

TOSCA-based orchestration of complex clusters at the IaaS level

Thursday, 13 October 2016 11:30 (15 minutes)

The INDIGO-DataCloud project's ultimate goal is to provide a sustainable European software infrastructure for science, spanning multiple computer centers and existing public clouds.

The participating sites form a set of heterogeneous infrastructures, some running OpenNebula, some running OpenStack. There was the need to find a common denominator for the deployment of both the required PaaS services and the end user applications. CloudFormation or Heat were technically viable options, but tied to specific implementations. The TOSCA Simple Profile in YAML v1.0 specification, on the other hand, is on its way to becoming a standard to describe the topology of Cloud applications, with growing support in different communities.

In the context of INDIGO-DataCloud, TOSCA Templates are used at the IaaS level to describe complex clusters of applications and services and to provide a way to express their automatic configuration via Ansible recipes. Within OpenStack, the TOSCA interface is implemented in the Heat orchestrator service, while in OpenNebula it is implemented using the Infrastructure Manager (IM), an open-source tool to deploy virtual infrastructures on multiple Clouds.

In a IaaS context both Heat and IM are very useful to ease portable provisioning of resources and deployment of services and applications on dynamically instantiated clusters. One of the advantages of using TOSCA and the IM/Heat approach is that the INDIGO PaaS layer can easily exploit it across different IaaS implementations, increasing the portability of the cluster definitions, and implementing the provisioning of the required services across multiple IaaS infrastructures through the INDIGO orchestrator.

In this contribution we will outline our enhancements for the TOSCA support in both OpenStack and OpenNebula, done together in close collaboration with industry partners such as IBM. These developments, emerging from the INDIGO requirements, have been contributed upstream to the relevant tools, as they have been considered of general interest. Moreover, we will showcase how it is possible to deploy an elastic cluster managed by a batch system like SLURM or Torque, where nodes are dynamically added and removed from the cluster to adapt to the workload.

Primary Keyword (Mandatory)

Cloud technologies

Secondary Keyword (Optional)

Tertiary Keyword (Optional)

Primary authors: MOLTÓ, Germán (Universidad Politécnica de Valencia); VELTEN, Mathieu (CERN); CABALLER, Miguel (UPV); BRITO DA ROCHA, Ricardo (CERN)

Co-author: DONVITO, Giacinto (INFN-Bari)

Presenter: BRITO DA ROCHA, Ricardo (CERN)

Session Classification: Track 6: Infrastructures

Track Classification: Track 6: Infrastructures