



Developing Artificial Intelligence in the Cloud: The AI_INFN platform



Rosa Petrini

rosa.petrini@fi.infn.it

Scope and objectives

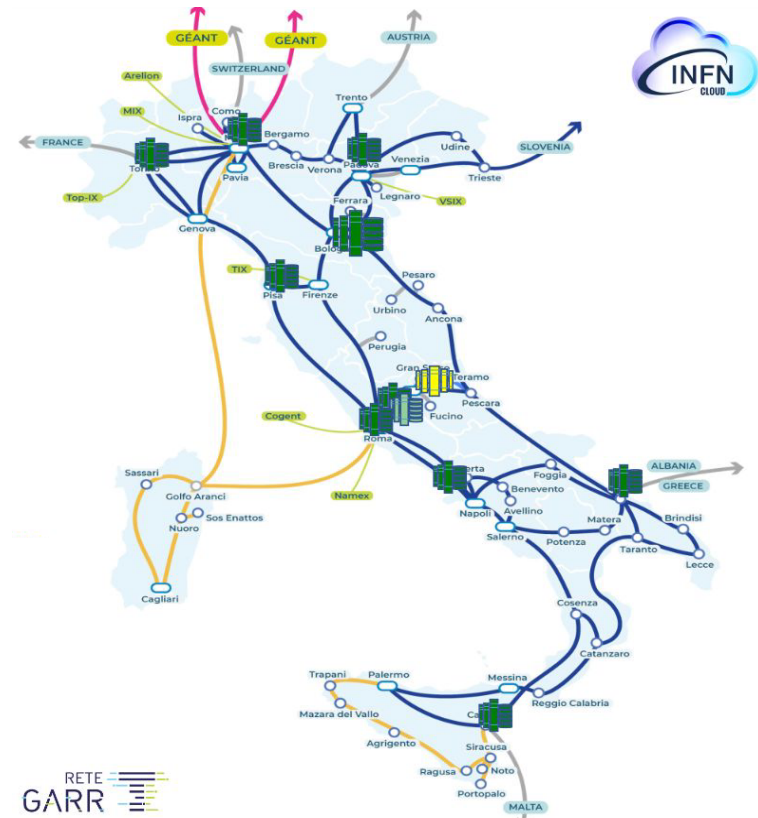
The provisioning of a **common, stable, and reliable** ground for researchers involved in ML to develop, review and share their applications, **crossing the borders between different communities**, INFN units, experiments and research domains

Provide a **centrally maintained cloud-based infrastructure** for interactive and batch ML fast prototyping, with access to modern hardware accelerators (GPU, FPGA...) and systems tuned for ML performance

DataCloud

In the framework of the current NRRP projects (ICSC, TeRABIT) INFN has a leading role in the creation of the **Italian Cloud Federation**, to access all Italian scientific computing resources through uniform interfaces

- **Tier-1** (CNAF)
- **Tier-2** (BA, CT, LNF, LNL/PD, NA, MI, PI, RM1, TO)
- Backbone and federated clouds
- HPC4DR (LNGS)
- **INFN Cloud:**
 - a data lake-centric, heterogeneous federated Cloud infrastructure spanning multiple sites across Italy, providing an extensible portfolio of solutions tailored to multidisciplinary scientific communities



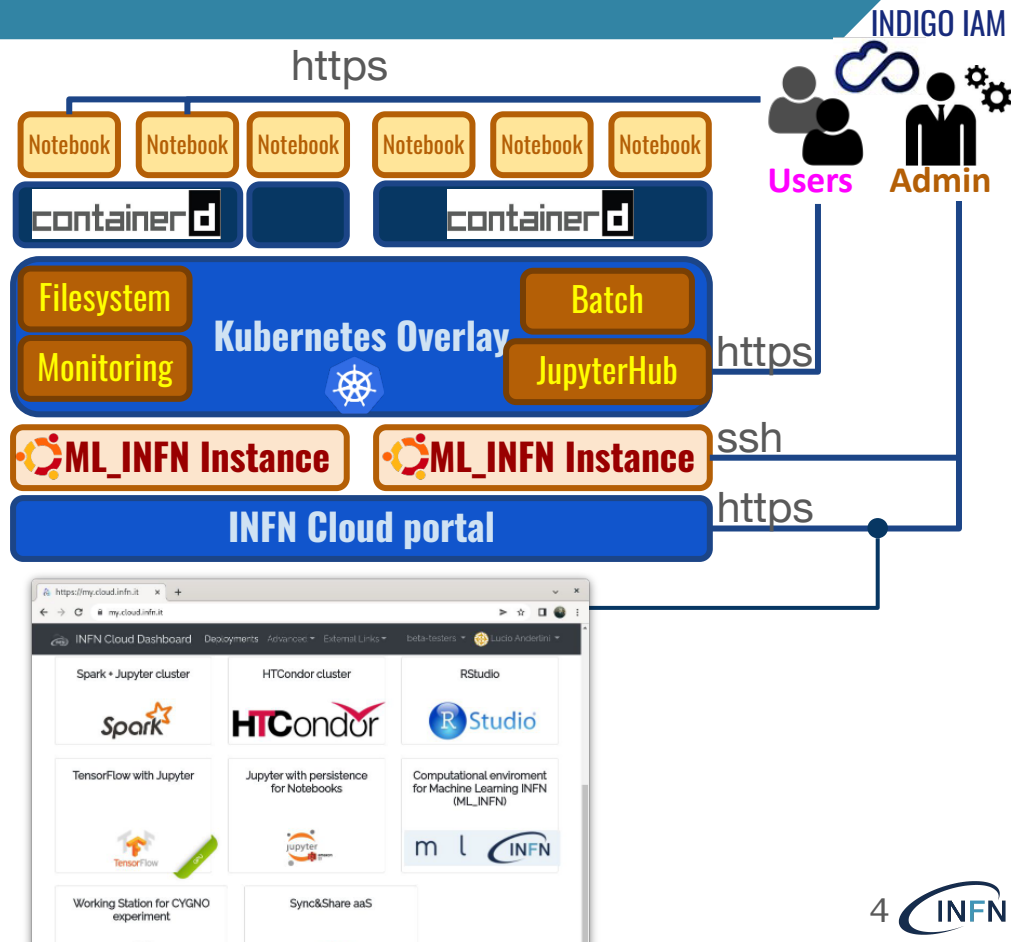
Computing resources available to AI_INFN are located in Room Tier-1 of CNAF and managed through a virtualization layer (**OpenStack of Cloud@CNAF**) in INFN Cloud.

INFN Cloud Resources: Architecture

- The ML_INFN outcome: **“sharing precious GPUs through the Cloud is feasible and effective!”**
- With **AI_INFN**, we improved on sharing by **decoupling data from computing resources**, with a **filesystem shared** across the VMs
- An additional abstract, elastic overlay is added on top of multiple VMs

Kubernetes Overlay:

- login via AAI → **INDIGO IAM**
 - Monitoring & Accounting
 - Managed software environments for ML
- Adding and removing VMs enables manual horizontal scaling



From interactive to batch jobs

- Once an analysis or the development of a model is mature, analysts want to scale it on more resources:
 - longer training time than available interactively;
 - freeing interactive resources for development;
 - parallel execution of multiple trials...




- We are developing a microservice ([vk-dispatcher](#)) translating an interactive session into a [Kubernetes Job](#), executed on the cluster resources.



Development is our priority!

Batch workloads must not affect the interactive use of the platform.



Need for a batch management system, **instantaneously evicting batch** jobs opportunistically.

From interactive to batch jobs

- Once an analysis or the development of a model is mature, analysts want to scale it on more resources:
 - longer training time than available interactively;
 - freeing interactive resources for development;
 - parallel execution of multiple trials...
- We are developing a microservice (`vk-dispatcher`) translating an interactive session into a Kubernetes Job, executed on the cluster resources.



Development is our priority!

Batch workloads must not affect the interactive use of the platform.



Kueue



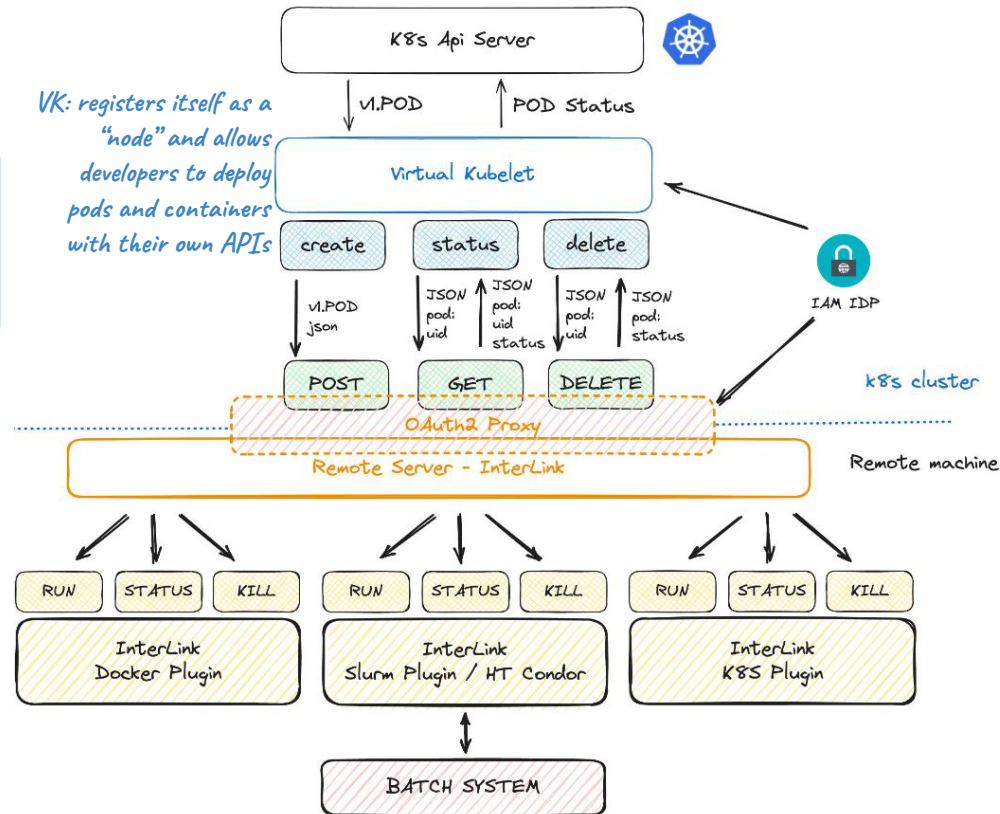
Offloading: InterLink & Virtual Kubelet

interLink

Provide an abstraction for the execution of a Kubernetes pod on any remote resource capable of managing a Container execution lifecycle.

The project consists of 3 main components:

- **Kubernetes Virtual Node:** based on the [VirtualKubelet](#) technology. Translating request for a kubernetes pod execution into a remote call to the interLink API server.
- **InterLink API server:** a modular and pluggable REST server that can handle requests coming from the VK and forwarding them to the sidecar;
- **Sidecar:** runs the containers on the infrastructure and returns the result. It can also communicate with the InterLink server.



Offloading: InterLink & Virtual Kubelet

interLink



VK: registers itself as a


v1.POD

POD Status



them to the sidecar;

- **Sidecar:** runs the containers on the infrastructure and returns the result. It can also communicate with the InterLink server.



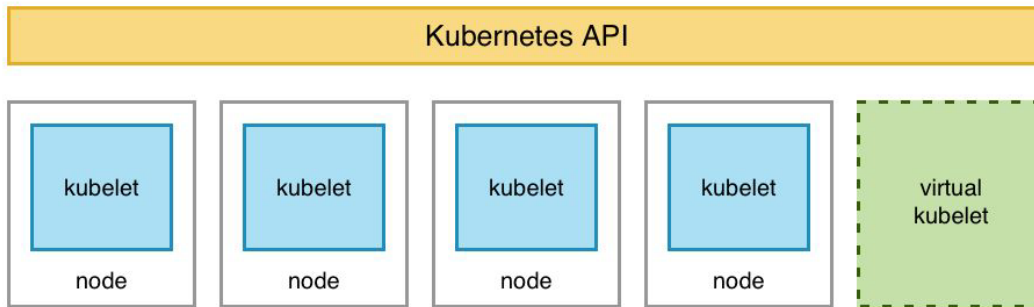
Conclusion

The AI_INFN platform is an intense R&D program with the aim of provide a more effective tool for developing Machine Learning and Artificial Intelligence for INFN researches.

Thank you for
your attention!



Offloading: InterLink & Virtual Kubelet



Typical kubeletes implement the pod and container operations for each node as usual.

