases. On the other side, scientific computing needs are in general orthogonal to each other in several dimensions.

We have been operating the Open Computing Cluster for Advanced data Manipulation (OCCAM), a multi-purpose heterogeneous HPC cluster, for more than one year adopting a cloud-like paradigm. Each computing application is mapped to a dynamically expandable virtual farm, tuned and configured to the application's needs and able to access special hardware like GPU accelerators or low-latency networks as needed, thus delivering computational frameworks that are well consolidated within the community (for a smooth end user experience) while leveraging modern computing paradigms.

By using mostly mainstream software tools like Docker (used throughout our architecture to run both service and computational tasks), Calico for virtual network management, Mesos and Marathon for orchestration and by exploiting some of the work made in the context of the INDIGO-DataCloud project, we aimed at minimising the development and maintenance effort, while using a high-quality software stack.

In this work we present the status of the system, operational experience, lessons learnt, and our outlook for further development. We will also present some preliminary performance comparisons between containerized and bare-metal scientific computing applications in an HPC environment.

**Primary authors:** BAGNASCO, Stefano (Istituto Nazionale di Fisica Nucleare, Torino); VALLERO, Sara (Università e INFN Torino (IT)); LUSSO, Stefano (INFN-TO); Mr CONCAS, Matteo (INFN e Politecnico di Torino (IT)); Prof. ALDINUCCI, Marco (Dept. of Computer Science, University of Torino ); Dr RABELLINO, Sergio (Dept. of Computer Science, University of Torino )

**Presenter:** BAGNASCO, Stefano (Istituto Nazionale di Fisica Nucleare, Torino)

**Session Classification:** Posters

**Track Classification:** Track 7 –Clouds, virtualization and containers