Contribution ID: **54** Type: **not specified**

Improving IaaS resources to accommodate scientific applications

Wednesday 14 October 2015 11:50 (20 minutes)

INDIGO-DataCloud aims at developing a data and computing platform targeted at scientific communities, integrating existing standards and open software solutions. INDIGO proposes:

- to build-up a PaaS solution leveraging existing resources and e-Infrastructures, since the mere access to IaaS resources has been demonstrated as not being a realistic option for most Research Communities
- to improve the virtualization layer provided by already diffed IaaS

In order to build such a Scientific oriented PaaS, several improvements are needed at the IaaS layer. However, the approach of the project is to ensure that those improvements are beneficial enough by their own and not focused on just enabling the upper PaaS layer. Therefore, we find that they can be interesting for the HEP community and for the WLCG in particular.

In this presentation we will try to give an overview of the developments that we consider of interest for the HEPiX community, namely:

- Container support in Cloud stacks: INDIGO will incorporate specific drivers into the most used cloud management frameworks (i.e. OpenNebula and OpenStack) to support the deployment and execution of containers as first-class resources on the IaaS.
- Scheduling improvements: Currently scheduling in most cloud middlewares is based in first-come, first-served policies. The project will develop new intra-cloud scheduling algorithms that are more appropriate for scientific scenarios. Namely we foresee two different and complimentary approaches: a fair-share based scheduling with a queuing mechanism similar to the current batch scheduling and a spot-instances based approach for opportunistic usage.
- IaaS orchestration trough standards: We will provide IaaS orchestration using a common standard language (TOSCA) in both OpenStack and OpenNebula, easing the deployment and management of (computing, storage and network) resources at the IaaS layer.
- Container integration in batch systems: We will study the possibility of executing containers through batch systems, with the aim of providing the access to specialized hardware (such as Infiniband networks and GPGPUs).

Length of presentation (max. 20 minutes)

20

Authors: LOPEZ GARCIA, Alvaro (Unknown); DONVITO, Giacinto (INFN-Bari); FUHRMANN, Patrick (DESY)

Presenter: CHIERICI, Andrea (INFN-CNAF)

Session Classification: Grid, Cloud and Virtualization

Track Classification: Grid, Cloud & Virtualisation