

Efficient management of INDIGO-IAM clients and S3 buckets via INDIGO PaaS Orchestrator in INFN Cloud (Remote Presentation)

Tuesday, March 26, 2024 4:00 PM (20 minutes)

The National Institute for Nuclear Physics (INFN) has been operating and supporting Italy's largest research and academic distributed infrastructure for several decades. In March 2021, INFN launched "INFN Cloud" which provides a federated cloud infrastructure and a customizable service portfolio tailored for the scientific communities supported by the institute. The portfolio assortment comprises standard IaaS options and more advanced PaaS and SaaS solutions, all crafted to suit the distinct needs of specific communities. All PaaS services in the portfolio are described through an Infrastructure as Code paradigm based on a declarative approach, via a combination of TOSCA templates (to model an application stack), Ansible roles (to manage the automated configuration of virtual environments), and Docker containers (to encapsulate high-level application software and runtime).

The federation middleware of the INFN Cloud platform is built upon the INDIGO PaaS Orchestration system, which consists of interconnected open-source microservices. Among them, there is the INDIGO PaaS Orchestrator which receives high-level deployment requests from users and coordinates the deployment process over IaaS platforms.

In this work, we address an issue within INFN Cloud concerning the proliferation of INDIGO Identity Access Management (INDIGO-IAM) clients and S3 buckets. Specifically, these resources are created during the on-demand deployment of high-level services, such as PaaS and SaaS offerings like Jupyter Hub or File Sync&Share solutions like ownCloud or NextCloud. These services require the creation of an INDIGO-IAM client to authenticate users and authorize them to access the service. Currently, this process is performed through Ansible recipes, with no transmission of client information to the INDIGO PaaS Orchestrator. This results in a scenario where, upon deployment deletion or failure, the related INDIGO-IAM client is not removed. This leads to an increasing number of unused clients, consequently causing a decrease in the performance of the INDIGO-IAM Service (an issue that has already been confirmed to exist).

Our proposed resolution involves delegating to the INDIGO PaaS Orchestrator the creation (and the deletion) of any INDIGO-IAM clients. The implemented feature offers users enhanced flexibility, enabling them to create multiple clients, select the identity provider, define scopes, and assign the client owner so that its configuration can be managed even later.

A similar problem was identified concerning the proliferation of S3 buckets. Some services that integrate S3 storage as a backend (such as Sync&Share as a Service), usually write a lot of data into these buckets. Also in this case, when the deletion of the deployment is triggered, the associated buckets (and their contents) persist, resulting in a continuous increase in disk space consumption. To address this, we found it convenient to delegate to the INDIGO PaaS Orchestrator both the creation and the deletion of the buckets.

This work presents the solution we developed for both challenges, along with a comprehensive overview of the new functionalities introduced in the code.

Primary authors: Dr GIOMMI, Luca (INFN CNAF); ANTONACCI, Marica (INFN)

Co-authors: VIANELLO, Enrico (INFN); Ms MICCOLI, Roberta (INFN); AGOSTINI, Federica (INFN-CNAF); FORNARI, Federico (INFN-CNAF); Mr VINO, Gioacchino (INFN); Mr SAVARESE, Giovanni (INFN); DONVITO, Giacinto (INFN); COSTANTINI, Alessandro (INFN CNAF)

Presenter: Dr GIOMMI, Luca (INFN CNAF)

Session Classification: Infrastructure Clouds & Virtualisation

Track Classification: Track 8: Infrastructure Clouds and Virtualizations