



Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



# The renewed PaaS Orchestration system architecture

Giovanni Savarese, INFN-BA

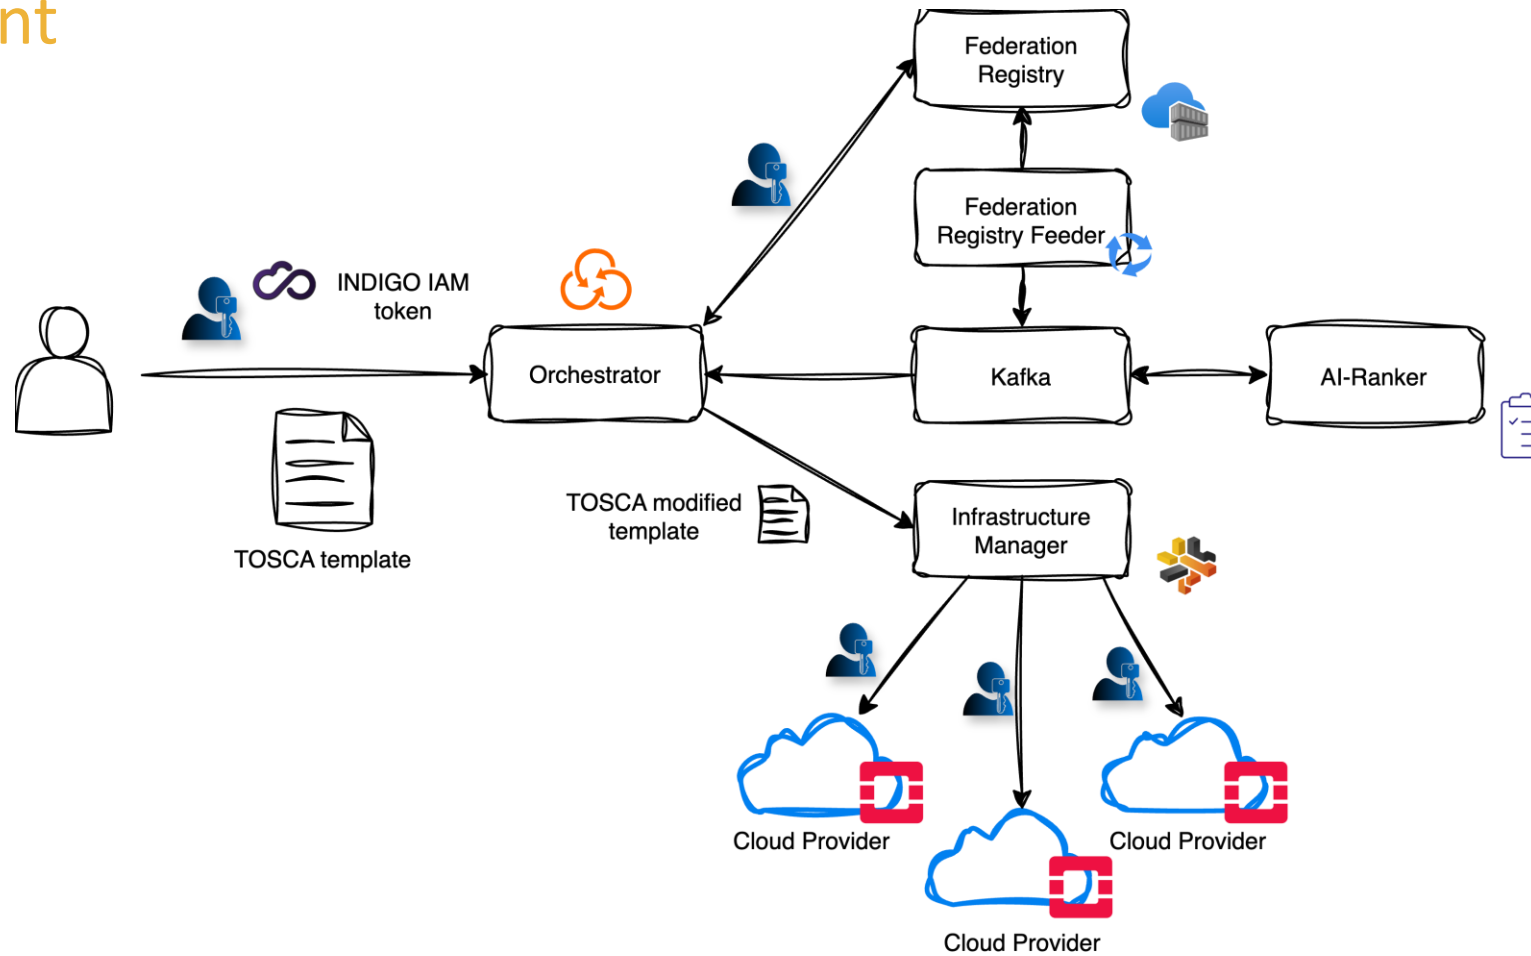
Luca Giommi, INFN-CNAF

Bari, 08.09.2025

Sviluppi PaaS, Bari, 08-10.09.2025



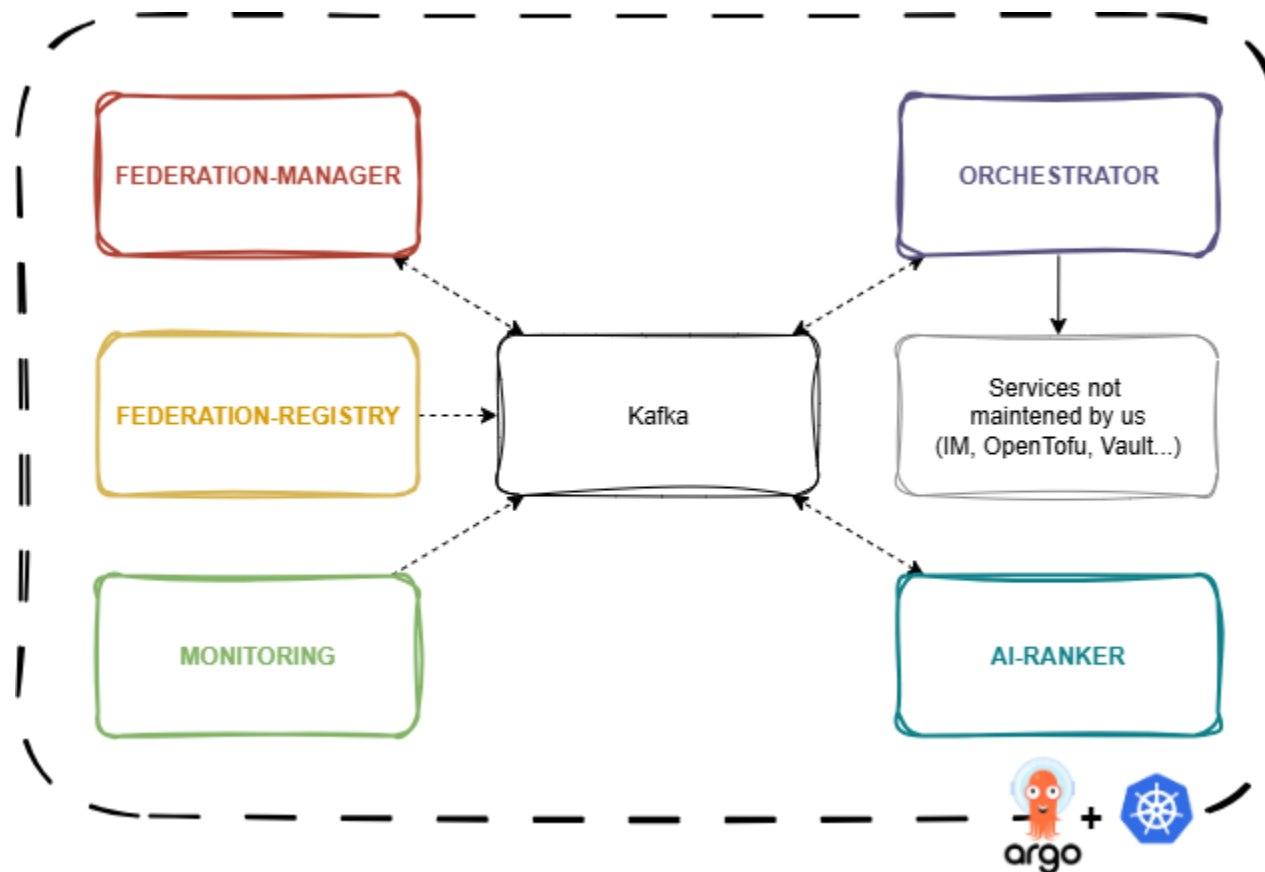
## Starting point





## Final destination

- Kafka
- Federation-Manager
- Federation-Registry
- Monitoring
- AI-Ranker
- Orchestrator
- Other services (not developed by us)
- Kubernetes cluster + ArgoCD





# Kafka

*Tool used to collect data to share among the micro-services and trigger automated tasks*

## Requirements:

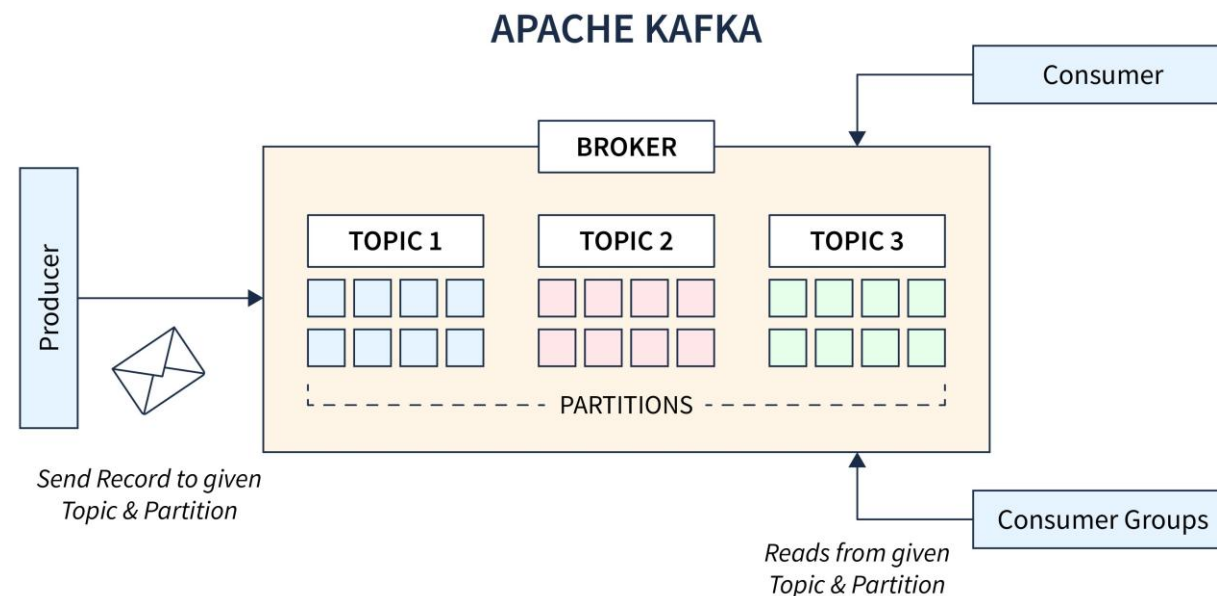
- High Availability

## Execution plan:

- Documentation
- Ansible role
- Deployment in production
- Define message retention time in topics
- Evaluate data dump on S3 storage for topics

## Known or possible problems:

- How to manage region replication/disaster recovery?





# Federation-Manager

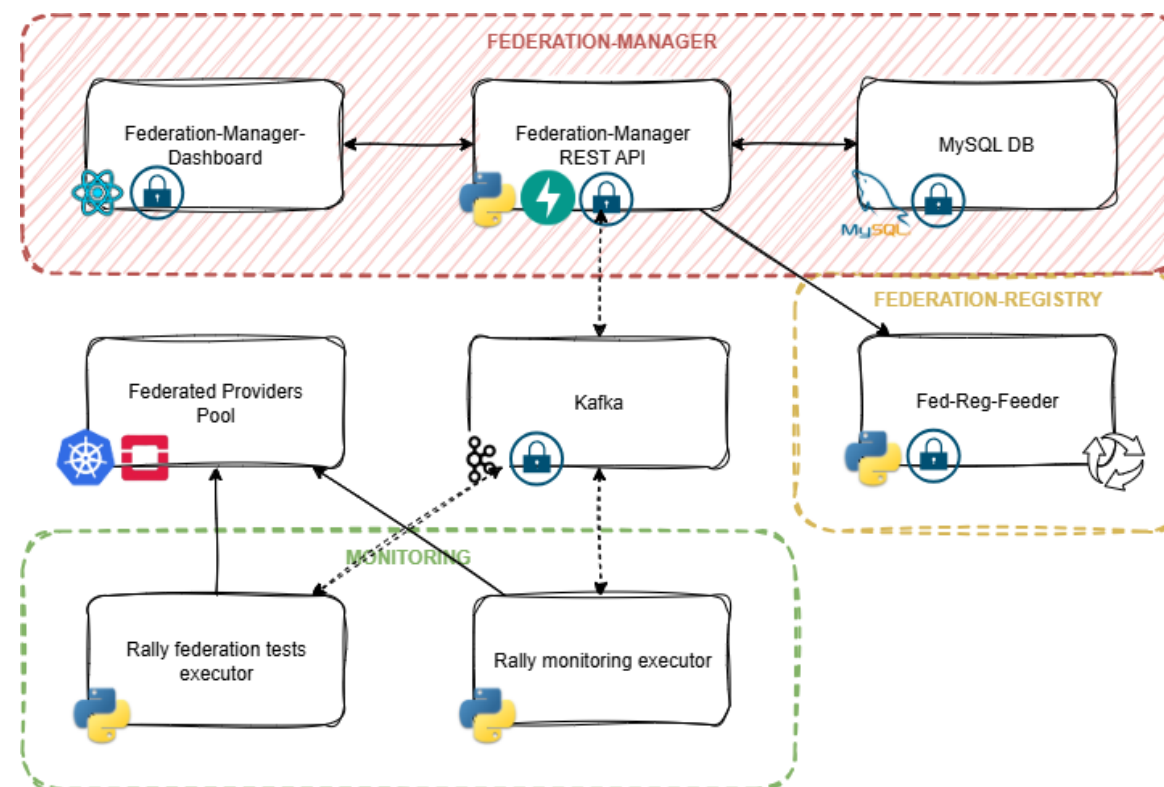
*Micro-service used to manage the provider federation procedure and new scientific groups on boarding*

## Requirements:

- Kafka
- OPA

## Execution plan:

- Federation-Manager-Dashboard implementation
- Federation-Manager REST API implementation
- Dedicated topics on Kafka
- Documentation

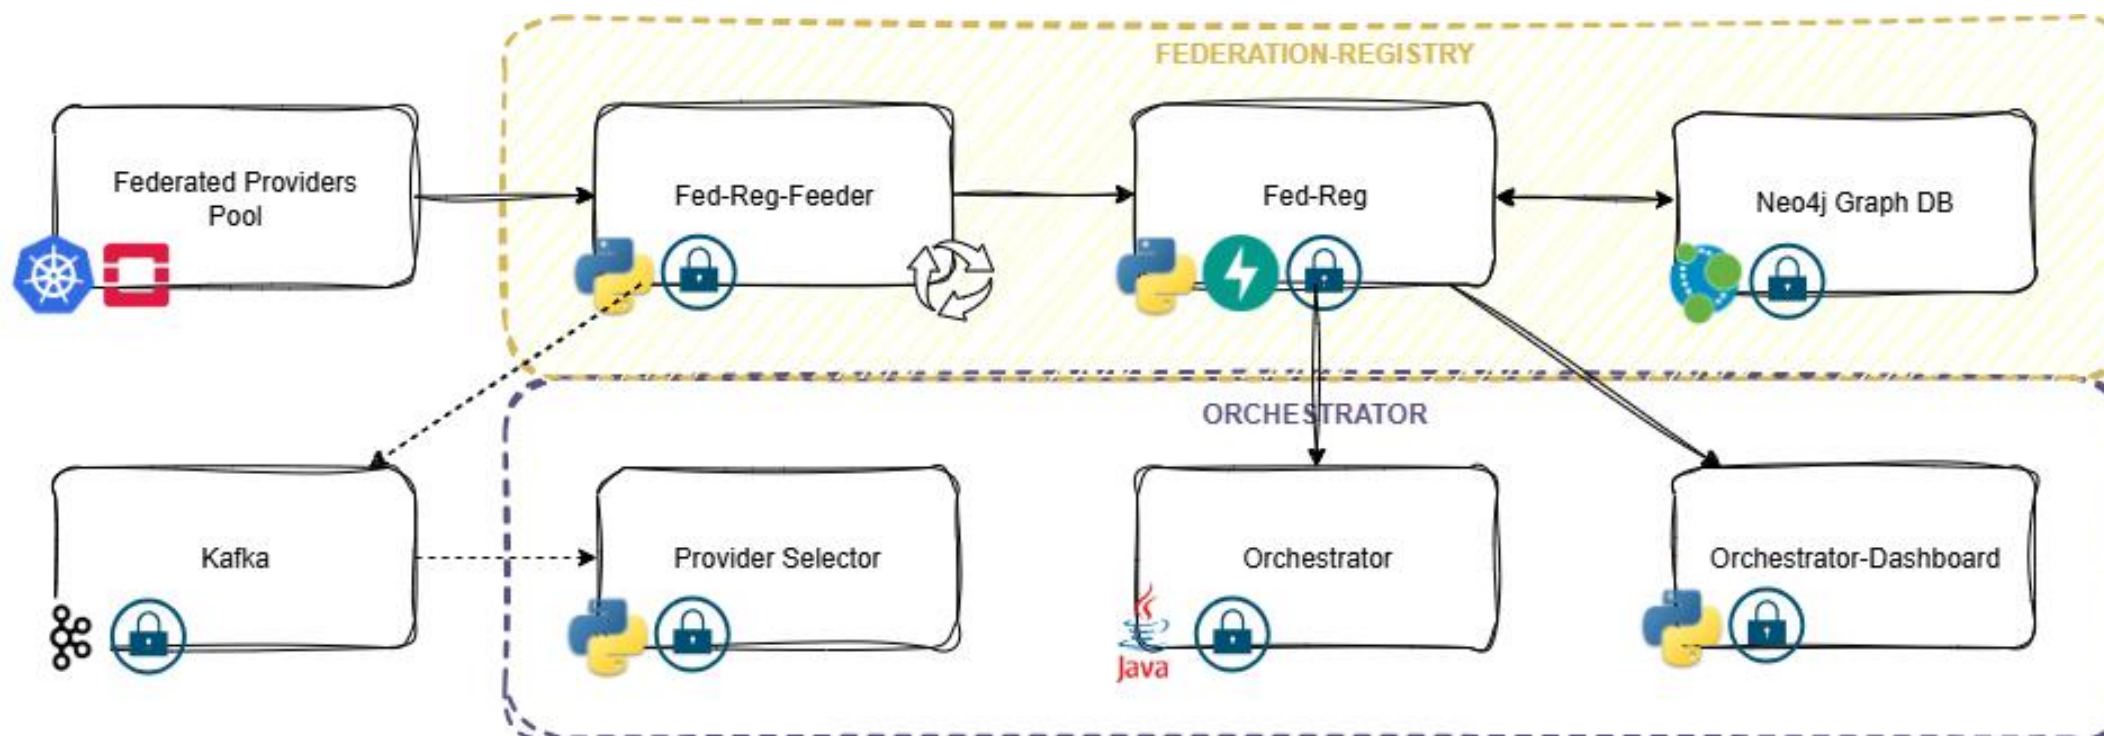






## Federation-Registry (State of the Art)

*Micro-service used to store federated provider configurations and SLAs details*





## Federation-Registry (Target architecture)

### Micro-service used to store federated provider configurations and SLAs details

#### Requirements:

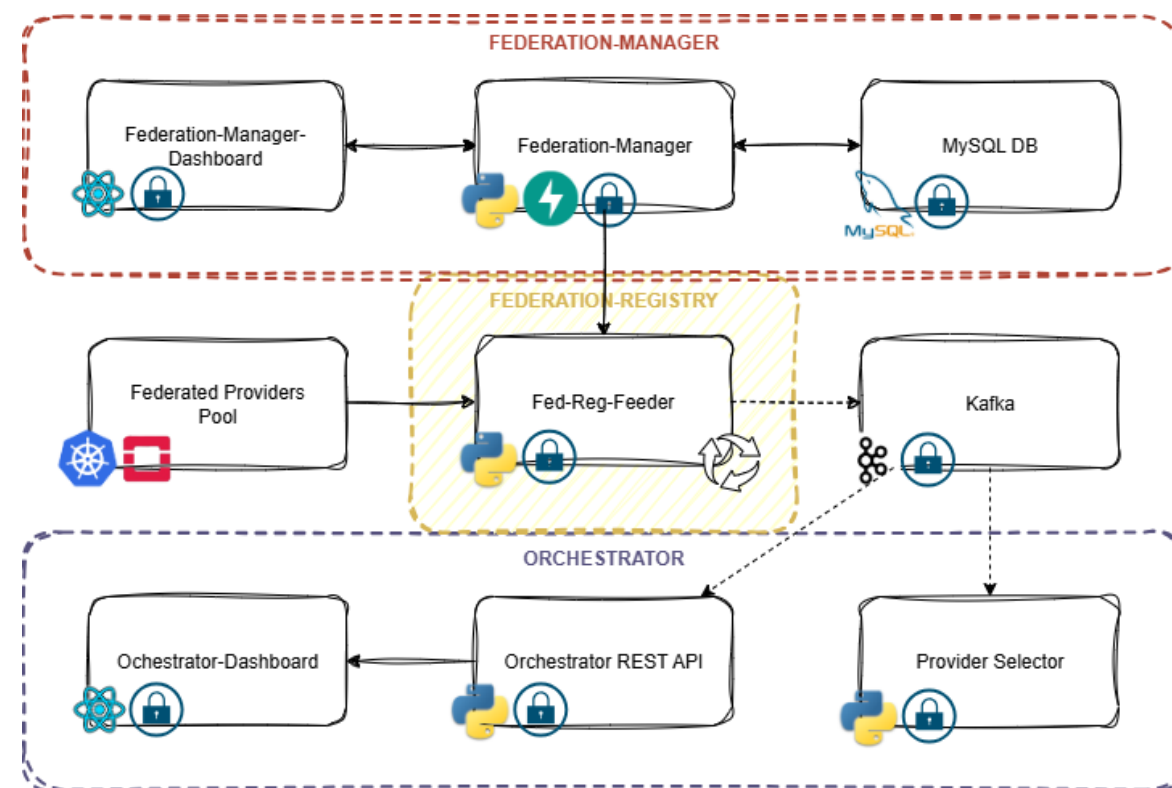
- Federation-Manager (currently using YAML files)
- Access to federated providers
  - Currently using **OIDC-Agent** and **ops** user.


#### Execution plan:

- Fed-Reg-Feeder: Complete support for Kubernetes clusters
- Fed-Reg-Feeder: Read configuration details from Fed-Mgr
- Fed-Reg-Feeder: Add support for RadosGW providers
- Fed-Reg-Feeder: Add support for Rucio services
- Remove Federation-Registry and Neo4j DB

#### Known or possible problems:

- Long running tokens used by the feeder generate problems with the application (restart required)
- OIDC-Agent needs account re-authentication (Client-credentials instead of user-based access to federated providers?)





## Monitoring (Federated resources)

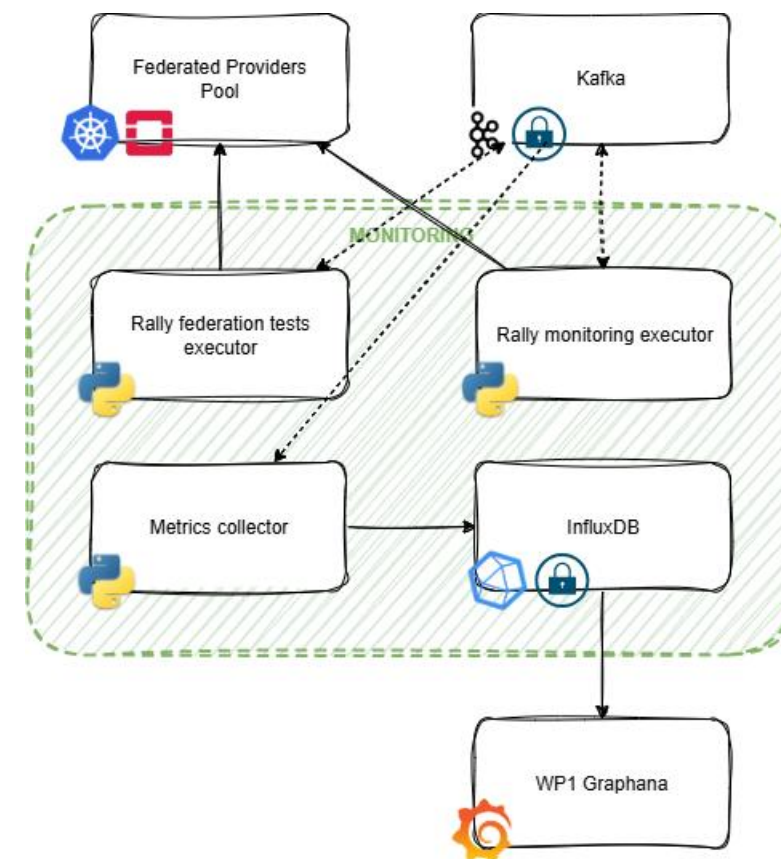
*Set of tools to monitor federated providers status and collect usage metrics*

### Requirements:


- Access to federated providers
- Kafka

### Execution plan:

- Rally federation tests executor implementation
- Rally monitoring executor implementation
- Metrics collector design and implementation
- InfluxDB installation and configuration
- WP1 Graphana connection

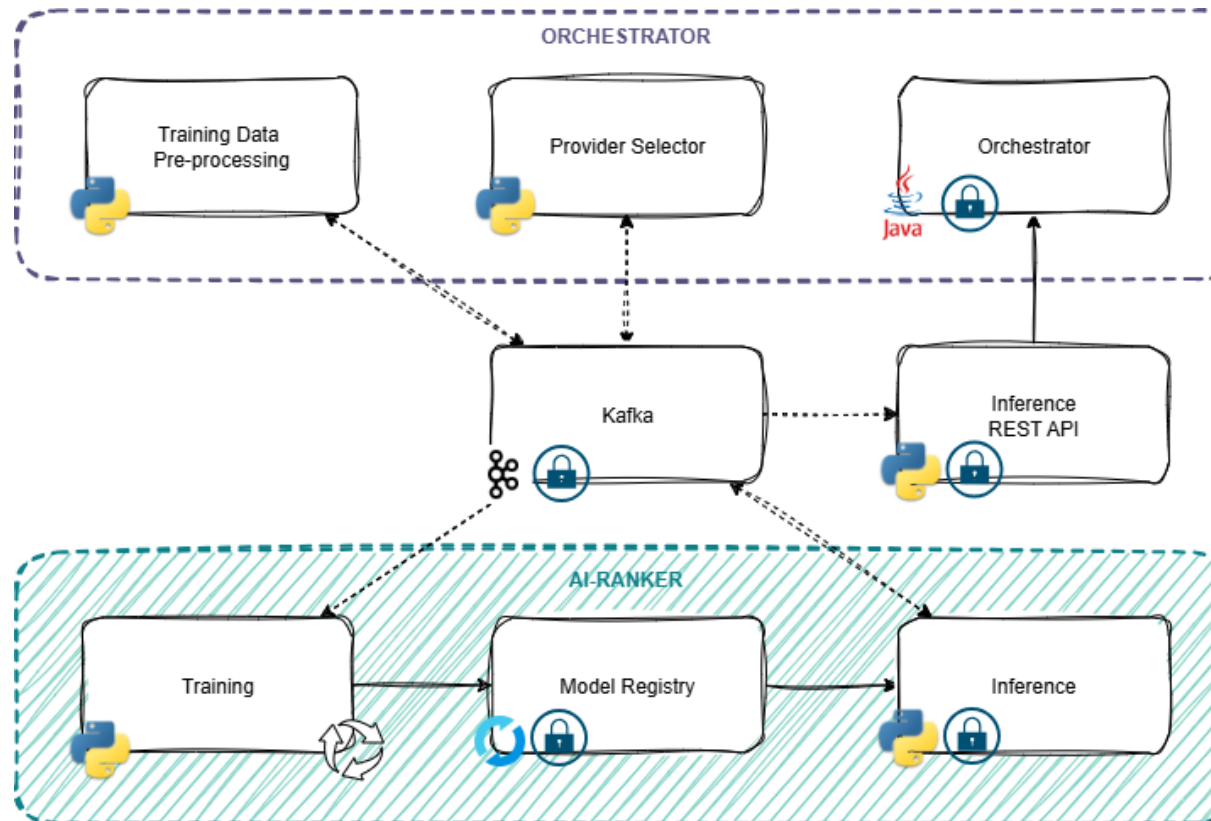







## AI-Ranker (State of the Art)

*Micro-service used to rank federated providers using ML techniques*





## AI-Ranker (Target Architecture)

*Micro-service used to rank federated providers using ML techniques*

### Requirements:

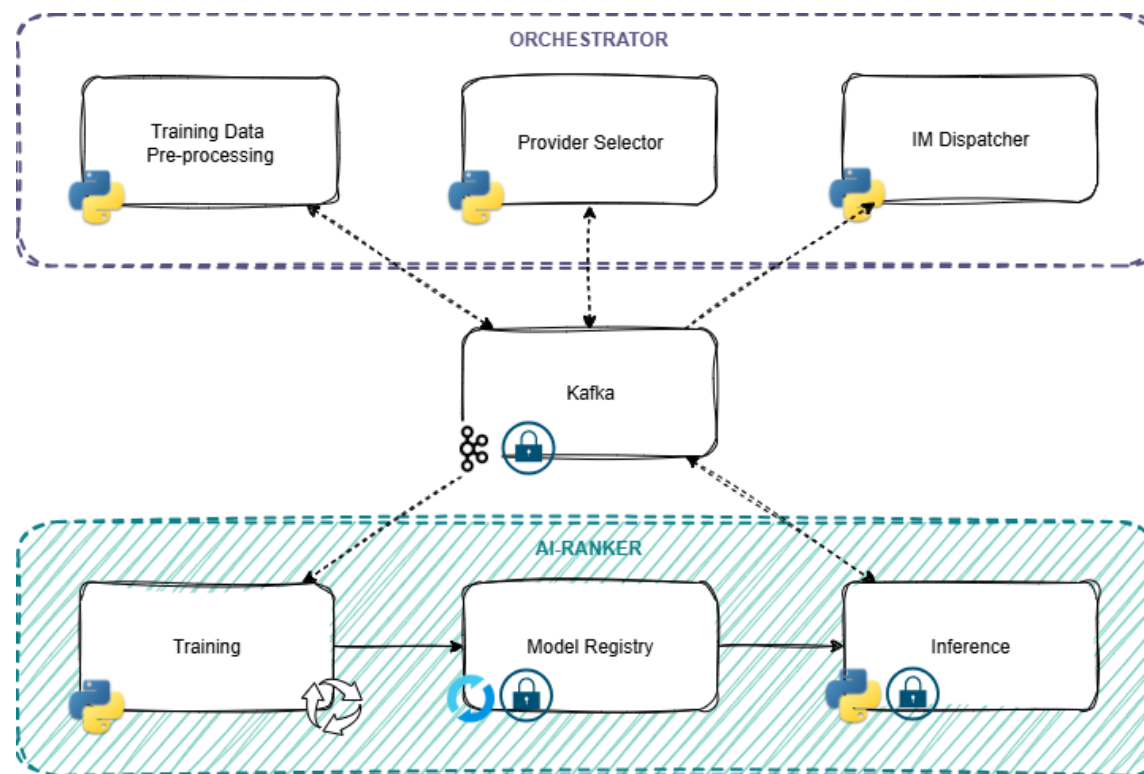
- Kafka

### Execution plan:

- Detect and manage edge cases
- Metrics definition for deployments targeting kubernetes providers

### Known or possible problems:

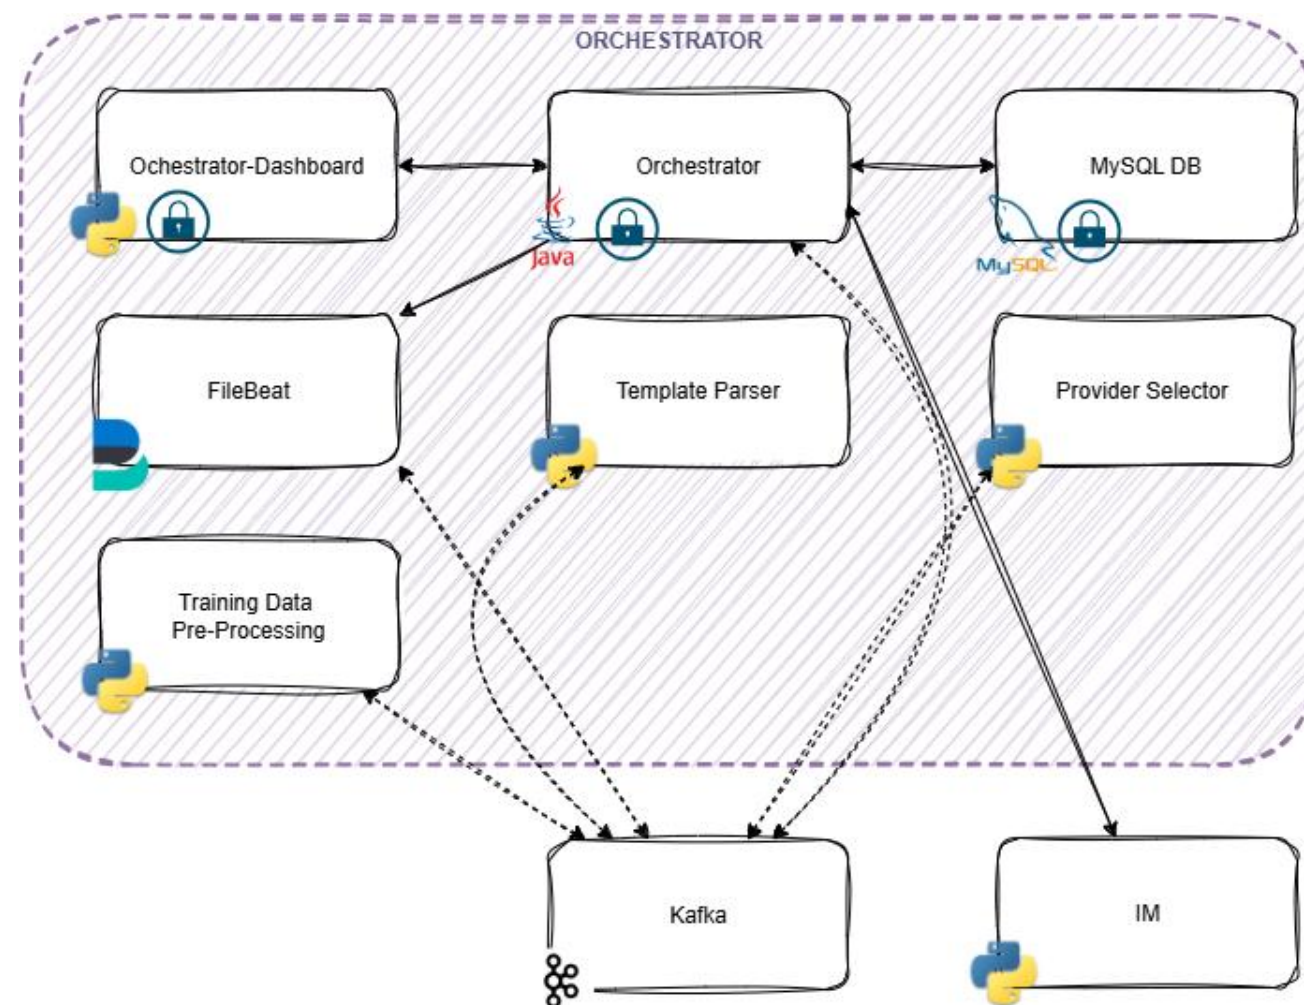
- Inference is reading again the whole training topic
- If the model registry is not available the Inference fails?






## Orchestrator (State of the Art)

*Set of components aiming to find a ranked list of federated providers where to deploy a user infrastructure and manage that deployment*







## Orchestrator (1)

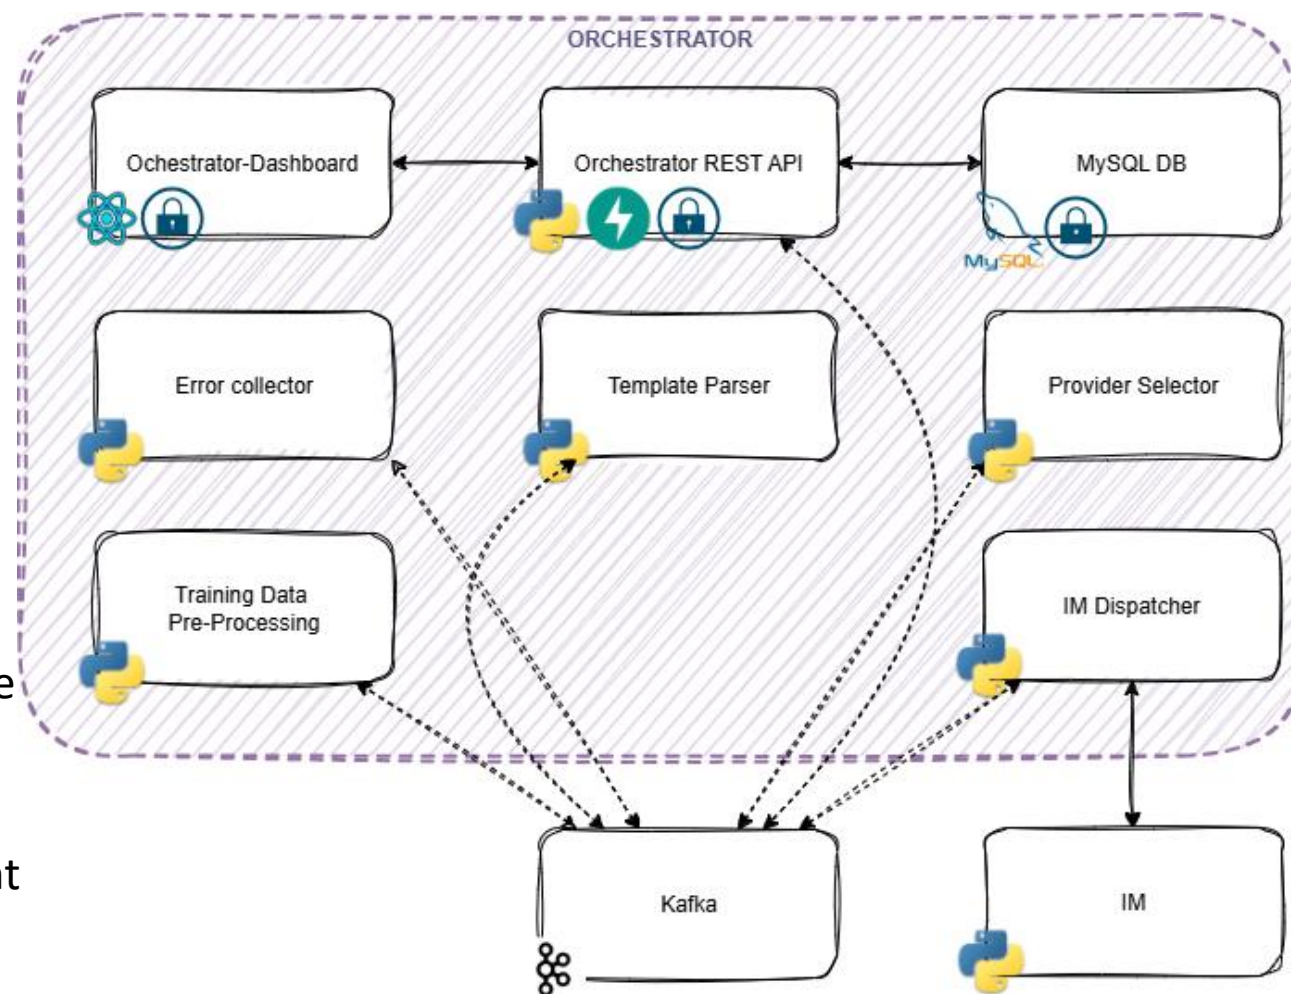
*Set of components aiming to find a ranked list of federated providers where to deploy a user infrastructure and manage that deployment*

### Requirements:

- Kafka
- OPA

### Execution plan:

- REST API revision and implementation
- Implementation of r/w logic on Kafka topics from the REST API
- Code refactoring of «PaaS-Mon» components
- Define and implement the error collector component
- Orchestrator-Dashboard (python) update





## Orchestrator (2)

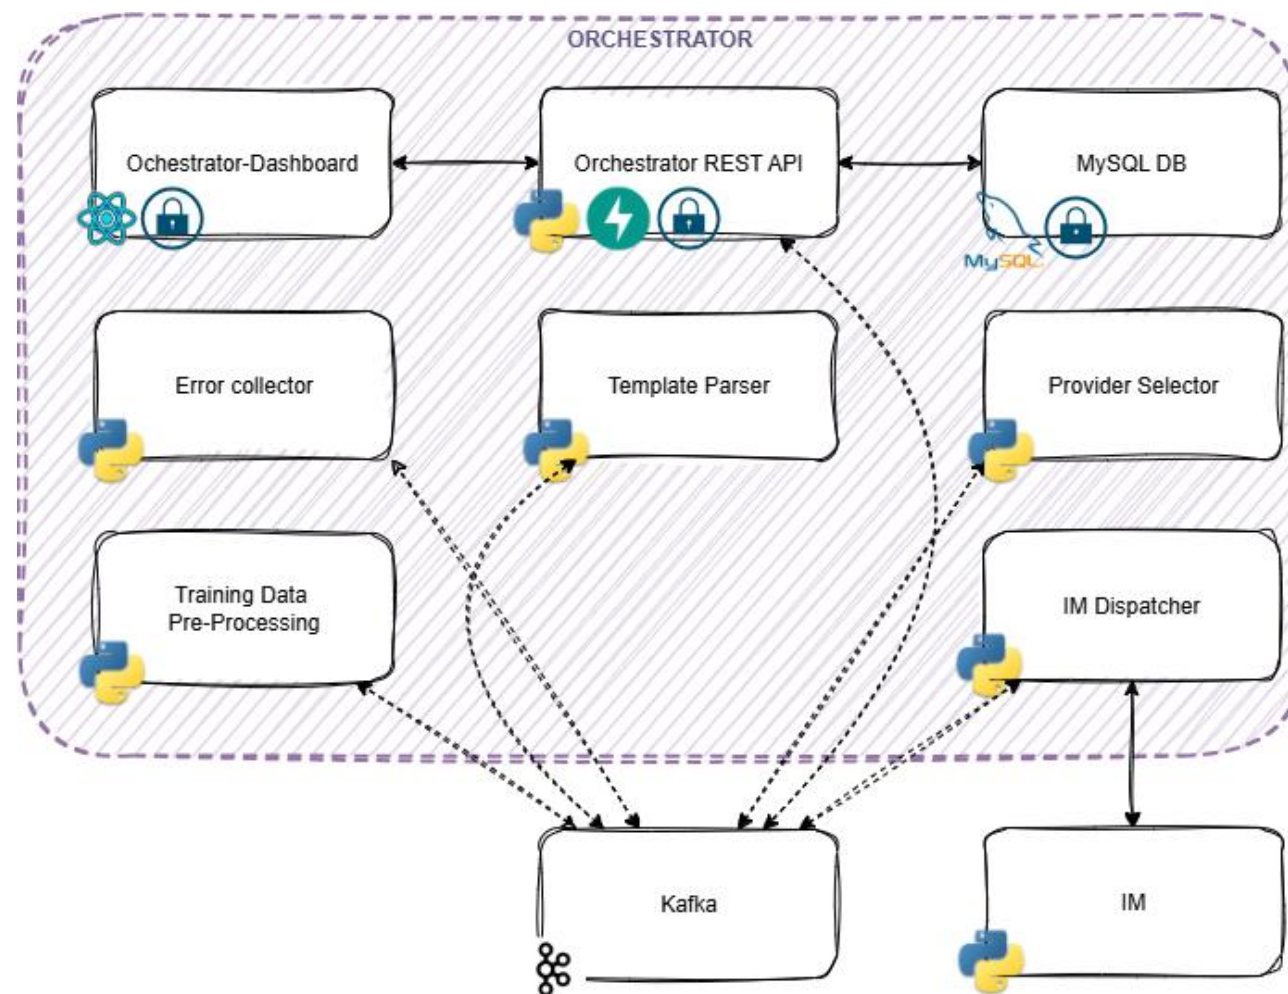
*Set of components aiming to find a ranked list of federated providers where to deploy a user infrastructure and manage that deployment*

### Execution plan:

- Integration of openstack-tosca-parser into the template parser component
- IM Dispatcher implementation for both Openstack and Kubernetes providers
- Use the typer library to develop the corresponding CLI (orchent no more needed)
- Orchestrator-Dashboard (TypeScript) implementation

### Known or possible problems:

- Improve message version management
- Avoid components unexpected termination
- Multithreading management on IM Dispatcher








## CVMFS Integration

### ***Requirements:***

- Communication with the Hashicorp vault
- Communication through rabbitmq with the CVMFS publisher

### ***Execution plan:***

- Edit the Orchestrator-Dashboard settings page to
  - Submit the creation of a new CVMFS repository
  - List CVMFS repositories



## Kubernetes + ArgoCD

- Cluster kubernetes with 3(more?) VM + ArgoCD on top of it
  - Production: one cluster on Bari and another one on CNAF
  - Pre-production
  - Development
  - ArgoCD must support OIDC authentication and authorization
    - To define access level
- New private repository on baltig with the **production** and **pre-production** configurations
- Service migration from VM + docker to Kubernetes
  - Convert docker-compose to kubernetes YAML (Already started for Fed-Mgr and Fed-Reg - [kompose](#))
  - Convert ansible-role and playbook into YAML for ArgoCD
- How to manage passwords and secrets ([Sealed Secrets](#))



## Before the end of the year

### ➤ Production

- Kafka + AI Ranker + «PaaS-Mon» components -> Gioacchino e Luca

### ➤ Pre-production

- ArgoCD + Kubernetes -> Francesco, Mauro e Luca
- Fed-mgr (deployed over k8s) -> Alessandra, Ettore, Jacopo e Giovanni
- IM Dispatcher (REST API, deployed over k8s) -> Caterina, Giovanni e Luca

### ➤ Prototypes

- REST API (Deployment creation through command line) -> Giovanni
- Dashboard python upgrade to communicate with the new REST API -> Michele



## Conferences and publications

### Conferences

- [ECAI](#) (25-30 Ottobre, Bologna)
- [ISGC](#) (16-21 Marzo, Taipei)
- [CHEP](#) (25-29 Maggio, Bangkok)
- Workshop di CCR (May-June)
- [ICCSA](#) (30 Giugno - 3 Luglio, Braga, Portogallo)

### Topics

- Fed-Mgr
- Deployment su Kubernetes+ArgoCD
- AI-Ranker evolution
- ...

*Goal: publish the new architecture of the PaaS Orchestration system in a peer-reviewed journal by July 2026. Possible journals:*

- ❖ [Transaction on cloud computing](#)
- ❖ [Journal of Cloud Computing Springer](#)
- ❖ [Computing Springer](#)
- ❖ [Future Generation Computer Systems](#)
- ❖ [Journal of Systems and Software](#)
- ❖ [Cluster Computing](#)



Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



# Thank you

[giovanni.savarese@ba.infn.it](mailto:giovanni.savarese@ba.infn.it),  
[luca.giommi@cnafr.infn.it](mailto:luca.giommi@cnafr.infn.it)

