The INFN Computing Infrastructure and Machine Learning

AI@INFN Workshop, 2-3/5/2022

Davide Salomoni (davide@infn.it)

Background: the main drivers for the next 10 years of scientific computing for INFN

- Infrastructure
  - Renew infrastructures to be ready for the High Luminosity-LHC (HL-LHC) era, up to ~2035 or more
  - Use more compact computing (from today's ~20 kW/rack to 80 or more)
  - Lower the PUE (*power usage effectiveness*), be greener
  - Extend and expand networking for a future-proof infrastructure
- Hardware and Software
  - Foster and simplify the utilization of more viable technologies (*Eur/task* or *J/task*), like GPUs, FPGA, ... down to Quantum when available
  - Be more efficient, elastic and resilient
    - Pervasive use of geographically distributed storage ("the datalake")
    - Abstract from physical machines, and **form a national pool of resources and services ("the Cloud")**
    - Extend elastically to external providers such as HPC@CINECA or other cloud providers (via "dynamic federations")

Background: An history of collaborations

- The development of **Scientific Computing within INFN** was originally driven by the needs of its own theoretical and experimental communities; however, being at the forefront of computing in research seeded **many projects with a much broader scope**.
- The **key overall driver** was always to let our users **effectively exploit all available resources and technologies**.

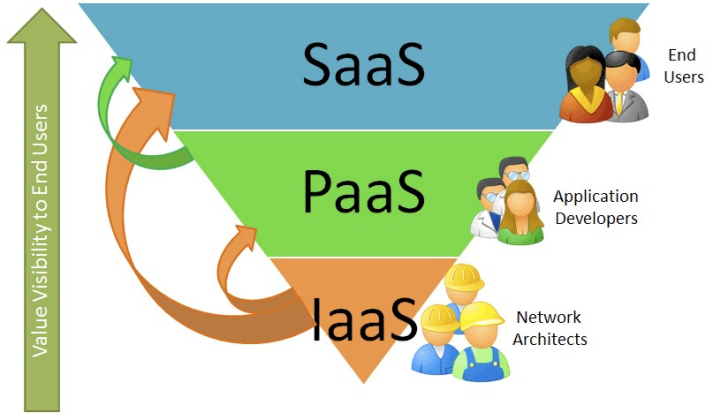

"Preparing the GRID"

"Preparing the Cloud"

"Expanding beyond HEP"

INFN and Computing / Big Data

- A long tradition in **state-of-the-art distributed IT technologies and solutions**, from the first small clusters to Grid and Cloud-based computing.
- INFN is not interested in computing per-se, but as an essential way to **support its research and mission**.
- INFN operates Grid and Cloud services based on its own:
  - 9 medium size centers
  - 1 large national center, at CNAF (Bologna) – **with an area certified ISO/IEC 27001, 27017, 27018**
- All the INFN centers are connected through 10-100 Gbit/s dedicated links via the GARR network.
- Collectively, our distributed infrastructure currently offers about **140,000 CPU cores, 120PB of enterprise-level disk space, 100PB of tape storage**.

The New INFN Computing Structure

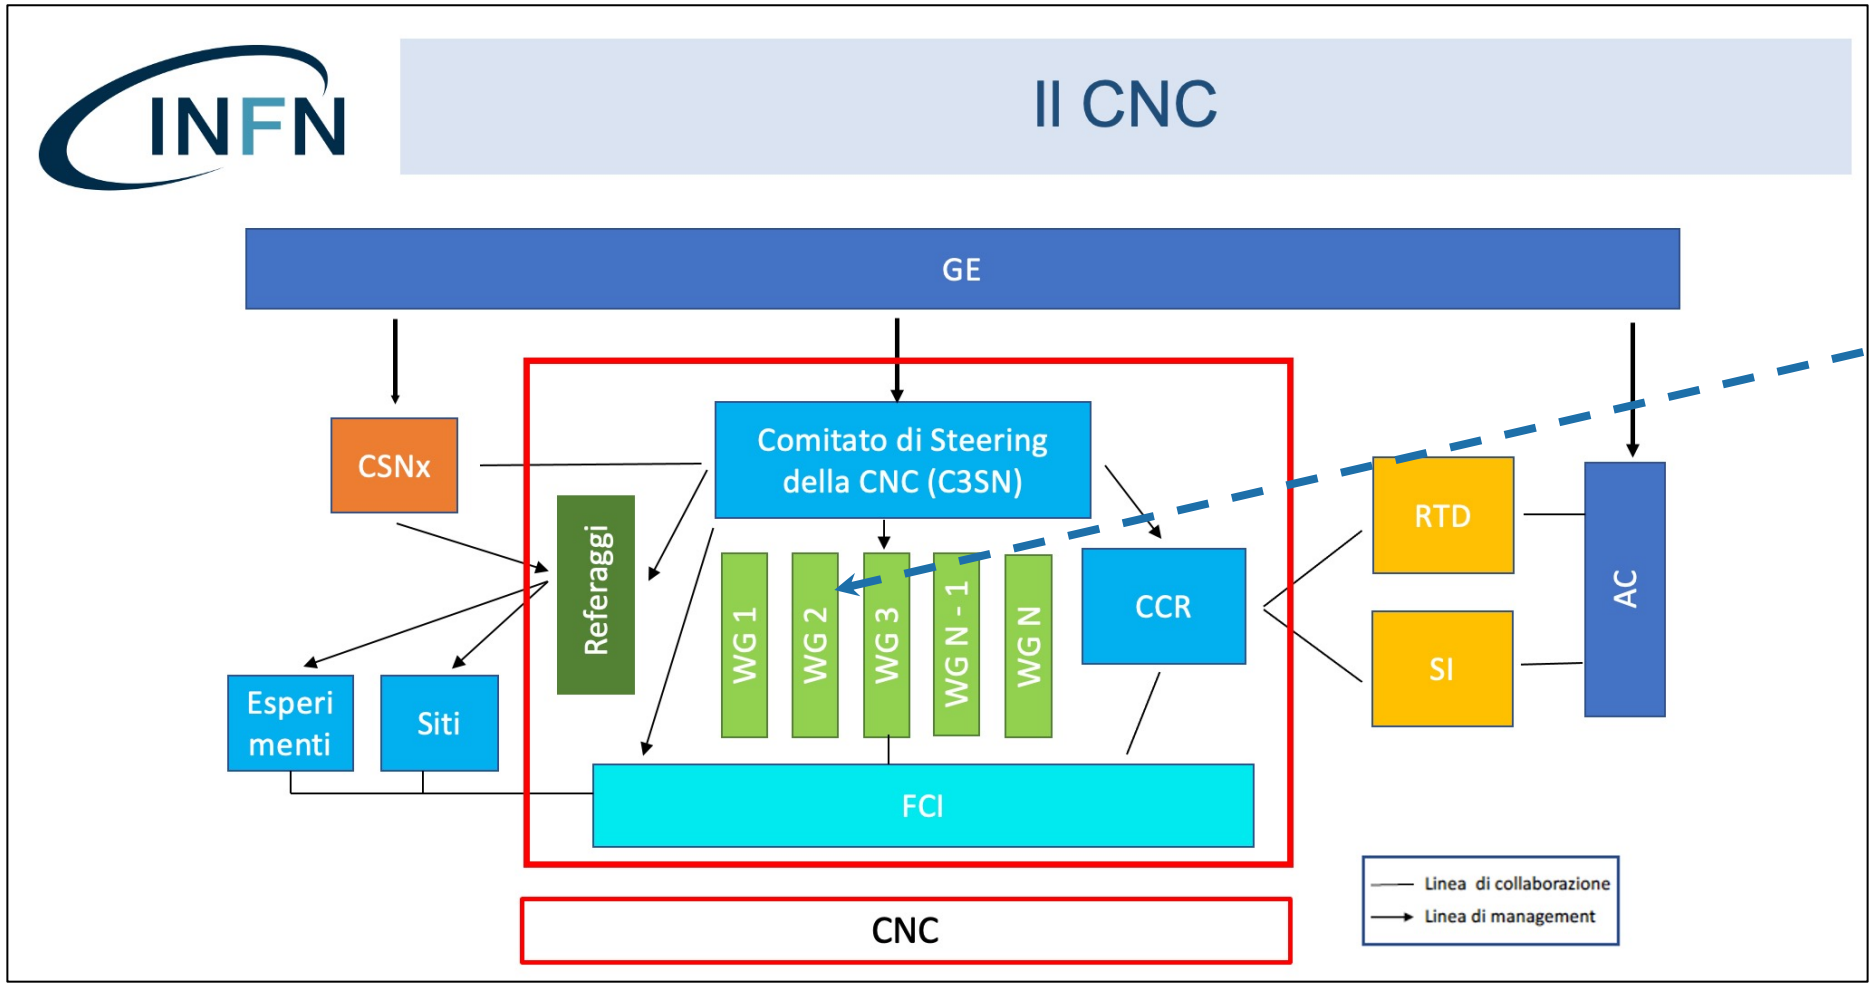

One of these WGs is the **WG Infrastruttura,** extending and evolving the current **INFN Cloud** project.

Il WG Infrastruttura

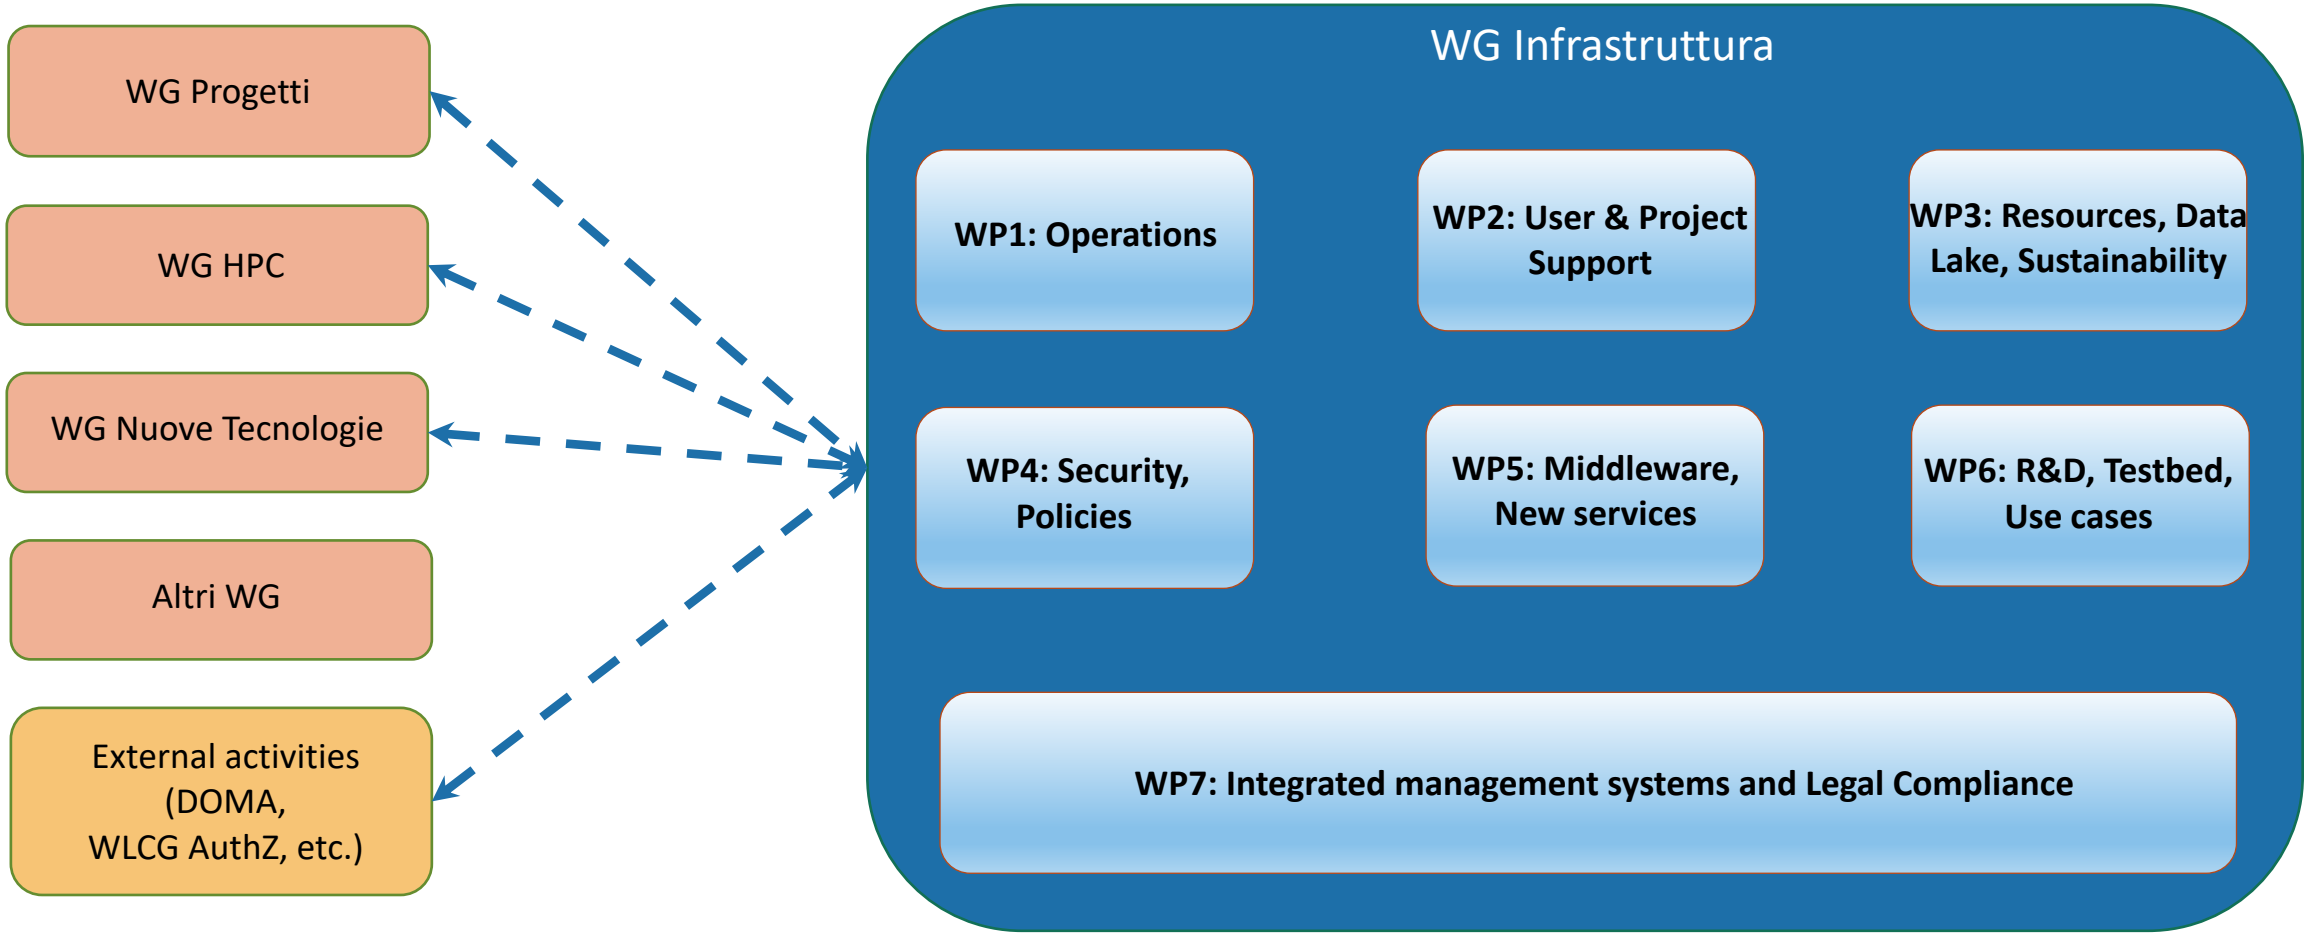

The Cloud national infrastructure for supercomputing

400 M€ Total Budget

188 M€ Cloud Infrastructure

40 M€ Open Call

40 M€ Innovation & TT

42% Investment South Regions

34 MUR Universities and Research institutions
15 Private Companies
1575 Researchers and Engineers
250 New Temporary positions
250 New PhD
40 % Female

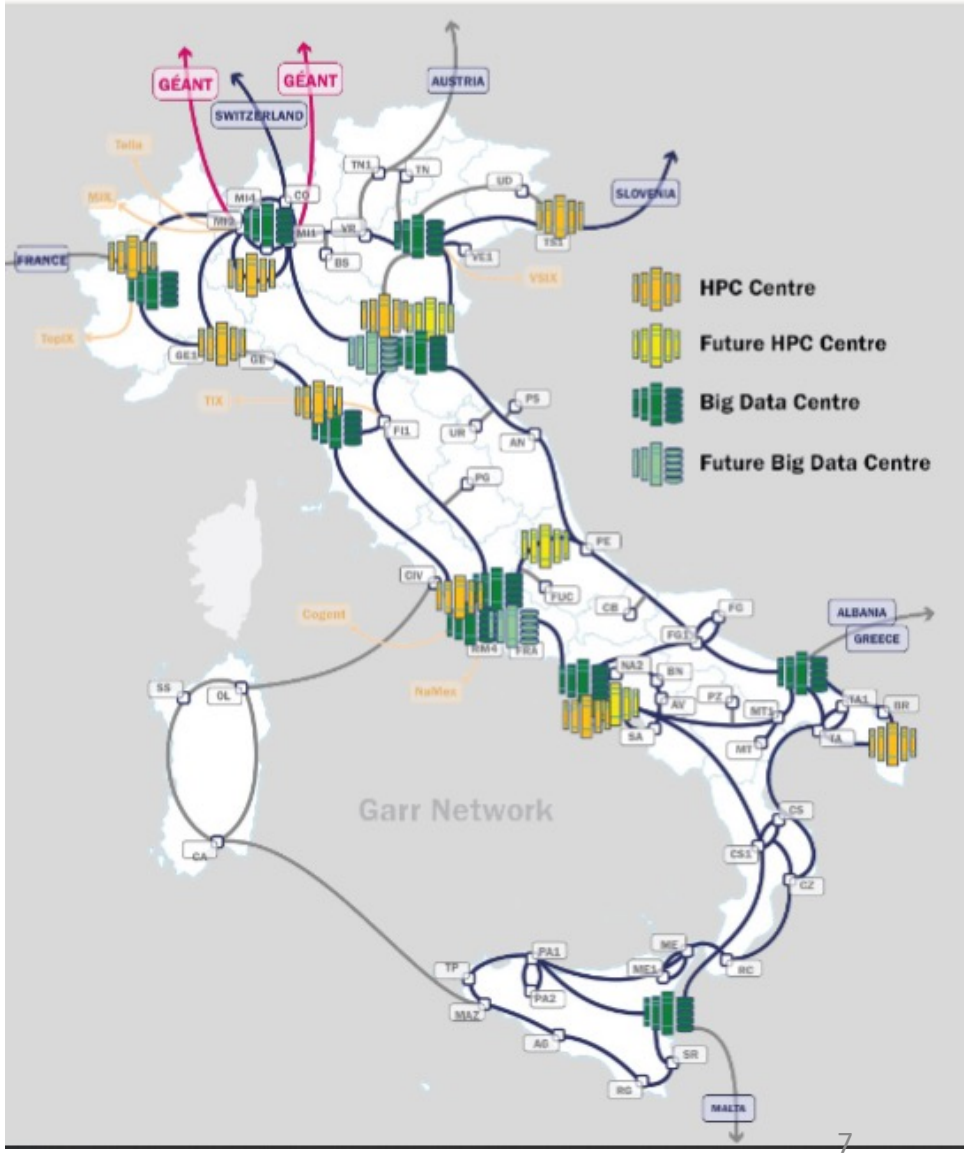

What about ML then?

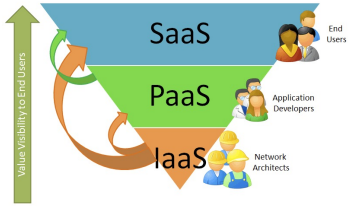

- **INFN Cloud** is the **starting point** for a National Datalake for research and beyond, building on (existing|renewed|new) e-Infrastructures.
  - Go to the INFN Cloud website (https://www.cloud.infn.it/) if you are interested in getting access.
- It is the **base of the evolution** of the INFN Distributed Computing vision, offering simple and complex services, **including ML-related ones**.
- Built on a **thin middleware layer** running on top of multiple federated clouds, decoupling physical and logical views via a **service composition** mechanism.
- In perspective, it will be the **Italian Node of the WLCG datalake** for HL-LHC.
- **INFN Cloud has >1y of production service** (plus about 1y of beta testing).
- It is based on **scalable open-source tools** and on a **twenty-year INFN experience** with distributed infrastructures, services and projects.

How INFN Cloud services work

- The **INFN Cloud Portfolio** is dynamically extendable thanks to a PaaS-level approach, centered on:
  1. A **distributed resource orchestration** framework (INDIGO Orchestrator).
  2. A standard—based **federated solution for identity access management** (INDIGO IAM).
  3. An **Infrastructure as Code** methodology, realized via a combination of:
     - TOSCA templates used 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.

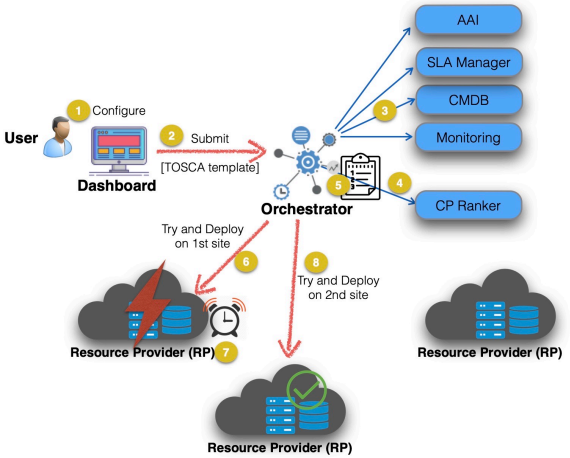

This is how all the INFN Cloud services are built, reused and expanded (Service Composition)

Excerpt from the INFN Cloud portfolio

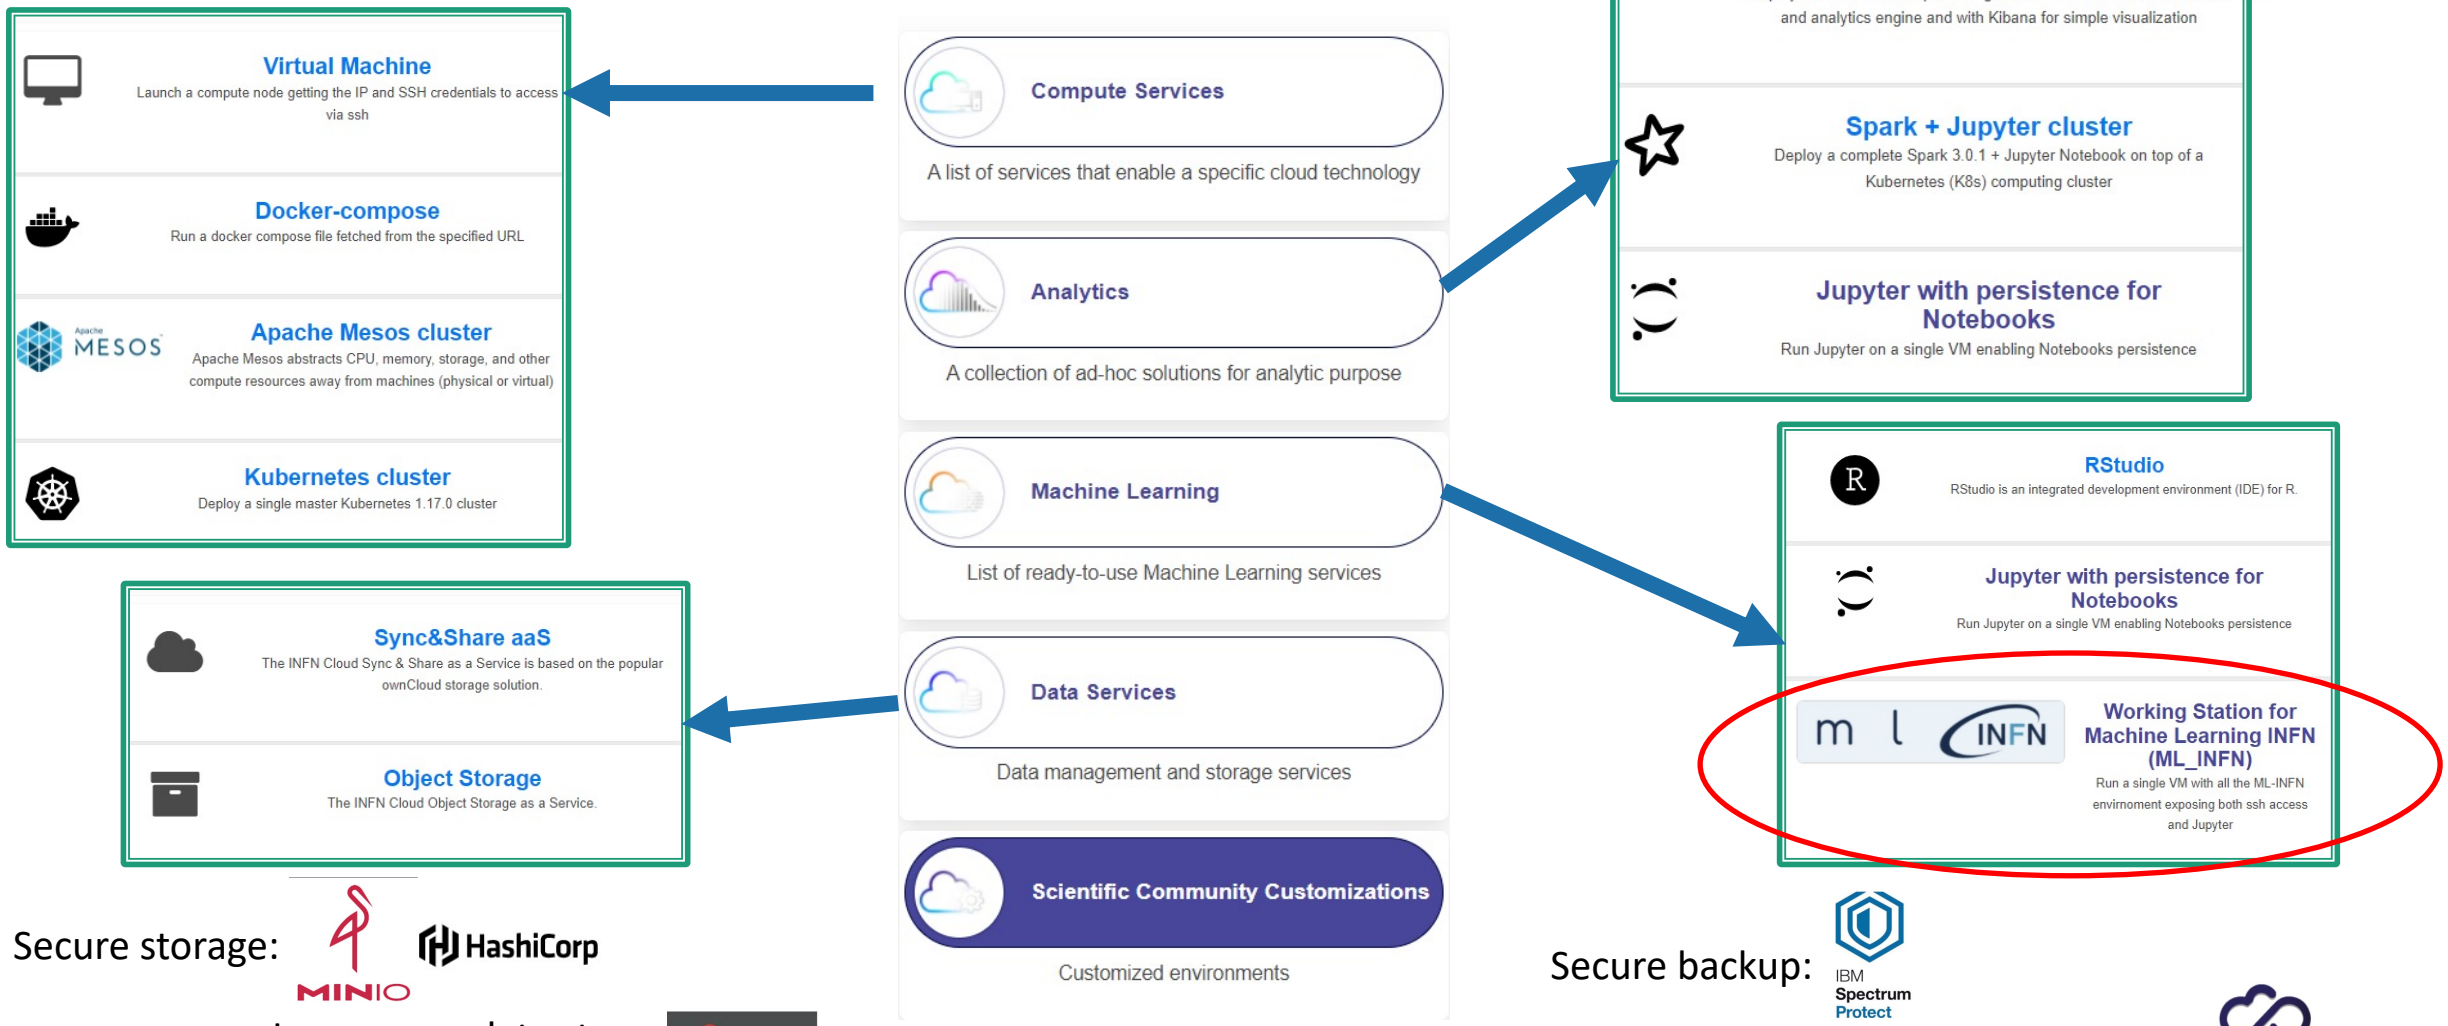

Secure storage: MINIO HashiCorp

In-memory data store: redis

Secure backup: IBM Spectrum Protect

PaaS Orchestrator & INDIGO IM

Selectable storage QoS levels: fast (SSD), normal (HDD), archive (tape-backed), remote replicas

A flexible federated datalake

Multiple ways to ingest and process data are possible. For example, for sensitive data (e.g., in the nation-wide Health Big Data project), we propose:

1. **Central harvesting** of data generated remotely
2. **Edge-level anonymization**, followed by central ingestion and analysis of data
3. **Edge-level feature extraction**, followed by central ingestion and analysis of features
4. **Federated learning** based on edge-level training, followed by publishing of the trained methods and by inference performed either centrally or at other edge locations.

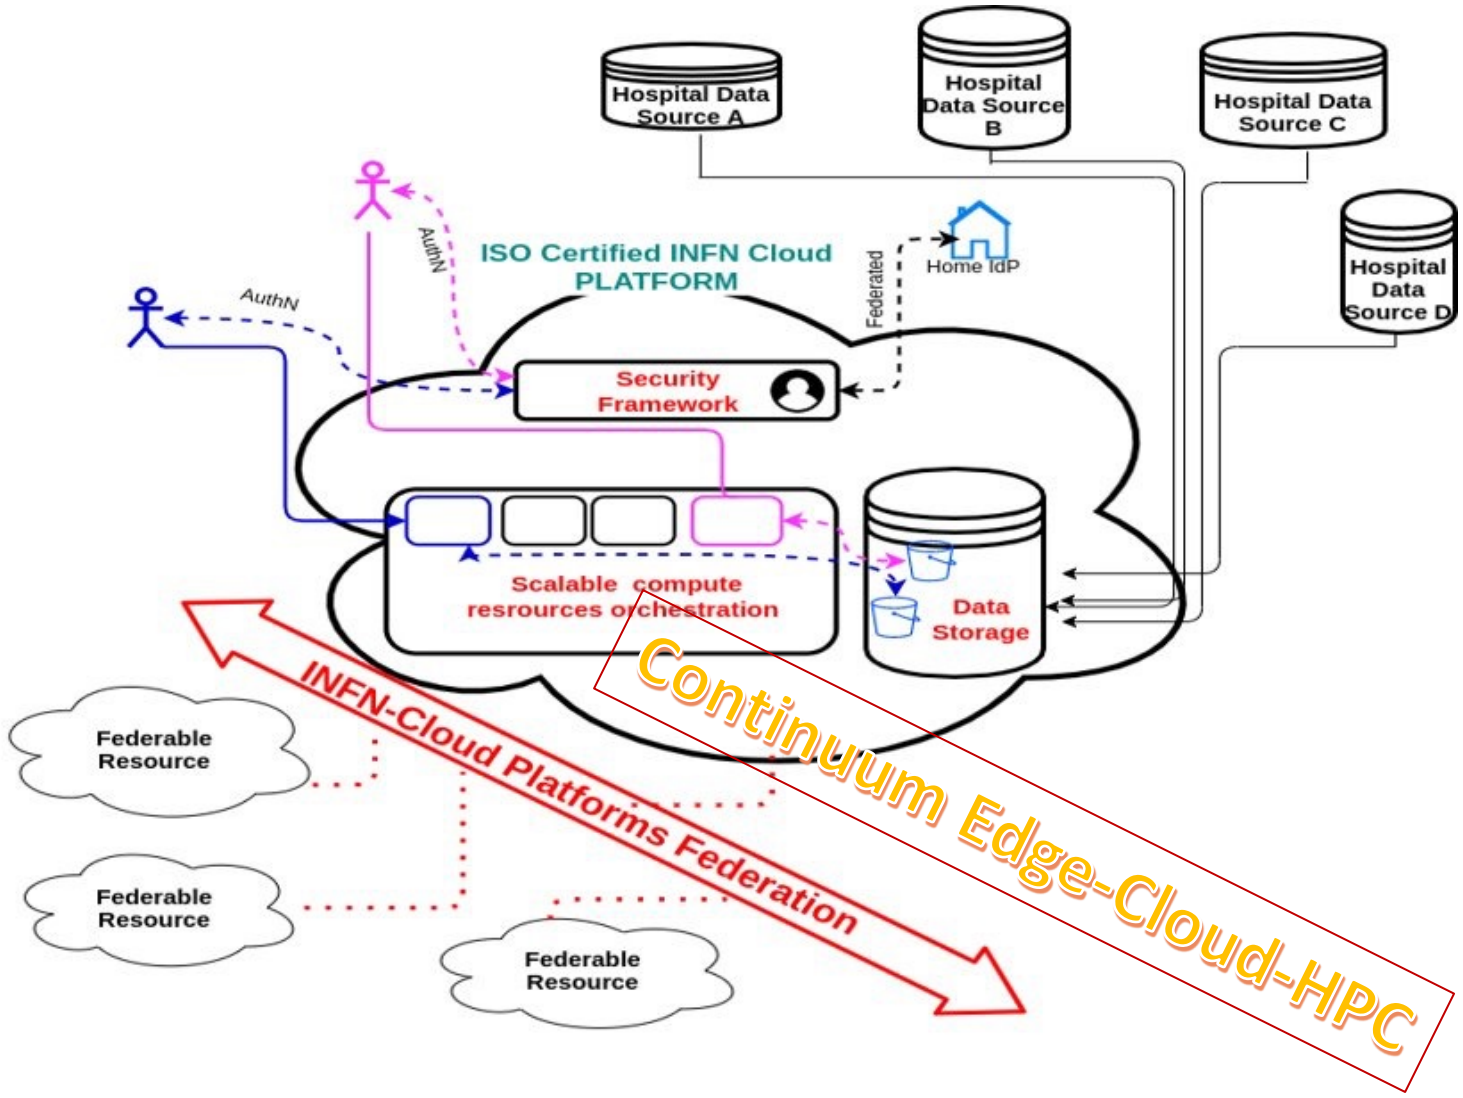

Main directions

- For existing ML_INFN solutions over INFN Cloud, see Lucio Anderlini's presentation.
- **Looking ahead**, we are working on the following topics of possibly direct relevance for the ML area:
  - The creation of Federated HPC Bubbles. Thanks to PNRR programs, we expect to provision multiple HPC resources to be deployed to INFN sites federated with INFN Cloud and integrate them into a national datalake.
  - Creation of distributed virtual availability zones to increase resiliency and to better exploit resources.
  - Creation of Content Distribution Points for distributed caches to reduce latency.
  - Data ingestion & analysis for IoT devices.
  - Expanded integration of heterogenous hardware, such as FPGA, GPU, CPU, storage with different QoS characteristics.
    - See for instance here for the first experiments on creating an "FPGA as a Service for ML", possibly transparently integrating also FPGA AWS resources into INFN Cloud.

ML-related services

- In order to build a distributed ML system, we need to effectively **manage a pipeline of steps that get composed in ML applications**. For this, we are currently looking to simplify services based on:
  - Support for data ingestion in both streaming (long-running services) and non-streaming (processing via batch jobs) modes.
  - Flexible and auto-scalable distributed resources, using Kubernetes.
  - Simplified building of ML and DL models, using TensorFlow.
  - Definition and orchestration of complex pipelines, using Argo Workflows.
  - Implementation and evolution of the DEEP-Hybrid DataCloud architectural model (described here).
- The INFN Cloud platform is **already available** to co-design and evolve these and other features so, if you are interested, contact us.

The DEEP-Hybrid DataCloud Model

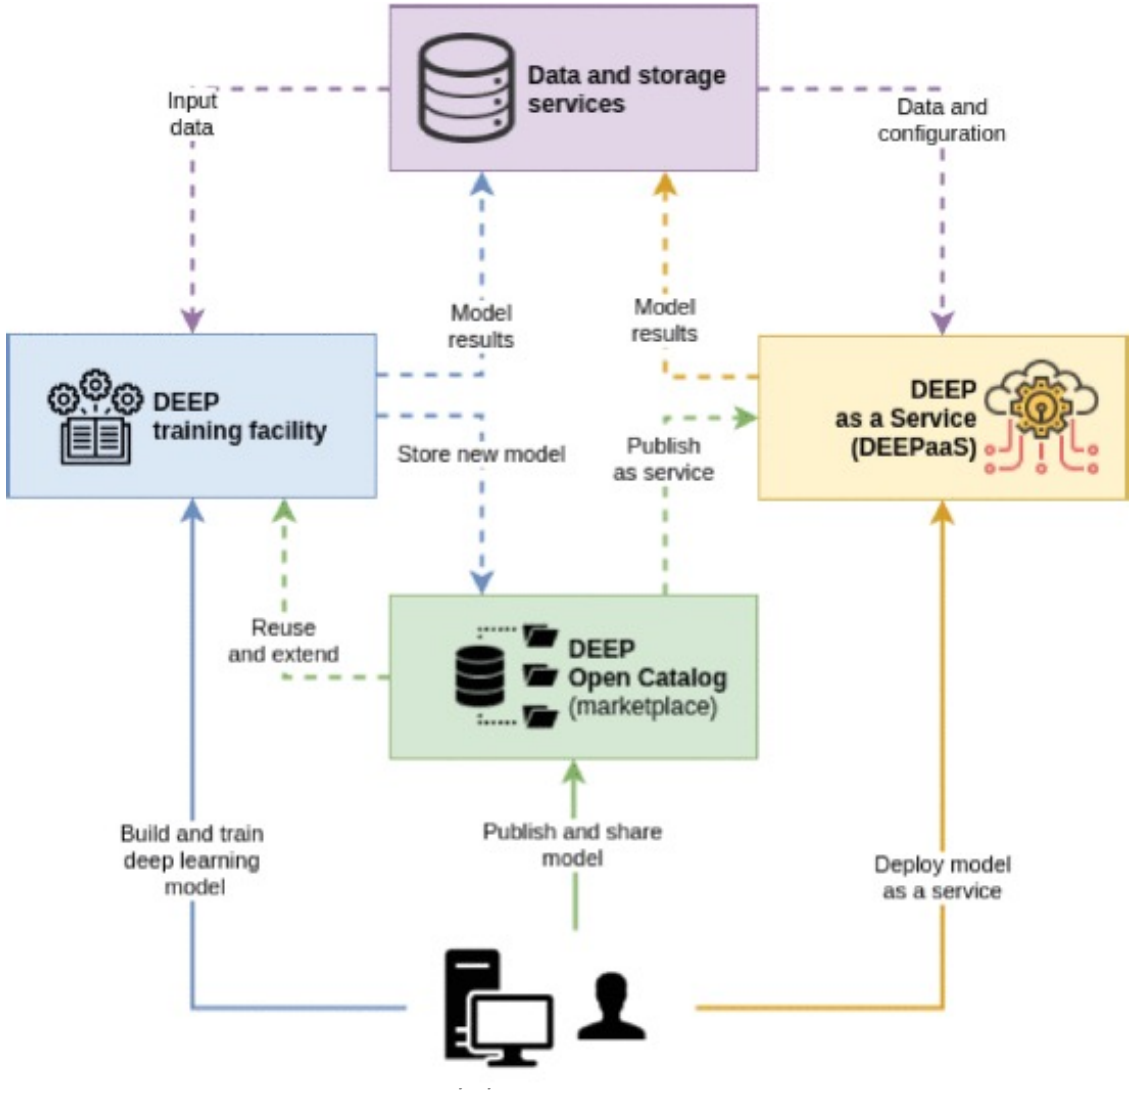

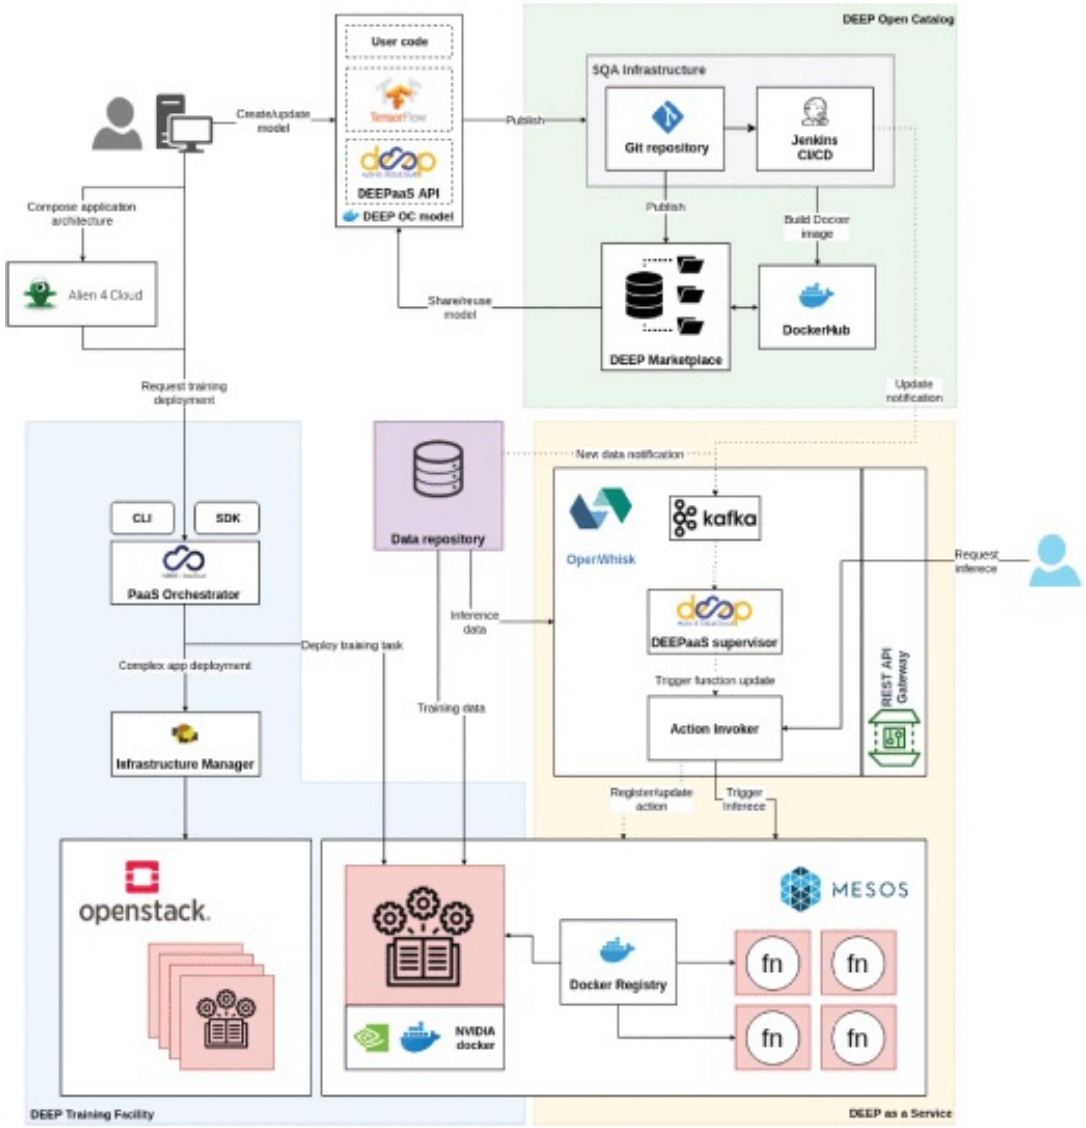

Some more things in the works

- Beside direct ML-related solutions, we are **introducing multiple new services** that should facilitate the use of the INFN distributed resources. Some of them are already used in INFN Cloud by various communities.
  - HTCondor as a Service.
  - Private Docker registry.
  - MinIO as a Service supporting both S3 and POSIX-like access.
  - Deployment of resources over dynamic private networks.
  - PaaS-level management of storage QoS levels, and integration of QoS requests into the PaaS orchestrator.
  - Flexible definition of IP ports for service configurations.
  - Selection of multiple GPU models.
  - Introduction of a Gitlab-as-a-Service solution.
  - Automated pipelines to test service templates and check for potential problems, such as vulnerabilities or configuration errors.
  - Definition of user-specific dashboards focusing on high-level managed services.
  - Definition of multi-level dashboards, directly interfacing for instance to Kubeapps.

Conclusions

- **The evolution of "e-infrastructures"** means on concrete terms the evolution of both the low-level hardware and the creation of dynamic, co-developed solutions to exploit state-of-the-art technology.
- **Scalable federation of resources** is key, there is no "one size fits all" solution. ML is no exception: both hardware and the solution portfolio should be considered jointly; service composition and teamwork are essential to come to optimal results.
- **INFN has been recently revisiting its computing organization** to support this vision. INFN Cloud is fully aligned to this and has already shown multiple results, with an ambitious roadmap for its evolution, covering ML and several other areas. The main drivers are and will be users and their needs.
- **The PNRR offers unprecedented opportunities** to exploit INFN know-how and to sustain the needs of both INFN and Italian scientific communities, through the creation of a nation-wide scientific datalake.