

Mesos exploitation within INDIGO-DataCloud for multi-site long-running services, computational workloads and advanced PaaS features

Thursday 13 October 2016 16:30 (15 minutes)

The European project INDIGO-DataCloud aims at developing an advanced computing and data platform. It provides advanced PaaS functionalities to orchestrate the deployment of Long-Running Services (LRS) and the execution of jobs (workloads) across multiple sites through a federated AAI architecture.

The multi-level and multi-site orchestration and scheduling capabilities of the INDIGO PaaS layer are presented in this contribution, highlighting the benefits introduced by the project to already available infrastructures and data centers.

User application/service deployment requests are expressed using TOSCA, an OASIS standard to specify the topology of services provisioned in IT infrastructures; the TOSCA template describing the application/service deployment is processed by the INDIGO Orchestrator, implementing a complex workflow aimed at fulfilling a user request using information about the health status and capabilities of underlying IaaS and their resource availability, QoS/SLA constraints, the status of the data files and storage resources needed by the service/application. This process allows to achieve the best allocation of the resources among multiple IaaS sites.

On top of the enhanced Cloud Management Frameworks scheduling capabilities developed by the project and described in other contributions, a two-level scheduling exploiting Apache Mesos has been implemented, where the Orchestrator is able to coordinate the deployment of applications/services on top of one or more Mesos clusters.

Mesos allows to share cluster resources (CPU, RAM) across different distributed applications (frameworks) organizing the cluster architecture in two sets of nodes: masters coordinating the work, and slaves executing it.

INDIGO uses and improves two already available Mesos frameworks: Marathon, which allows to deploy and manage LRS, and Chronos, which allows to execute jobs. Important features that are currently missing in Mesos and that are being added by INDIGO include: the elasticity of a Mesos cluster so that it can automatically shrink or expand depending on the tasks queue, the automatic scaling of the user services running on top of the Mesos cluster and a strong authentication mechanism based on OpenID-Connect. Docker containers are widely used in order to simplify the installation and configuration of both services and applications. A further application of this architecture and of these enhancements addresses one of the objectives of the INDIGO project, namely to provide a flexible Batch Systems as a Service, i.e. the possibility to request and deploy a virtual cluster on-demand for submitting batch jobs. To this purpose, the INDIGO team is implementing the integration of HTCondor with Mesos and Docker, as described in detail in another contribution.

Tertiary Keyword (Optional)

Secondary Keyword (Optional)

Distributed workload management

Primary Keyword (Mandatory)

Cloud technologies

Primary authors: SALOMONI, Davide (INFN CNAF); AIFTIMIEI, Doina Cristina (INFN - CNAF, IFIN-HH); MOLTÓ,

Germán (Universidad Politécnica de Valencia); DONVITO, Giacinto (INFN-Bari); BLANQUER, Ignacio (UPV); ANTONACCI, Marica (INFN Bari); SZEPIENIEC, Tomasz (Cyfronet)

Session Classification: Posters B / Break

Track Classification: Track 6: Infrastructures