

Federation-registry: the renovated Configuration Management Database for dynamic cloud federation (Remote Presentation)

Tuesday, 26 March 2024 16:40 (20 minutes)

The **INDIGO PaaS orchestration system** is an open-source middleware designed to seamlessly federate heterogeneous computing environments, including public and private clouds, container platforms, and more. Its primary function lies in orchestrating the deployment of virtual infrastructures, ranging from simple to intricate setups. These virtual infrastructures can implement high-level services, such as Jupyter Hub, Kubernetes, Spark, and HTCondor clusters, providing users with convenient access and operational control.

At the heart of the orchestration system lies its core component, the Orchestrator, supported by a suite of micro-services. These micro-services play a crucial role in assisting the Orchestrator by facilitating the selection of the optimal provider from the federated environments, based on the specific deployment request.

Within this architecture, a pivotal micro-service is dedicated to implement the information system of the federation. This crucial component records comprehensive details about all the providers, encompassing their characteristics and capabilities. The information stored plays a central role in the matchmaking process between user deployment requests and available providers.

For instance, if a deployment request specifies the allocation of one or more GPUs, the Orchestrator relies on the information system to identify which providers within the federation, for which the user is entitled to allocate resources, offer GPU capabilities.

Currently, this functionality is implemented by the Configuration Management Database (CMDB) service, which stores and organizes information about resource providers, and the Service Level Agreement Manager (SLAM) which retains SLAs signed by users and resource provider administrators.

We have opted to replace the existing services due to the discontinuation of maintenance for the CMDB developed during the INDIGO-DataCloud project, which relies on outdated components. The forthcoming solution, the **Federation-Registry**, is a state-of-the-art web application built on the FastAPI framework. It features a REST API secured by OpenID-Connect/OAuth2 authentication and authorization technologies and policies. This upgrade ensures a more robust and secure foundation for managing federation-related information.

The Federation-Registry leverages *neo4j*, a highly flexible graph database, as opposed to the legacy *CouchDB* - a non-relational database - for storing and organizing data related to resource providers. Additionally, it adopts S3 object storage to securely store the signed SLA agreements. The implementation of a new population script becomes essential to retrieve information from the target resource providers and feed the database with relevant data.

This upgrade promises several advantages, including improved data organization, independence from outdated and unmaintained software, adherence to test-driven code practices, enhanced flexibility for accommodating various types of providers, and simplified database structure updates for the incorporation of new provider types. This contribution will outline the architectural decisions and delve into the specifics of the implementation.

The newly implemented Federation-registry service will be integrated into the **INFN Cloud platform**, which is already exploiting the INDIGO PaaS middleware to provide INFN scientific communities with a portfolio of high-level services supplied on-demand across geographically distributed cloud sites.

Primary authors: SAVARESE, Giovanni (INFN); ANTONACCI, Marica (INFN)

Co-author: GIOMMI, Luca (INFN CNAF)

Presenter: SAVARESE, Giovanni (INFN)

Session Classification: Infrastructure Clouds & Virtualisation

Track Classification: Track 8: Infrastructure Clouds and Virtualizations