Contribution ID: **232** Type: **Oral**

OCCAM: a flexible, multi-purpose and extendable HPC cluster

Monday, October 10, 2016 11:30 AM (15 minutes)

Obtaining CPU cycles on an HPC cluster is nowadays relatively simple and sometimes even cheap for academic institutions. However, in most of the cases providers of HPC services would not allow changes on the configuration, implementation of special features or a lower-level control on the computing infrastructure and networks, for example for testing new computing patterns or conducting research on HPC itself. The variety of use cases proposed by several departments of the University of Torino, including ones from solid-state chemistry, high-energy physics, computer science, big data analytics, computational biology, genomics and many others, called for different and sometimes conflicting configurations; furthermore, several R&D activities in the field of scientific computing, with topics ranging from GPU acceleration to Cloud Computing technologies, needed a platform to be carried out on.

The Open Computing Cluster for Advanced data Manipulation (OCCAM) is a multi-purpose flexible HPC cluster designed and operated by a collaboration between the University of Torino and the Torino branch of the Istituto Nazionale di Fisica Nucleare. It is aimed at providing a flexible, reconfigurable and extendable infrastructure to cater to a wide range of different scientific computing needs, as well as a platform for R&D activities on computational technologies themselves. Extending it with novel architecture CPU, accelerator or hybrid microarchitecture (such as forthcoming Intel Xeon Phi Knights Landing) will be as a simple as plugging a node in a rack.

The initial system counts slightly more than 1100 cpu cores and includes different types of computing nodes (standard dual-socket nodes, large quad-sockets nodes with 768 GB RAM, and multi-GPU nodes) and two separate disk storage subsystems: a smaller high-performance scratch area, based on the Lustre file system, intended for direct computational I/O and a larger one, of the order of 1PB, to archive near-line data for archival purposes. All the components of the system are interconnected through a 10Gb/s Ethernet layer with one-level topology and an InfiniBand FDR 56Gbps layer in fat-tree topology.

A system of this kind, heterogeneous and reconfigurable by design, poses a number of challenges related to the frequency at which heterogeneous hardware resources might change their availability and shareability status, which in turn affect methods and means to allocate, manage, optimize, bill, monitor VMs, virtual farms, jobs, interactive bare-metal sessions, etc.

This poster describes some of the use cases that prompted the design ad construction of the HPC cluster, its architecture and a first characterization of its performance by some synthetic benchmark tools and a few realistic use-case tests.

Secondary Keyword (Optional)

High performance computing

Primary Keyword (Mandatory)

Computing facilities

Tertiary Keyword (Optional)

Primary authors: ALDINUCCI, Marco (Computer Science Department, University of Torino); PASTERIS, Paolo (Computer Science Department, University of Torino); RABELLINO, Sergio (Computer Science Department, University of Torino); BAGNASCO, Stefano (Universita e INFN Torino (IT)); LUSSO, Stefano (Istituto Nazionale di Fisica Nucleare, Torino)

Presenter: BAGNASCO, Stefano (Universita e INFN Torino (IT))

Session Classification: Track 6: Infrastructures

Track Classification: Track 6: Infrastructures