Finanziato dall'Unione europea NextGenerationEU | Ministero dell'Università e della Ricerca | Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA | terabit

The INFN Cloud platform: state of the art and plan for a renewed PaaS Orchestration system

L. Giommi and G. Savarese

On behalf of many people, not only from WP5

Workshop sul Calcolo nell'INFN – La Biodola | 26-30 Maggio 2025

The INFN Cloud ecosystem

INFN decided to implement a **national Cloud computing infrastructure** for research

- as a **federation** of existing distributed Cloud infrastructures
- as an “user-centric” infrastructure which makes available to the final users a dynamic **set of services** tailored on specific use cases
- leveraging the outcomes of several national and European Cloud projects where INFN actively participated, e.g. INDIGO DataCloud

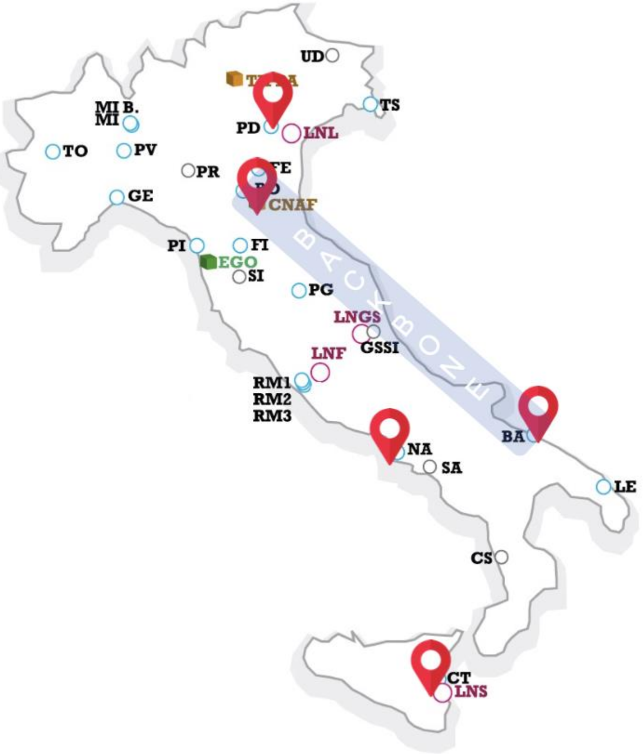

INFN Cloud was officially made available to users in **March 2021**

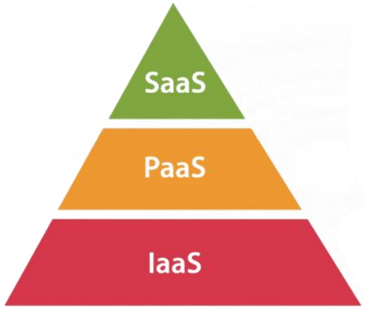

e.g. Notebook as a Service

e.g. Virtual Machine, Docker compose

e.g. Start & Stop, Hostname choice

The Infrastructure as Code paradigm

All PaaS services are defined using an **Infrastructure as Code** paradigm, based on a procedural paradigm that aims to reduce manual processes and increase flexibility and portability across environments, via a combination of:

- **TOSCA** (**T**opology and **O**rchestration **S**pecification for **C**loud **A**pplications) templates, to model an application stack
- **Ansible** roles, to manage the automated configuration of virtual environments
- **Docker** containers, to encapsulate high-level application software and runtime
- **Helm** charts, to manage the deployment of an application in Kubernetes clusters

Finanziato dall'Unione europea NextGenerationEU | Ministero dell'Università e della Ricerca | Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA | terabit

The current INDIGO PaaS Orchestration system of INFN Cloud

- The federative middleware of INFN Cloud is based on the **INDIGO PaaS Orchestration system**, consisting of interconnected open-source microservices
- The Orchestrator receives high-level deployment requests in the form of TOSCA templates and coordinates the deployment process by using the **Infrastructure Manager (IM)** to interact with provider services and deploy complex, customized virtual infrastructures on the IaaS platforms offered by the federated providers
- A central activity was the introduction of the **Federation Registry** and **Feeder**, which replaced old components and were integrated with the entire PaaS Orchestration system

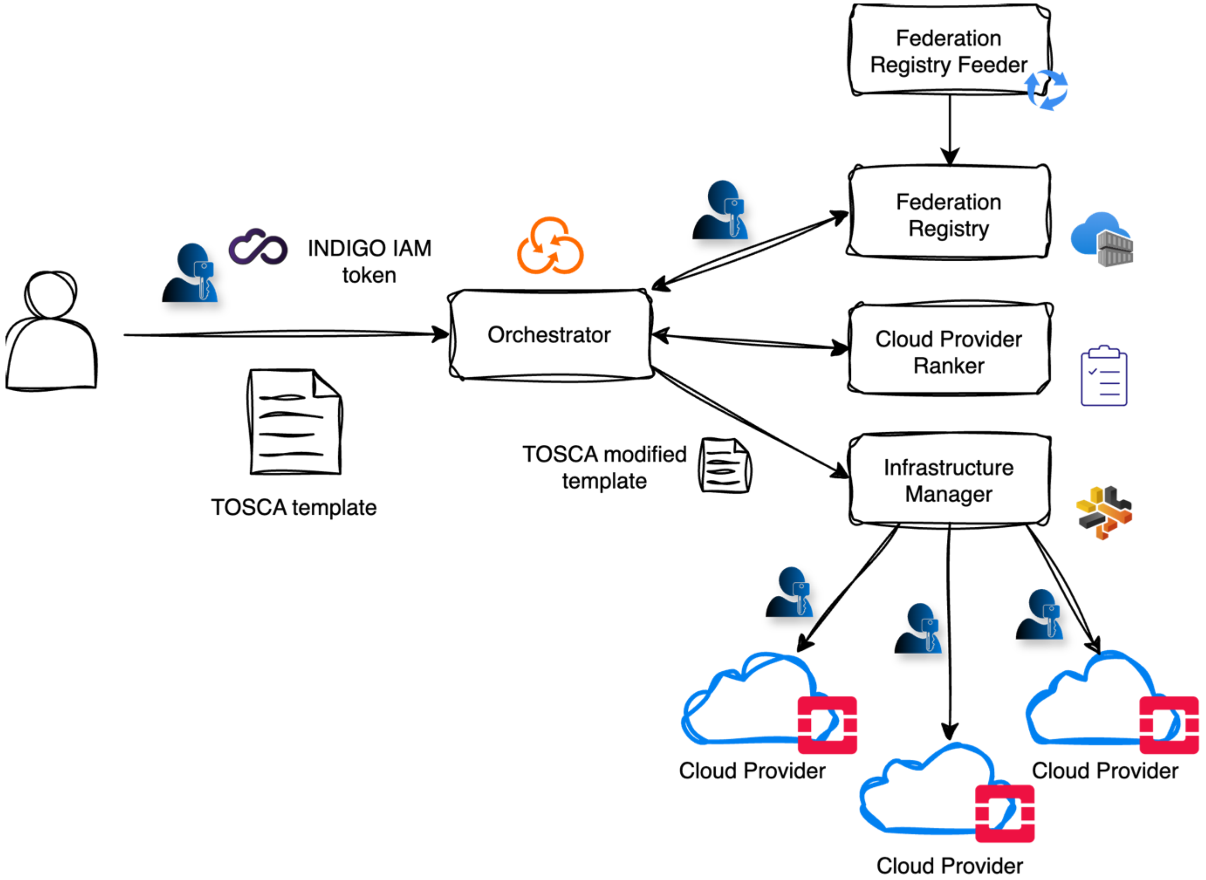

Introduction of the Federation Registry

Federation Registry Feeder

- Periodic Python script
- Based on YAML configuration file to connect to federated providers
- Update the Federation Registry with up to date information (flavors, images, networks, quotas and more) retrieved directly from federated providers

Federation Registry

- Python REST API based on FastAPI
- Support for OAuth2/OIDC authentication and authorization
- Uses Neo4j as graph database

Operations

- Jenkins pipelines to: build and push docker images and test code
- Containerized services deployed through ansible role and playbooks
- Service replication between CNAF and BARI

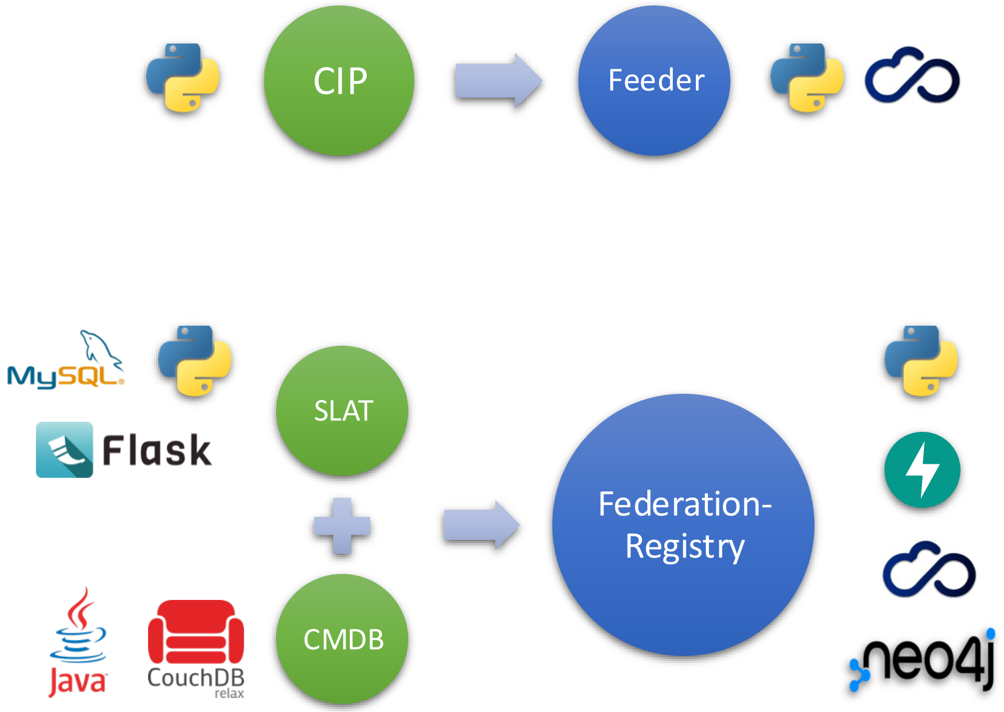

The PaaS Orchestrator Dashboard

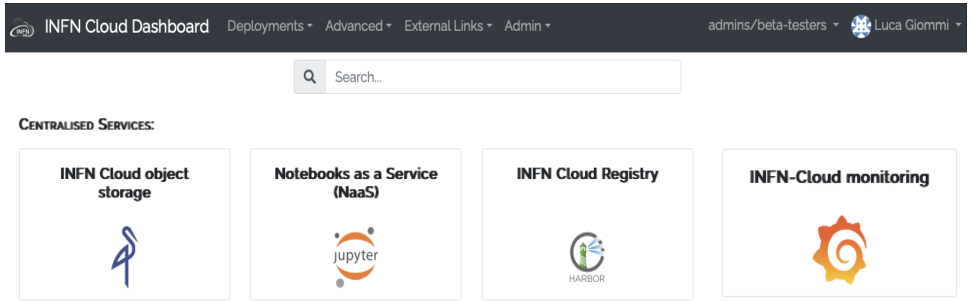

CENTRALISED SERVICES:

INFN Cloud object storage

Notebooks as a Service (NaaS)

INFN Cloud Registry

INFN-Cloud monitoring

Old style

New style

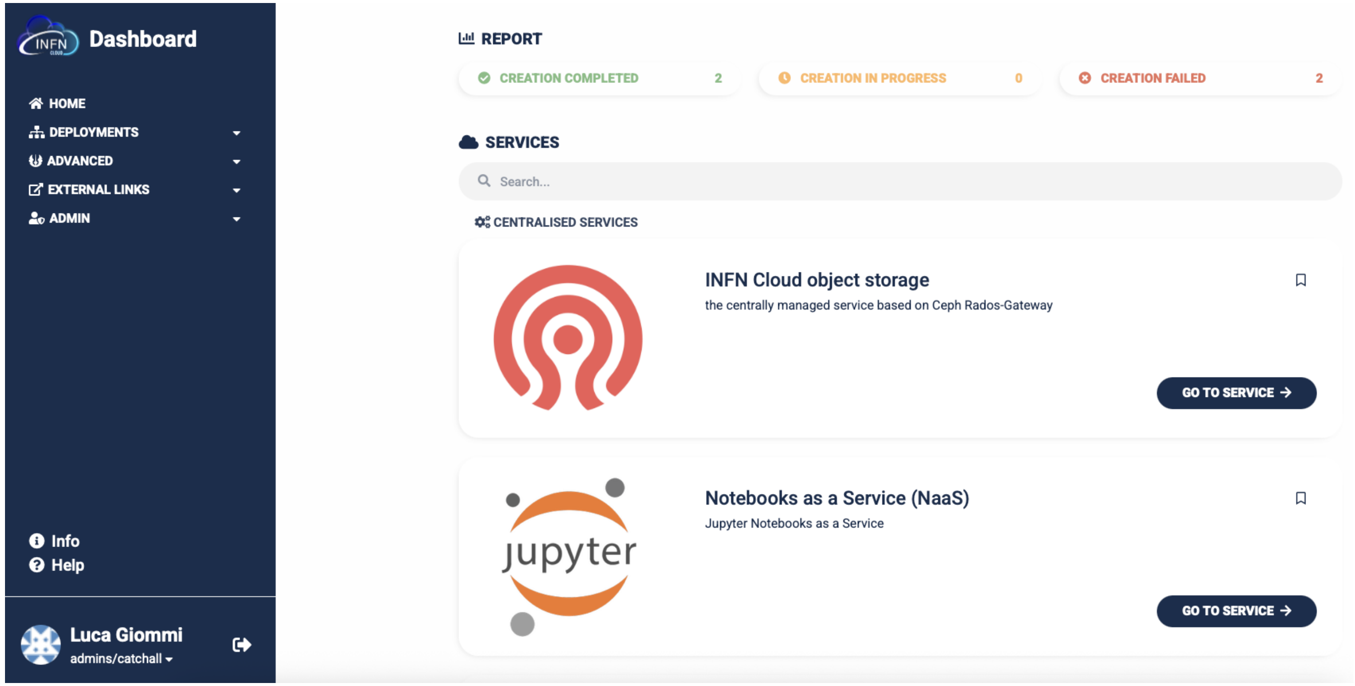

REPORT

CREATION COMPLETED 2

CREATION IN PROGRESS 0

CREATION FAILED 2

SERVICES

CENTRALISED SERVICES

INFN Cloud object storage
the centrally managed service based on Ceph Rados-Gateway

Notebooks as a Service (NaaS)
Jupyter Notebooks as a Service

https://my.cloud.infn.it

Updates on the PaaS Orchestrator Dashboard: Admins

- **Admins** can
  - Manage deployments of other users: deletion of deployments and full logs visualization

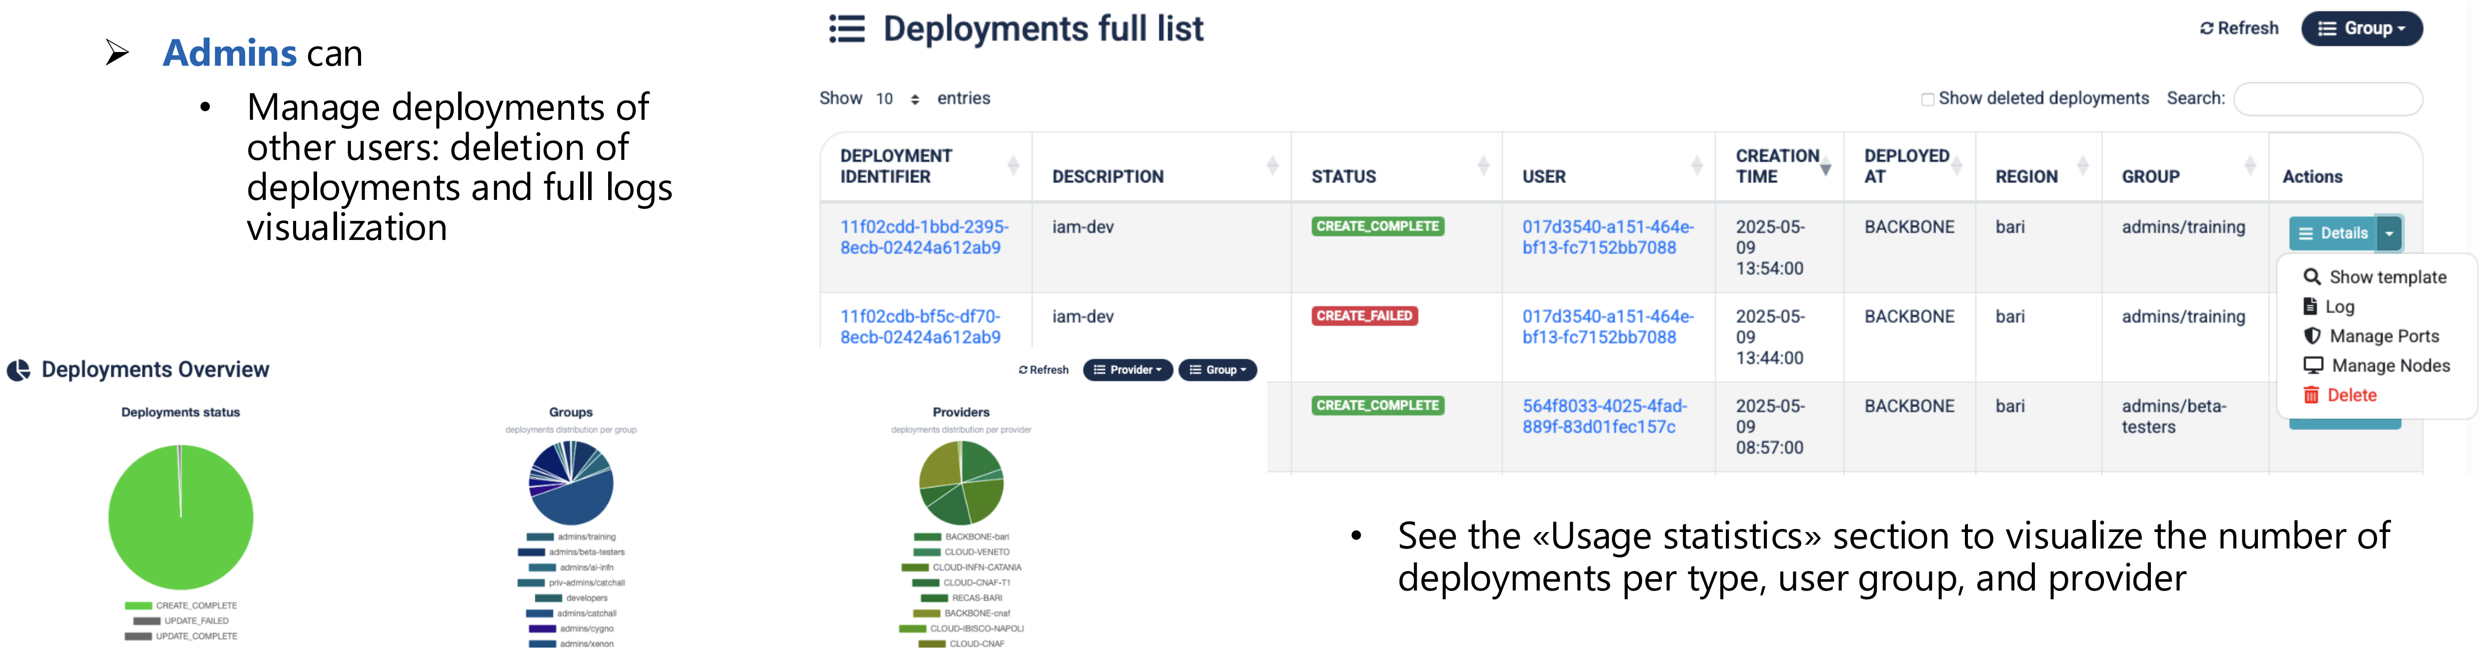

Deployments full list

Show 10 ⇕ entries

□ Show deleted deployments Search:

DEPLOYMENT IDENTIFIER	DESCRIPTION	STATUS	USER	CREATION TIME	DEPLOYED AT	REGION	GROUP	Actions
11f02cdd-1bbd-2395-8ecb-02424a612ab9	iam-dev	CREATE_COMPLETE	017d3540-a151-464e-bf13-fc7152bb7088	2025-05-09 13:54:00	BACKBONE	bari	admins/training	☰ Details ▾
11f02cdb-bf5c-df70-8ecb-02424a612ab9	iam-dev	CREATE_FAILED	017d3540-a151-464e-bf13-fc7152bb7088	2025-05-09 13:44:00	BACKBONE	bari	admins/training	
		CREATE_COMPLETE	564f8033-4025-4fad-889f-83d01fec157c	2025-05-09 08:57:00	BACKBONE	bari	admins/beta-testers	

Deployments Overview

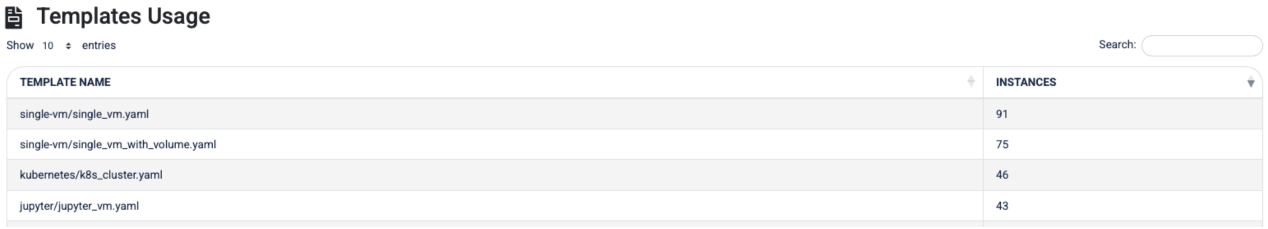

Templates Usage

Show 10 ⇕ entries

Search:

TEMPLATE NAME	INSTANCES
single-vm/single_vm.yaml	91
single-vm/single_vm_with_volume.yaml	75
kubernetes/k8s_cluster.yaml	46
jupyter/jupyter_vm.yaml	43

- See the «Usage statistics» section to visualize the number of deployments per type, user group, and provider

Updates on the PaaS Orchestration system

➢ News for the **users**

- added **icons** to the buttons in the sidebar, along with a "Home" button that redirects to the dashboard homepage

- during the creation of a new deployment, the **scheduling selection** is now presented first: this allows users to pre-select the cloud site where they want the deployment to be created. In the following pages, only the specific flavours and operating systems available on the selected site will be shown

- if the deletion of a deployment fails, a "**Delete (force)**" option will be available in the action list, allowing forced deletion of the deployment

- in case of a deployment creation failure, a "**Retry**" option will be available in the drop-down menu, enabling users to resubmit the deployment request with the same parameters

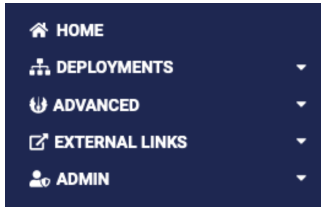

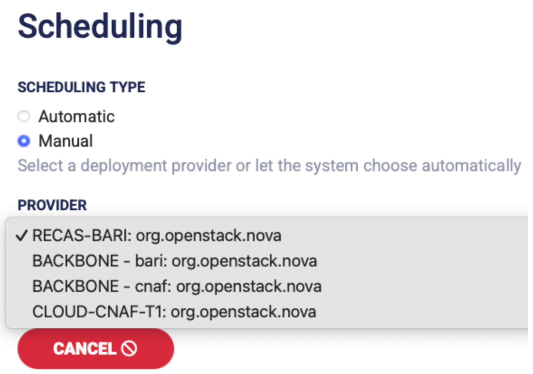
Scheduling

SCHEDULING TYPE

○ Automatic
● Manual

Select a deployment provider or let the system choose automatically

PROVIDER

✓ RECAS-BARI: org.openstack.nova
BACKBONE - bari: org.openstack.nova
BACKBONE - cnaf: org.openstack.nova
CLOUD-CNAF-T1: org.openstack.nova

CANCEL

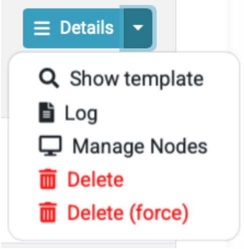

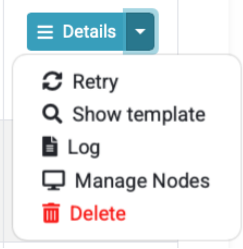

Updates on the PaaS and SaaS services

- **PaaS services**
  - Changed the default OS to Debian 12
  - Updated and improved the TOSCA template of the «Kubernetes cluster» service
  - Added a new variant of the PaaS service «Kubernetes cluster» with an InterLink Virtual Node

- **SaaS services**
  - Restored and updated the Jupyter Notebook as a Service (NaaS) service
  - Added a new service based on Healthchecks.io
  - Developed and in production a Web App used as a GUI for the Object Storage Service of INFN Cloud

The PaaS new architecture

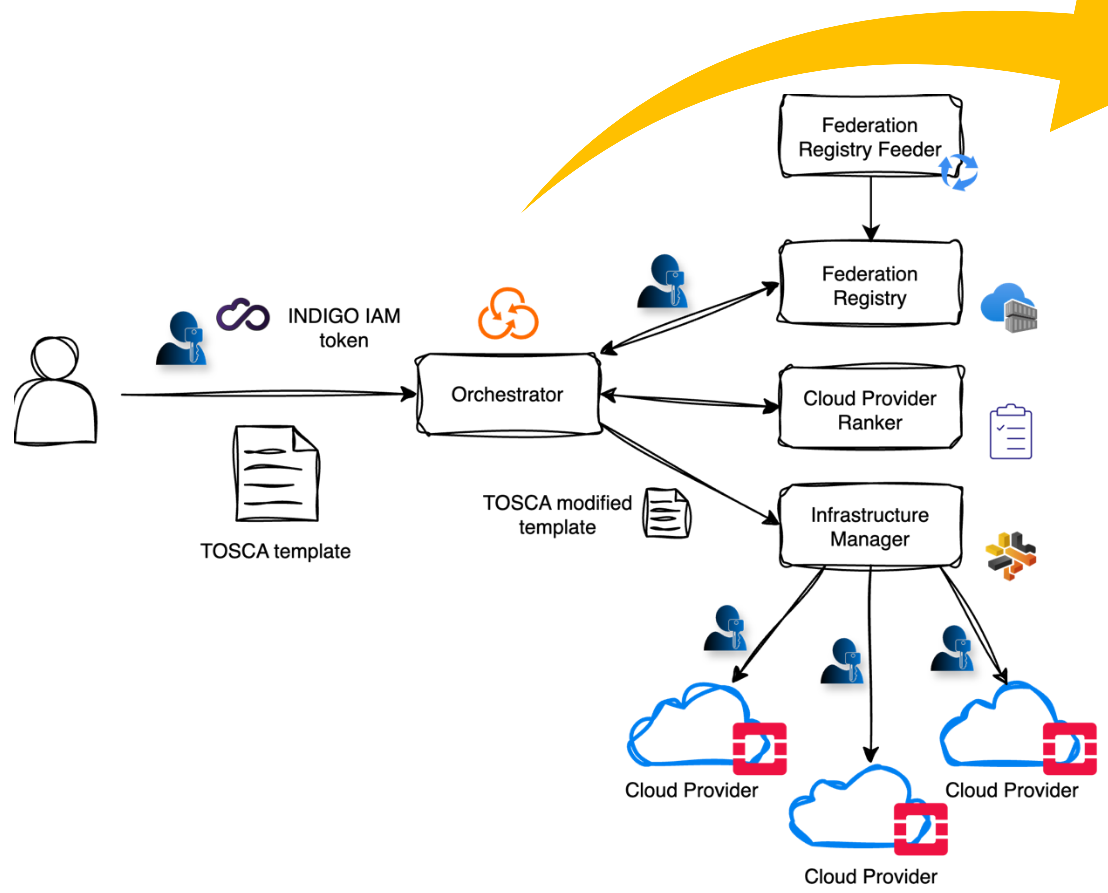

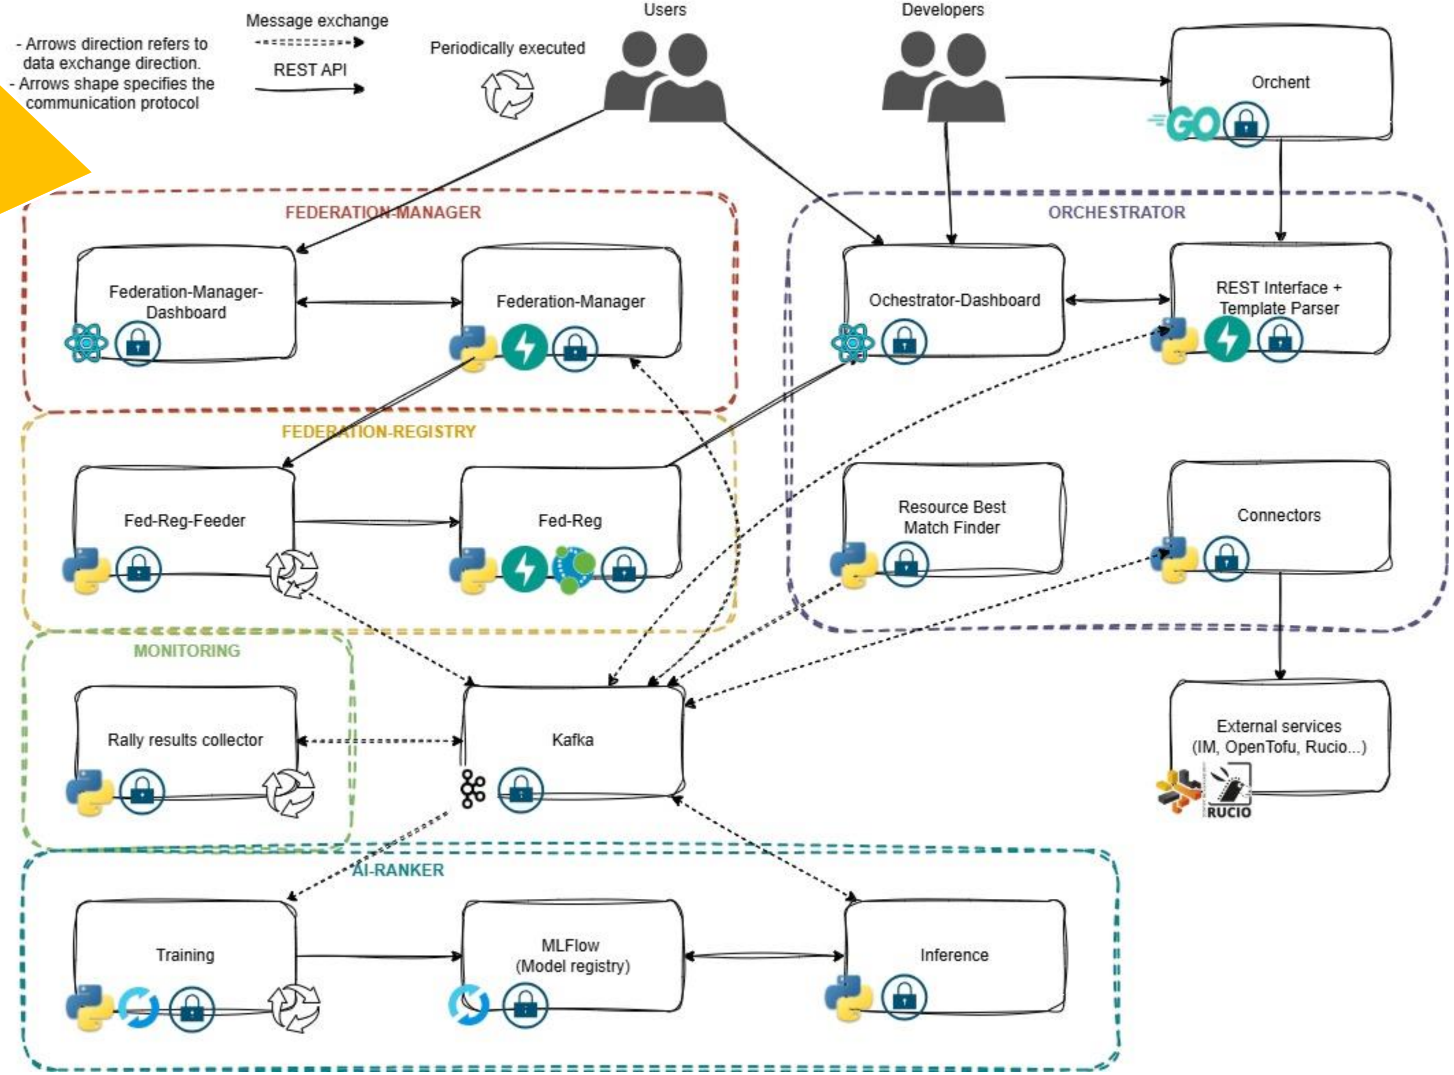

The monitoring and AI-Ranker service

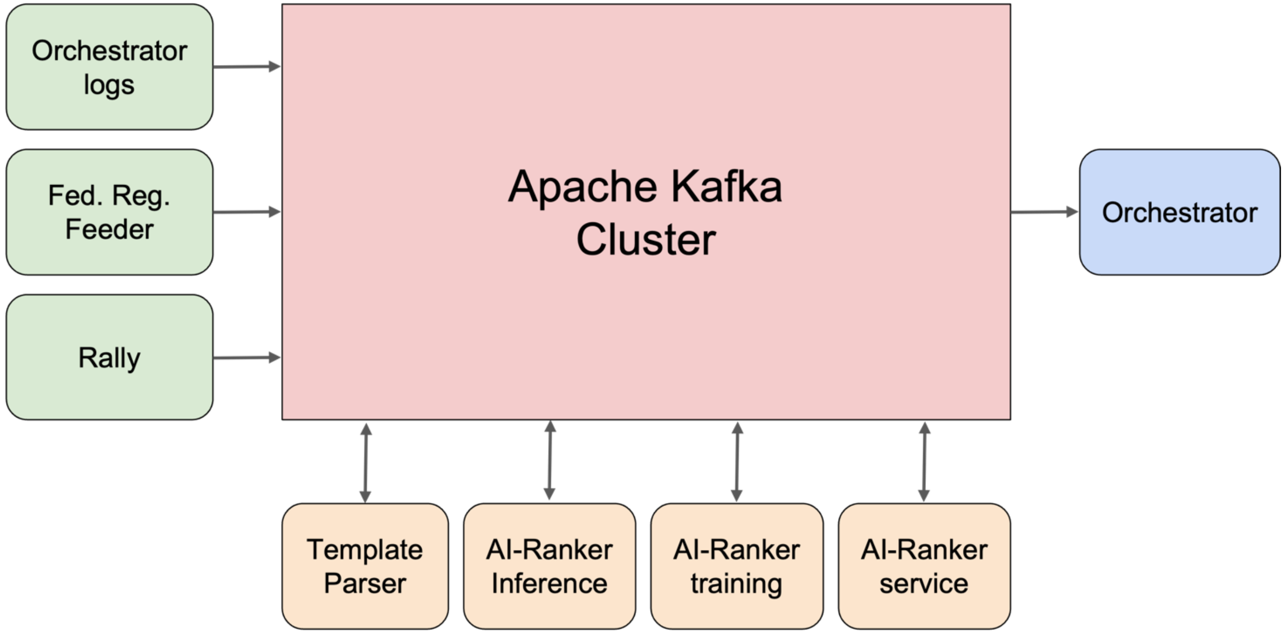

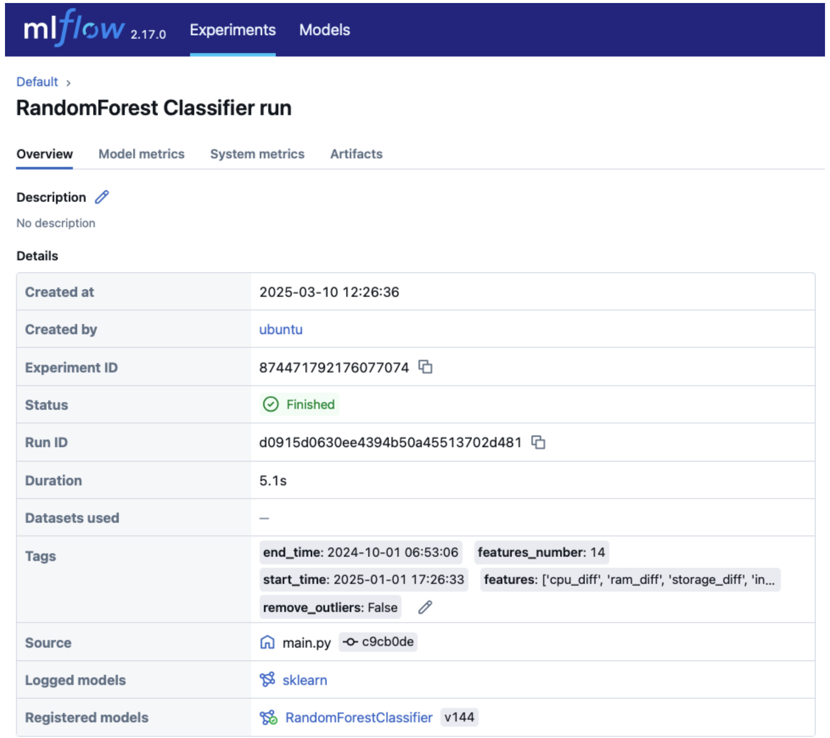

RandomForest Classifier run

Description

No description

Details

Created at	2025-03-10 12:26:36
Created by	ubuntu
Experiment ID	874471792176077074
Status	Finished
Run ID	d0915d0630ee4394b50a45513702d481
Duration	5.1s
Datasets used	—
Tags	**end_time**: 2024-10-01 06:53:06 **features_number**: 14 **start_time**: 2025-01-01 17:26:33 **features**: ['cpu_diff', 'ram_diff', 'storage_diff', 'in... **remove_outliers**: False
Source	main.py c9cb0de
Logged models	sklearn
Registered models	RandomForestClassifier v144

See the poster Improving the Cloud Provider Ranking in the INDIGO PaaS Orchestration System with AI Techniques

Federation-Manager

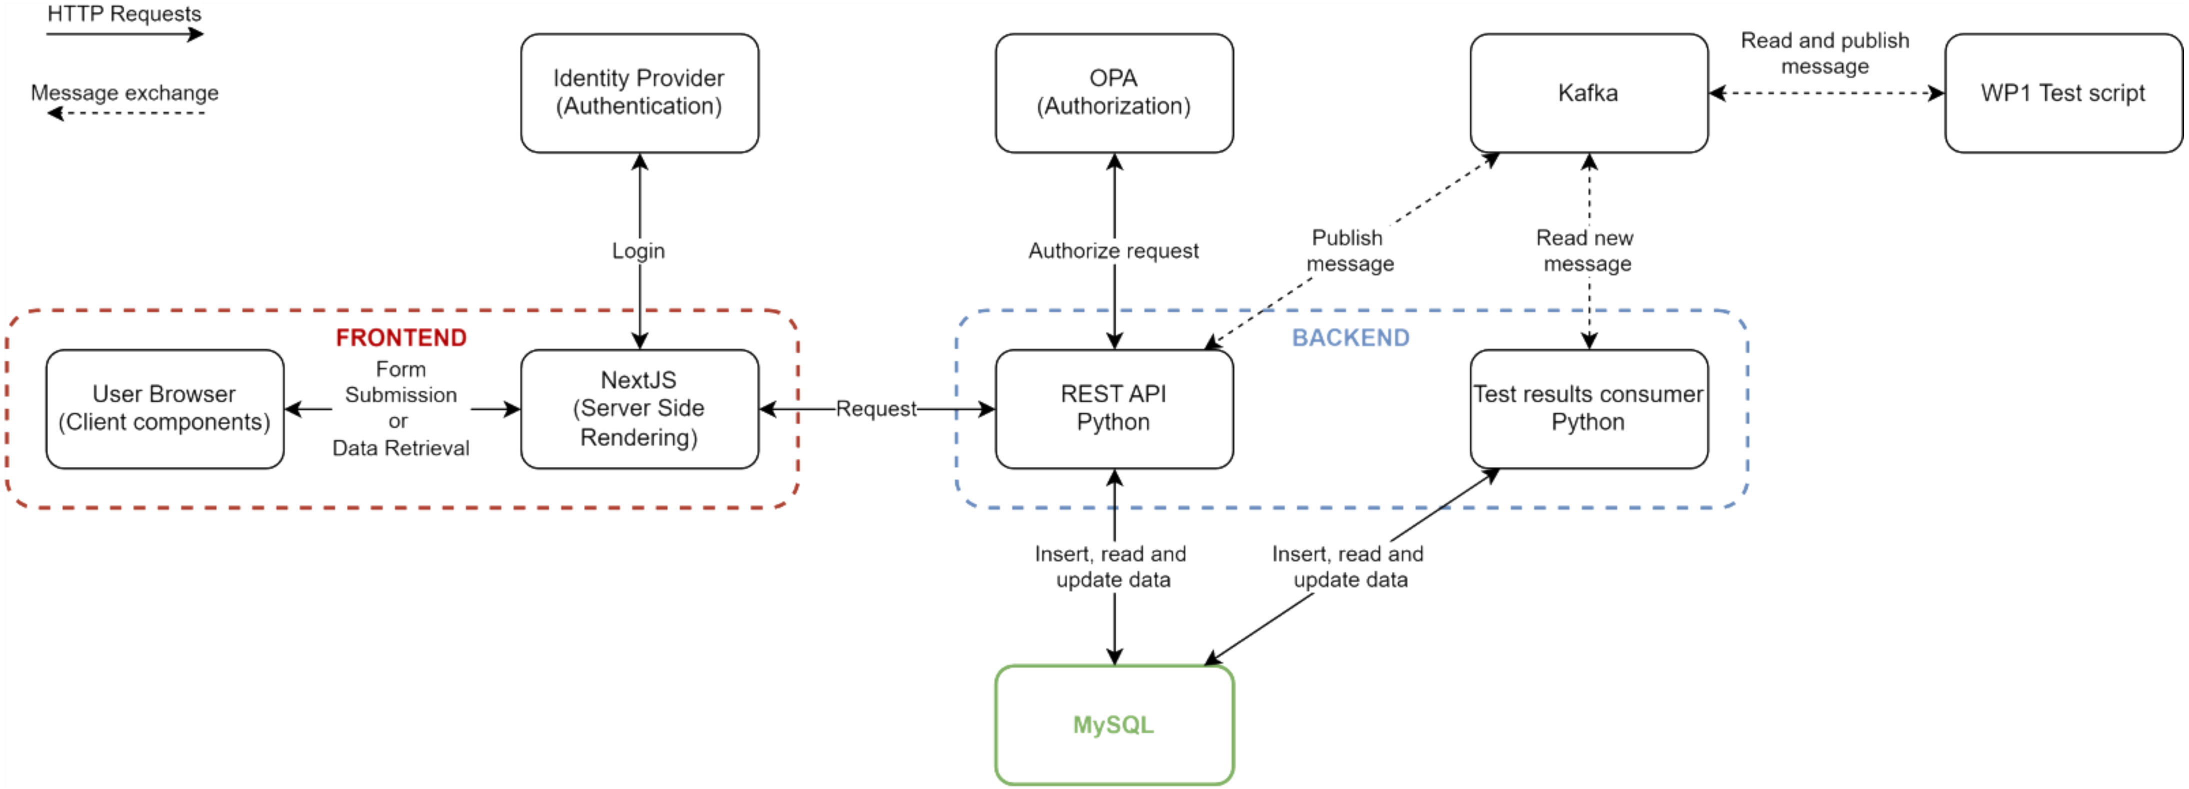
HTTP Requests
Message exchange
Identity Provider (Authentication)
OPA (Authorization)
Kafka
Read and publish message
WP1 Test script
Login
Authorize request
Publish message
Read new message
FRONTEND
BACKEND
User Browser (Client components)
Form Submission or Data Retrieval
NextJS (Server Side Rendering)
Request
REST API Python
Test results consumer Python
Insert, read and update data
Insert, read and update data
MySQL

See the poster Federation-Manager: Un nuovo strumento per l'integrazione di nuovi provider di risorse e comunità nel progetto DataCloud

Federation-Manager

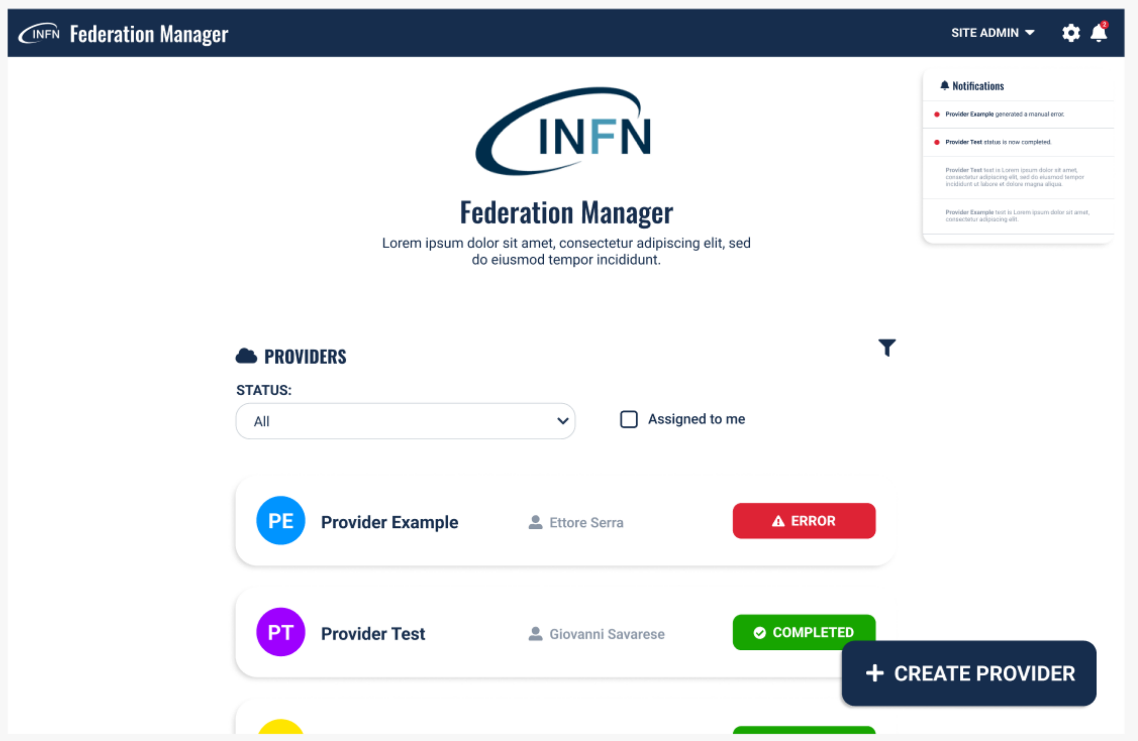

INFN Federation Manager
SITE ADMIN ▾

Notifications

- ● **Provider Example** generated a manual error.
- ● **Provider Test** status is now completed.
- **Provider Test** text is Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt ut labore et dolore magna aliqua.
- **Provider Example** text is Lorem ipsum dolor sit amet, consectetur adipiscing elit.

INFN

Federation Manager

Lorem ipsum dolor sit amet, consectetur adipiscing elit, sed do eiusmod tempor incididunt.

PROVIDERS

STATUS: All ▾   □ Assigned to me

PE	Provider Example	Ettore Serra	ERROR
PT	Provider Test	Giovanni Savarese	COMPLETED

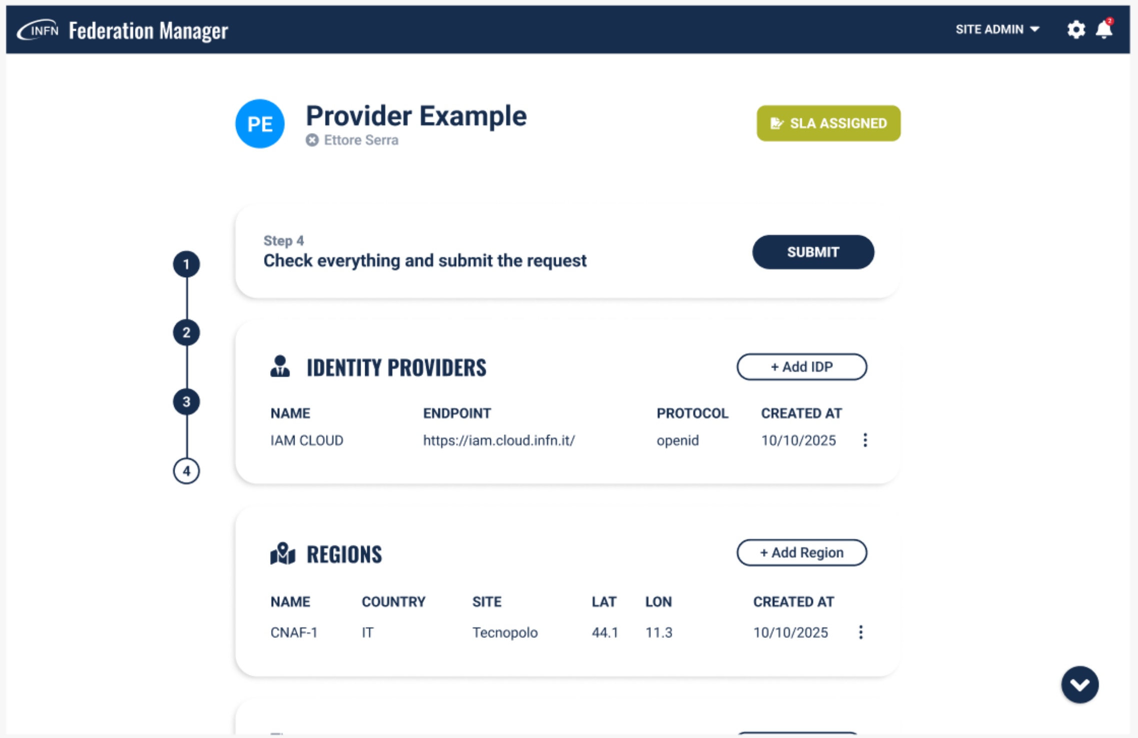

INFN Federation Manager
SITE ADMIN ▾

PE Provider Example
Ettore Serra

Step 4
Check everything and submit the request

IDENTITY PROVIDERS

NAME	ENDPOINT	PROTOCOL	CREATED AT
IAM CLOUD	https://iam.cloud.infn.it/	openid	10/10/2025

REGIONS

NAME	COUNTRY	SITE	LAT	LON	CREATED AT
CNAF-1	IT	Tecnopolo	44.1	11.3	10/10/2025

See the poster Federation-Manager: Un nuovo strumento per l’integrazione di nuovi provider di risorse e comunità nel progetto DataCloud

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
terabit

The new Orchestrator component

- **REST API (python)**
  - Based on FastAPI
  - Authentication and authorization through OAuth2/OIDC and OPA
  - TOSCA template validation
  - Produce messages to Kafka to start automatic procedures
  - Retrocompatibily (as much as possible) with orchent and the current dashboard
- **Request pre-processing (python)**
  - Retrieve from kafka, users available providers (with their configuration), providers status, users requests
  - Upload the message used by the AI-Ranker to rank providers
- **Dispatchers and other compontents (python)**
  - **IM Connector** forward the user request to create/update/delete deployments to the IM
    - Support both Openstack and Kubernetes deployments
  - **Opentofu Connector** create/update/delete deployments (to be studied and evaluated)
  - **Rucio Connector** for data management (replica, migration and more)
- **Dashboard (typescript)**
  - Rewrite the dashboard component with the technologies used by RGW Web App and IAM

New DevOps strategies

- **What has been done**
  - Standardization of CI/CD via Jenkins pipelines
  - Update ansible playbooks and roles
  - Reviewed access permission and security rules on pre-production VM instances
  - New VPN based on OIDC authentication to access the pre-production VM instances
- **Future upgrades**
  - Migrate services from docker containers to k8s pods
  - Exploit ArgoCD to automate deployment to operations procedures

Thank you

luca.giommi@cnaf.infn.it,

giovanni.savarese@ba.infn.it

New PaaS service: Kubernetes cluster with an InterLink Virtual Node

- It uses the interLink plugin developed within the interTwin project
  - **interTwin**: project funded by the EU for the development of an open-source platform, called Digital Twin Engine (DTE), to handle "digital twins" of selected scientific communities
  - **interLink**: allows transparent offloading of Kubernetes workloads to remote computation systems

- Workload offloading
  - specify requirements, e.g. the number of GPUs
  - resources may not be available on the local cluster
  - workload can be opportunistically offloaded to a remote cluster where resources are available
- InterLink main components
  - **Virtual Node**: translate requests for a Kubernetes POD execution into a remote call to the InterLink API server.
  - **InterLink API Server**: a pluggable REST server that talks to the remote cluster

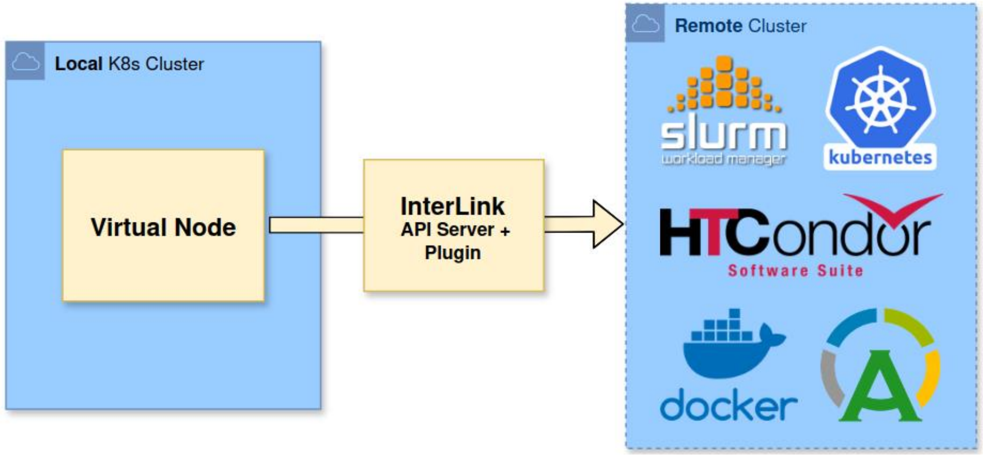

Renovation of the S3 Web App for the Object Storage Service

- The object storage service of INFN Cloud is now based on **CEPH / Rados Gateway** (previously MinIO)
  - Uses Open Policy Agent (OPA) for fine-grained Authorization
  - Two zones (CNAF and Bari) with independent Ceph Storage Clusters
  - Three instances of RADOS Gateway run in high availability within each zone

- Developed a new version (v1.0.0) of the **Web App** used as a GUI for the Object Storage Service
  - Built on Next.js and React.js
  - OIDC protocol with IAM to generate Json Web Token (JWT)
  - Uses IAM Access Token to perform STS with RGW
  - S3 operations using AWS SDK library
  - New graphical UI to be consistent with other INFN web applications

https://s3webui.cloud.infn.it

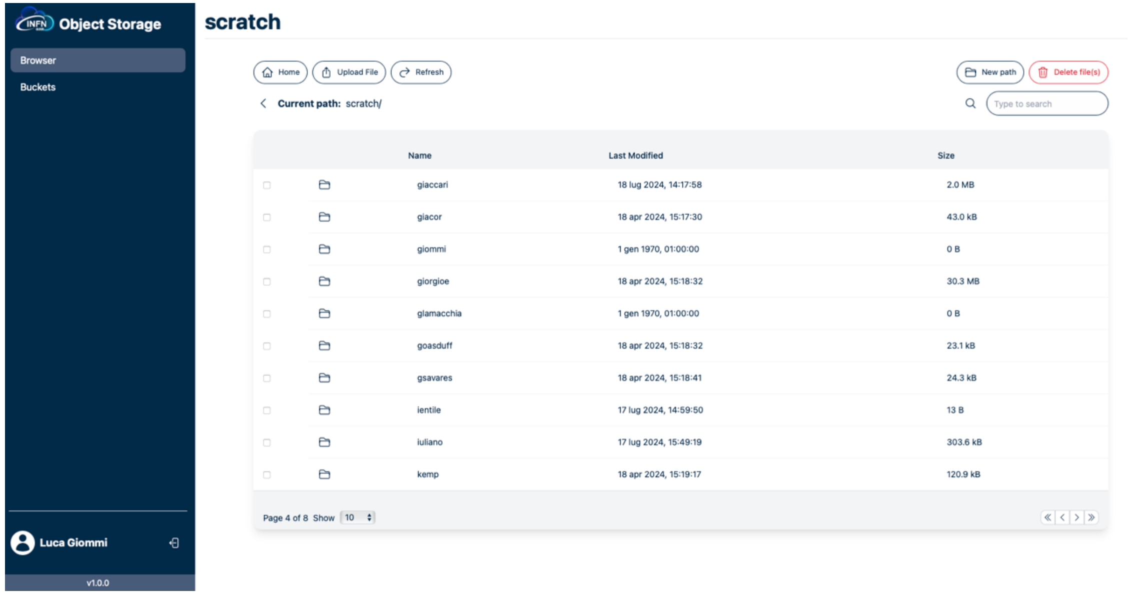

scratch

Current path: scratch/

	Name	Last Modified	Size
□	giaccari	18 lug 2024, 14:17:58	2.0 MB
□	giacor	18 apr 2024, 15:17:30	43.0 kB
□	giommi	1 gen 1970, 01:00:00	0 B
□	giorgioe	18 apr 2024, 15:18:32	30.3 MB
□	glamacchia	1 gen 1970, 01:00:00	0 B
□	goasduff	18 apr 2024, 15:18:32	23.1 kB
□	gsavares	18 apr 2024, 15:18:41	24.3 kB
□	ientile	17 lug 2024, 14:59:50	13 B
□	iuliano	17 lug 2024, 15:49:19	303.6 kB
□	kemp	18 apr 2024, 15:19:17	120.9 kB

Page 4 of 8 Show 10

Updates on the PaaS Orchestrator Dashboard

- **Admins** can:
  - Manage deployments of other users: deletion of deployments and full logs visualization

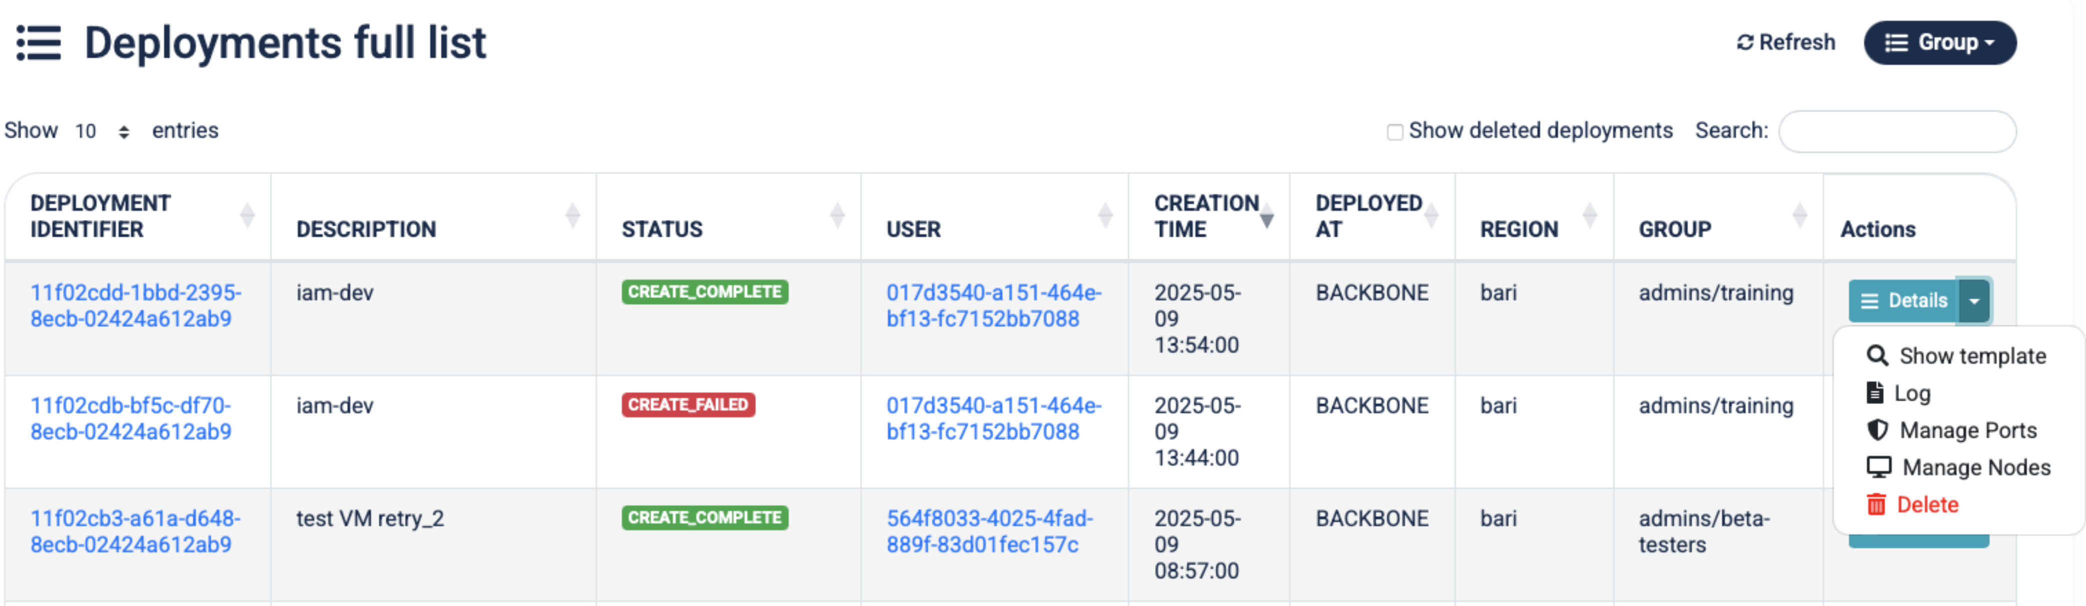

Deployments full list

DEPLOYMENT IDENTIFIER	DESCRIPTION	STATUS	USER	CREATION TIME	DEPLOYED AT	REGION	GROUP	Actions
11f02cdd-1bbd-2395-8ecb-02424a612ab9	iam-dev	CREATE_COMPLETE	017d3540-a151-464e-bf13-fc7152bb7088	2025-05-09 13:54:00	BACKBONE	bari	admins/training	Details
11f02cdb-bf5c-df70-8ecb-02424a612ab9	iam-dev	CREATE_FAILED	017d3540-a151-464e-bf13-fc7152bb7088	2025-05-09 13:44:00	BACKBONE	bari	admins/training	
11f02cb3-a61a-d648-8ecb-02424a612ab9	test VM retry_2	CREATE_COMPLETE	564f8033-4025-4fad-889f-83d01fec157c	2025-05-09 08:57:00	BACKBONE	bari	admins/beta-testers	

Updates on the PaaS Orchestrator Dashboard

- **Admins** can:
  - See the «Usage statistics» section to visualize the number of deployments per type, user group, and provider

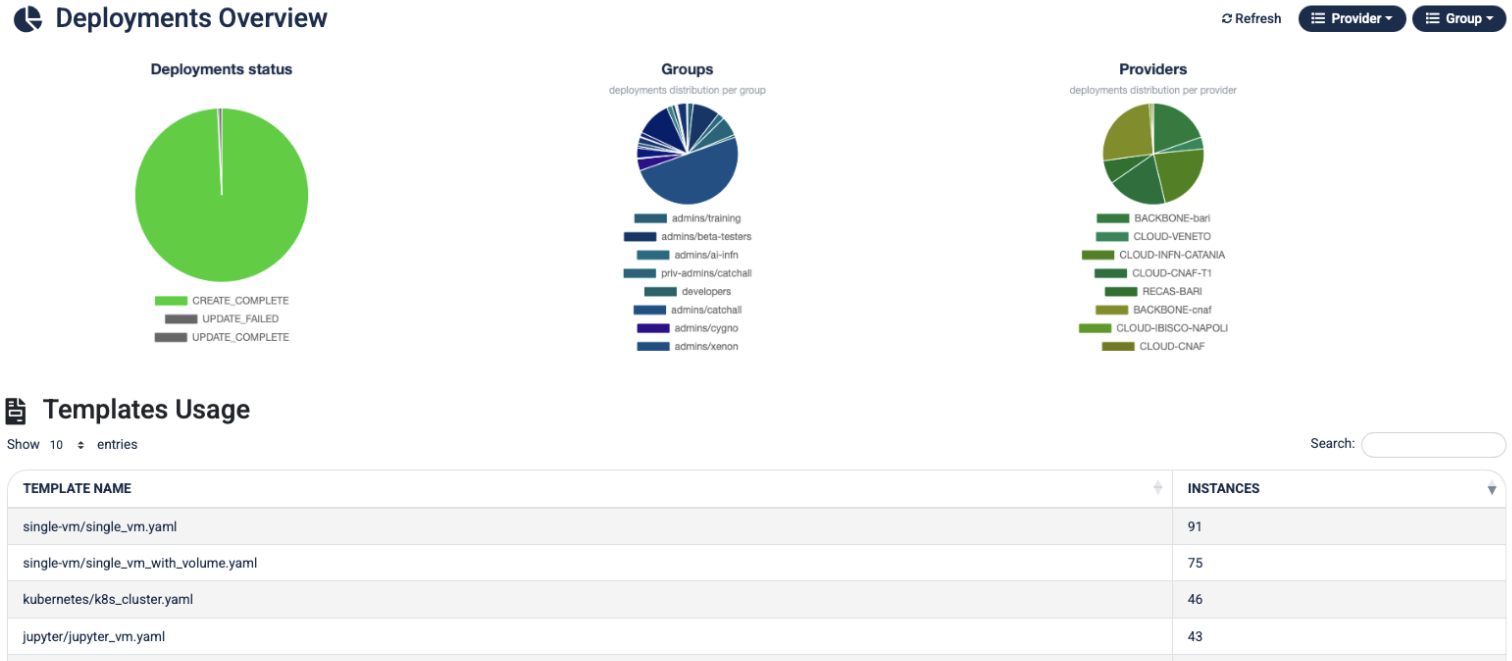

Templates Usage

Show 10 entries

Search:

TEMPLATE NAME	INSTANCES
single-vm/single_vm.yaml	91
single-vm/single_vm_with_volume.yaml	75
kubernetes/k8s_cluster.yaml	46
jupyter/jupyter_vm.yaml	43