

Leveraging TOSCA orchestration to enable fully automated cloud-based research environments on federated heterogeneous e-infrastructures (Remote Presentation)

Wednesday, 22 March 2023 16:30 (30 minutes)

In the last years cloud computing has opened up interesting opportunities in many fields of scientific research. Cloud technologies allow to scale applications and adapt quickly, ease the adoption of new software development methods (e.g. DevOps), accelerating time to value.

However, the lack of integration of the existing infrastructures and the consequent fragmentation of the resources are still a barrier to a broader adoption of these technologies.

Starting from the times of the INDIGO-Datacloud project (2015-2017) we have been developing a set of solutions for implementing a seamless and transparent access to geographically distributed compute and storage resources, mainly the INDIGO IAM, a modern authentication and authorization system, and the INDIGO PaaS, a suite of microservices that allow to federate multiple providers and orchestrate cloud deployments via TOSCA.

At the beginning of 2021, INFN inaugurated a national multi-site cloud infrastructure (INFN Cloud), that is currently exploiting and extending the INDIGO solutions to provide an extensible portfolio of services tailored to multi-disciplinary scientific communities, spanning from traditional IaaS to more elaborate PaaS and SaaS solutions. Some examples are: data analytics and visualisation environments based on Elasticsearch and Kibana, file sync & share solution based on OwnCloud with replicated backend storage, web-based multi-user interactive development environment for notebooks, code and data built on JupyterLab, Kubernetes clusters, HTCondor on-demand clusters, Spark clusters integrated with Jupyter, Cloud storage solutions, etc. Moreover, the INFN Cloud service catalogue includes integration with the Kubernetes ecosystem and customizations for specific use-cases, e.g. the exploitation of GPUs for machine learning projects or pre-installed experiment software for data analysis.

The topology of each service is described through a TOSCA template, whereas the provisioning of the cloud resources is orchestrated through the INDIGO PaaS, that is able to schedule the request on the best provider of the federation; finally, the configuration of the resources is fully automated through Ansible roles. All these technical details are hidden to the final users that can request the instantiation of the services through a user-friendly web portal.

Security is another key aspect that is carefully addressed in our platform. First of all, we adopt consistent authentication and authorization rules defined at the different IaaS, PaaS and SaaS levels, providing user/group isolation. The recipes used for automating the installation are then developed by IT experts that take care also of implementing secure configurations and updating the recipes as soon as a vulnerability is discovered. Finally, the INDIGO PaaS system allows also to perform deployments on private networks (through a bastion properly configured by the provider) where the deployed services can be reached through dedicated VPNs.

In this contribution we will provide details about both the platform architecture, the high-level service implementation strategy and the expected lines of further development.

Primary authors: ANTONACCI, Marica (INFN); SALOMONI, Davide (INFN)

Co-authors: CIASCHINI, Vincenzo; COSTANTINI, Alessandro (INFN-CNAF); DONVITO, Giacinto (INFN); DUMA, Doina Cristina (INFN - CNAF); FANZAGO, Federica; GIORGIO, Emidio; ITALIANO, Alessandro; SGARAVATTO, Massimo; SPIGA, danielle (INFN-PG); SPINOSO, Vincenzo; STALIO, Stefano; VERLATO, Marco; TRACOLLI, Mirko

Presenters: ANTONACCI, Marica (INFN); SALOMONI, Davide (INFN)

Session Classification: VRE

Track Classification: Track 5: Virtual Research Environment (including tools, services, workflows, portals, ... etc.)