

Abstracting application deployment on Cloud infrastructures

Tuesday 11 October 2016 16:30 (15 minutes)

Deploying a complex application on a Cloud-based infrastructure can be a challenging task. Among other things, the complexity can derive from software components the application relies on, from requirements coming from the use cases (i.e. high availability of the components, autoscaling, disaster recovery), from the skills of the users that have to run the application.

Using an orchestration service allows to hide the complex deployment of the application components, the order of each of them in the instantiation process and the relationships among them. In order to further simplify the application deployment to users not familiar with Cloud infrastructures and above layers, it can be worthwhile to provide an abstraction layer on top of the orchestration one.

In this contribution we present an approach for Cloud-based deployment of applications and its implementation in the framework of several projects, such as “!CHAOS: a cloud of controls, a project funded by MIUR (Italian Ministry of Research and Education) to create a Cloud-based deployment of a control system and data acquisition framework, “INDIGO-DataCloud”, an EC H2020 project targeting among other things high-level deployment of applications on hybrid Clouds, and “Open City Platform”, an Italian project aiming to provide open Cloud solutions for Italian Public Administrations.

Through orchestration services, we prototyped a dynamic, on-demand, scalable platform of software components, based on OpenStack infrastructures. A set of Heat templates developed ad-hoc allow to automatically deploy all the application components, minimize the faulty situations and guarantee the same configuration every time they run. The automatic orchestration is an example of Platform as a Service, that can be instantiated both via command-line or OpenStack dashboard, presuming a certain level of knowledge of OpenStack usage.

On top of the orchestration services we developed a prototype of a web interface exploiting the Heat APIs, that can be related to a specific application provided that ad-hoc Heat templates are available. The user can start an instance of the application without having knowledge about the underlying Cloud infrastructure and services. Moreover, the platform instance can be customized by choosing parameters related to the application such as the size of a File System or the number of instances of a NoSQL DB cluster. As soon as the desired platform is running, the web interface offers the possibility to scale some infrastructure components.

By providing this abstraction layer, users have a simplified access to Cloud resources and data center administrators can limit the degrees of freedom granted to the users.

In this contribution we describe the solution design and implementation, based on the application requirements, the details of the development of both the Heat templates and of the web interface, together with possible exploitation strategies of this work in Cloud data centers.

Tertiary Keyword (Optional)

Secondary Keyword (Optional)

Primary Keyword (Mandatory)

Cloud technologies

Authors: SALOMONI, Davide (Universita e INFN, Bologna (IT)); AIFTIMEI, Doina Cristina (INFN - CNAF, IFIN-HH); FATTIBENE, Enrico (INFN - National Institute for Nuclear Physics); Mr PANELLA, Matteo (INFN); Mr

GARGANA, Riccardo (INFN)

Session Classification: Posters A / Break

Track Classification: Track 6: Infrastructures