Contribution ID: **63** Type: **Full presentation: long (25 mins.)**

RECAS-BARI: new high-level services for eScience researchers

Tuesday, 3 November 2020 10:25 (20 minutes)

In the last 5 years the RECAS-BARI datacenter has been offering to its users an increasing amount of compute and storage resources through the local batch system and the on-premise Openstack-based cloud infrastructure. Our users come from different scientific communities (HEP, bioinformatics, medical physics, etc.) and some local SMEs: they need to execute their workloads with different requirements and ask for different levels of technical support.
In the last period, the emergence of ML techniques applied to diverse research areas has increased the need to access specialised hardware devices like GPUs and Infiniband.
In parallel, containers are gaining traction among users as this lightweight virtualization technology dramatically simplifies the distribution and deployment of their software encapsulating the runtime dependencies in a single package.
Since GPUs are not available yet in our cloud, users have been using our batch system (slurm-based) to access these specialized hw devices. They had to learn how to interact with the system and access the datacenter LAN from a bastion in order to submit their jobs.
One of the most common wishes of our users is of course to access resources in a transparent and easy way. On the other side, our admins and support team desire to limit the manual operations and configurations needed to operate the cluster and support the users in their daily activities. To this purpose we have decided to install and manage some of the new compute nodes (GPU equipped) acquired with the IBISCO PON under a Mesos cluster and to adopt the INDIGO/DEEP solutions to provide high-level interfaces to the end users. We have developed a set of ansible roles for setting up the Mesos cluster with the needed configuration, including the support for the OpenID connect authentication and for exploiting GPUs. The Mesos cluster has been integrated with the RECAS-BARI PaaS Orchestration system in order to facilitate the user interaction: the Orchestrator and its dashboard hide the complexity of managing Mesos tasks offering a transparent access to almost all the functionalities provided by the cluster. Currently we are working on hardening the user isolation in an environment that is multi-tenant natively and addressing the security aspects related to the use of docker containers.

Primary authors: DONVITO, Giacinto (INFN); VINO, Gioacchino (INFN); ANTONACCI, Marica (INFN); PERNIOLA, Michele (UNIBA)

Presenter: ANTONACCI, Marica (INFN)

Session Classification: Data analytics and thematic services - part 2