

INFN Cloud, an easy to use, distributed, user-centric Cloud infrastructure and solutions toolbox

Davide Salomoni (davide@infn.it) & Doina Cristina Duma (aiftim@infn.it),
for the INFN Cloud Team (cloud-info@infn.it)

EGI Conference 2020, #EGI2020

INFN (National Institute for Nuclear Physics)

www.infn.it

- A long tradition in **state-of-the-art distributed IT technologies**, 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**.
- For the past 10 years, this mainly meant supporting the experiments @ CERN (LHC), although the scope is now widening very quickly to other communities.
- Currently, INFN operates:
 - 9 medium size centers (Tier-2s in the LHC Computing Grid lingo)
 - 1 large Tier-1 center, at CNAF (Bologna) – certified ISO-27001
- All the INFN centers are connected with 10-100 Gbit/s dedicated connections through the GARR network.
- Collectively, our main centers have about 65,000 CPU cores, 50PB of enterprise-level disk space, 60PB of tape storage.

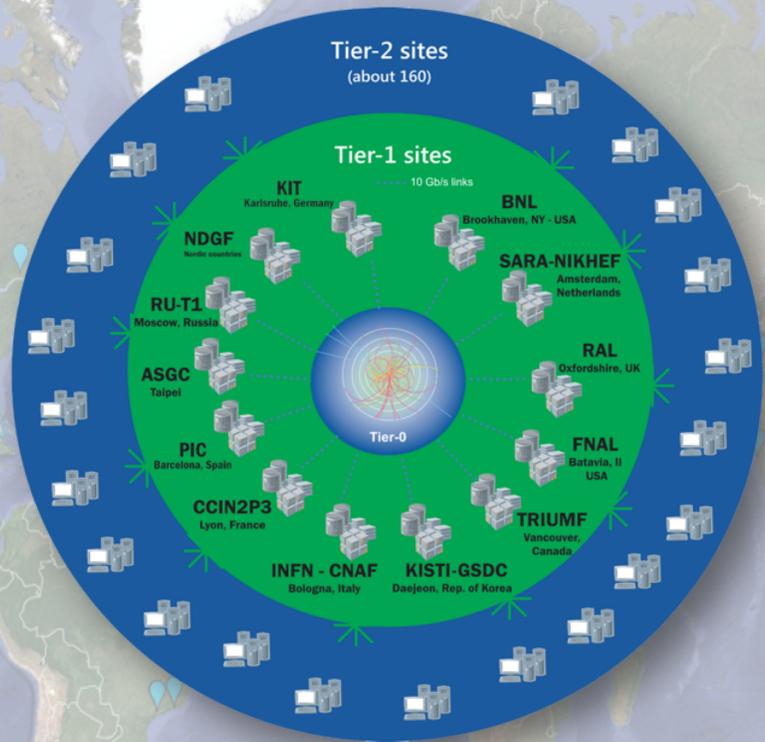


Istituto Nazionale di Fisica Nucleare



Where are we coming from?

HEP computing embraced a large scale distributed model since early 2000s
Based on grid technologies, federating national and international grid initiatives



167 sites, 42 countries

~1M CPU cores

~1 EB of storage

> 2 million jobs/day

10-100 Gb links

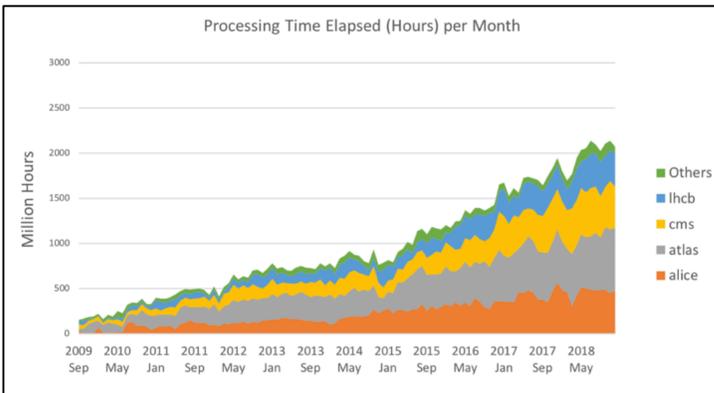
1 TB ~ 10-100 CHF

1 core ~ 100 CHF

HW lifetime: 3-5 years

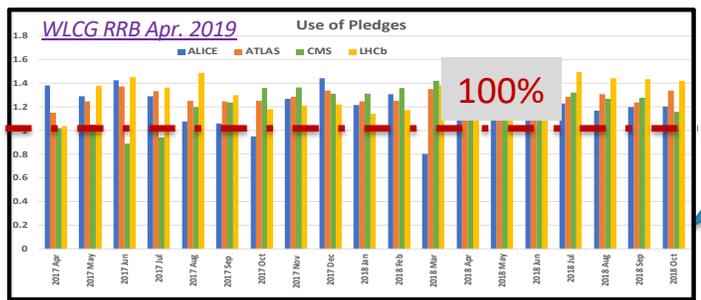
WLCG: an International collaboration to distribute and analyse LHC data

Integrates computer centres worldwide that provide computing and storage resource into a single infrastructure accessible by all LHC physicists.



The data explosion

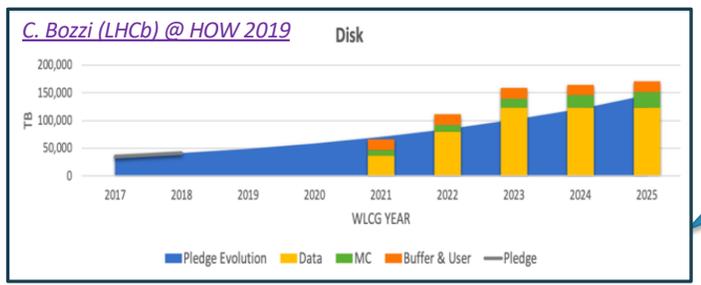
2018



Today LHC generally gets the **requested** computing resources. Extra **opportunistic** compute capacity is available

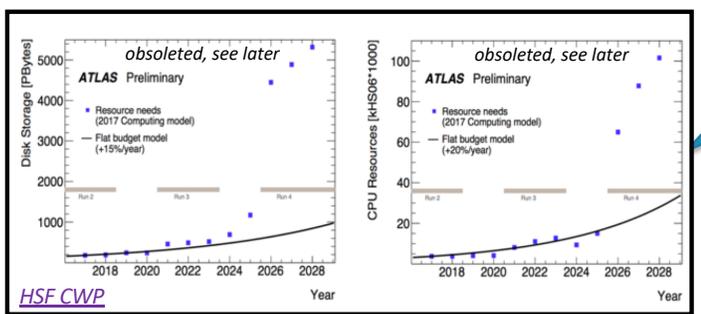
The LHC data volume challenge

Run-3



ALICE and LHCb will increase considerably the data rates in 2022 (LHC Run-3)

HL-LHC



ATLAS and CMS will increase the event rates by a factor 10 and the event complexity in Run-4 (HL-LHC)

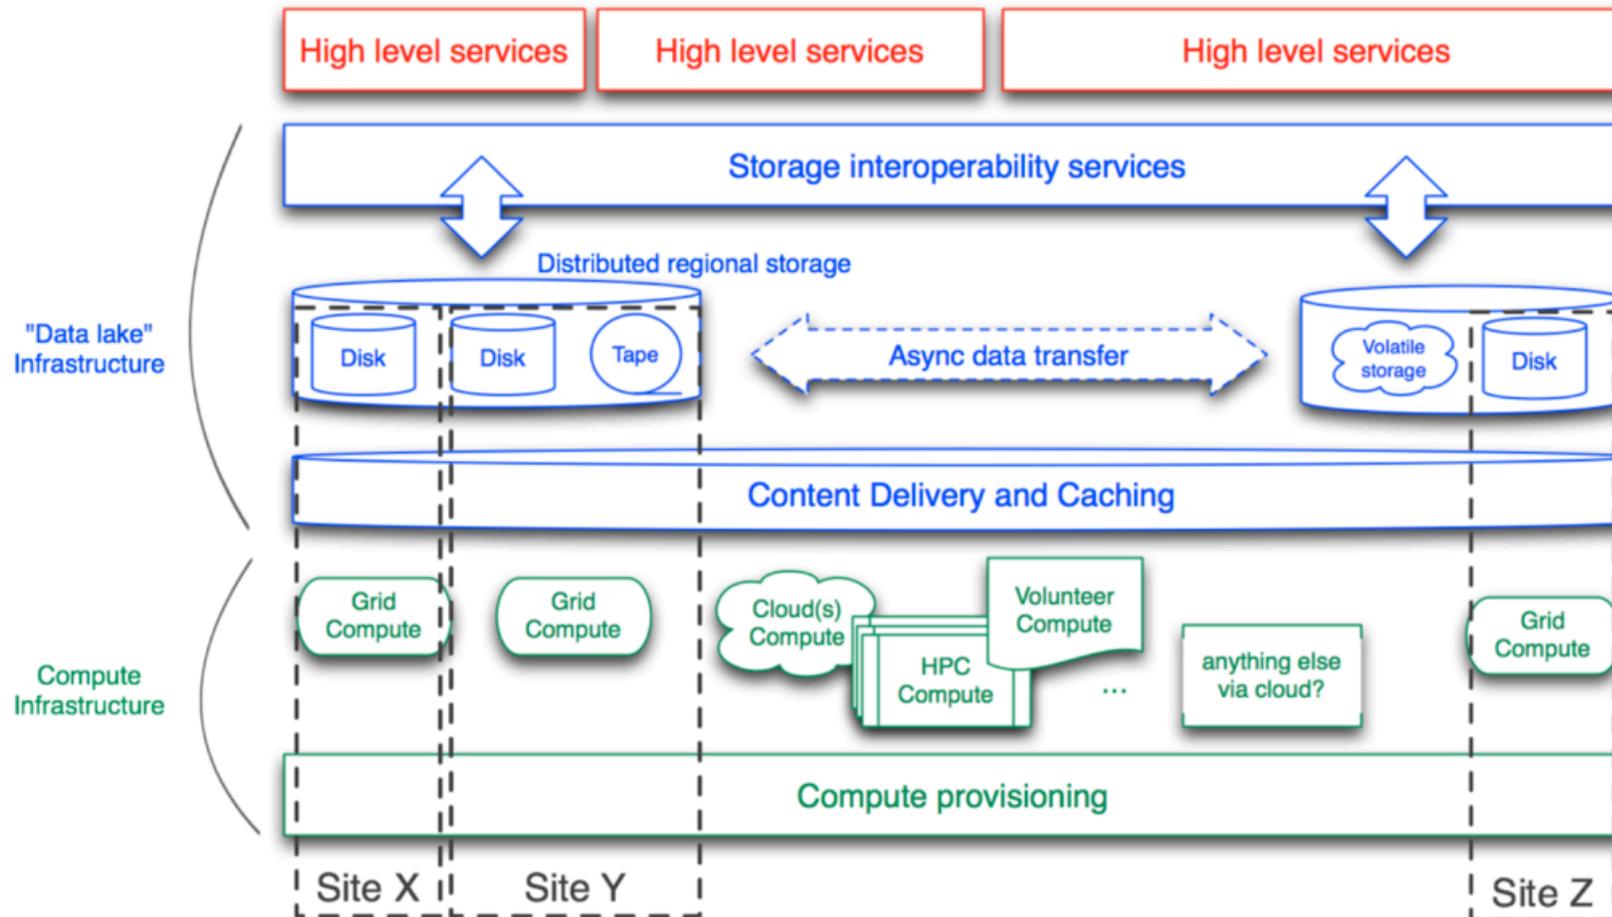
No expected increase of funding for computing: "flat budget"





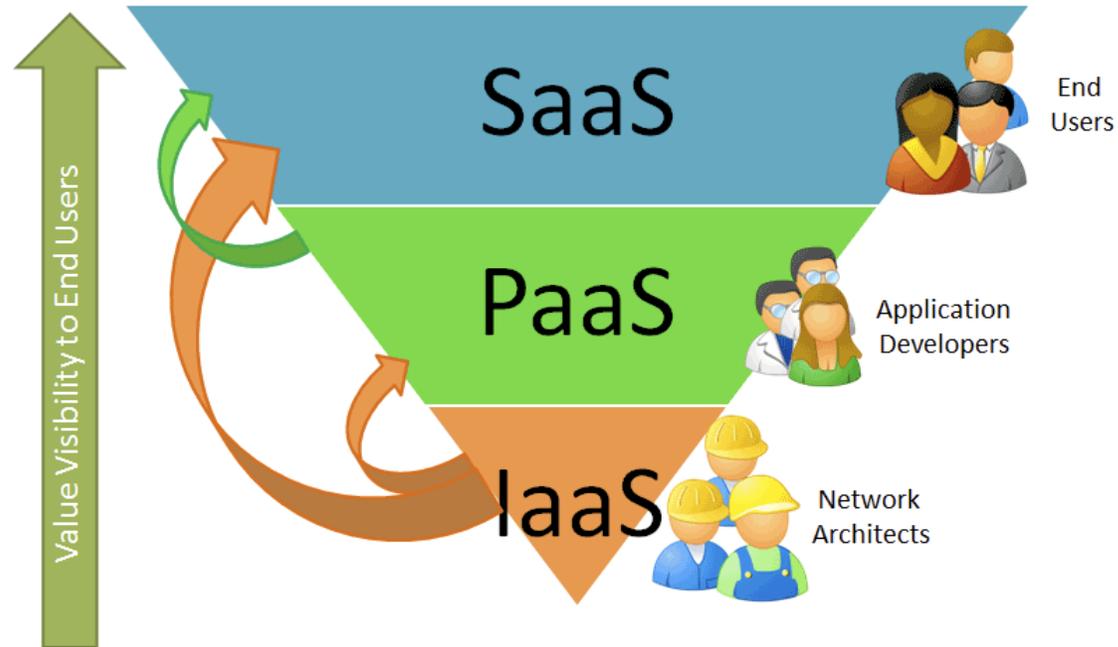
A global technological approach: Data Lake federations

An overarching theme: Service Composability





From a different perspective



In summary

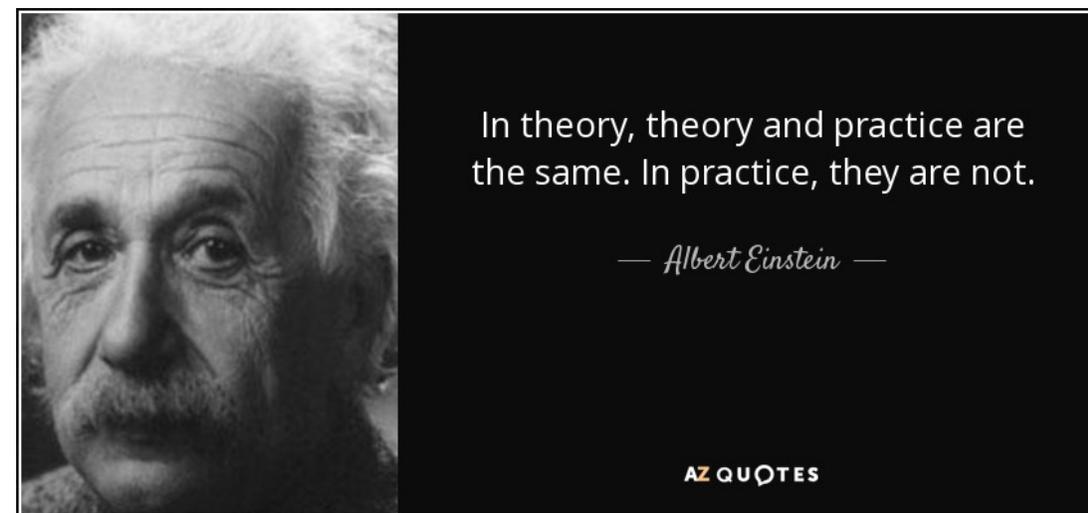
- Use hardware where it is available, possibly also in an opportunistic way, exploiting infrastructures, agreements, projects, as well as regional, national or international grants.
- Besides hardware, what matters most is *know-how* and *availability of high-level solutions*, i.e. toward the applications (“from the PaaS up”).

Right. But how?



How about a typical scientific workflow?

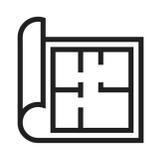
- In a **naïve set of assumptions**, I *have*:
 - A data set somewhere, that I want to analyze.
 - Some algorithms that I want to apply to this data.
 - Some software that can run these algorithms.
 - Some computing resources that can run this software and produce some output.
 - Some space where I can store the output.
- I *assemble everything together* and off I am.





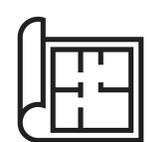
How do we go about this?

- **High-level Objective #1: build added value services on top of IaaS and PaaS infrastructures**
 - Note that, due to the nature of many scientific endeavors (but also public services and commercial industries), these infrastructures may often be *hybrid* (public + private).
- **High-level Objective #2: lower the entry barrier for non-skilled scientists**
 - Transparent (“ZeroOps”) execution on e-Infrastructures.
 - Offer ready-to-use modules, components or services through a catalog or a configurable marketplace.
 - Enable flexible service composition.
 - Implement common software development techniques also for scientists’ applications (DevOps).



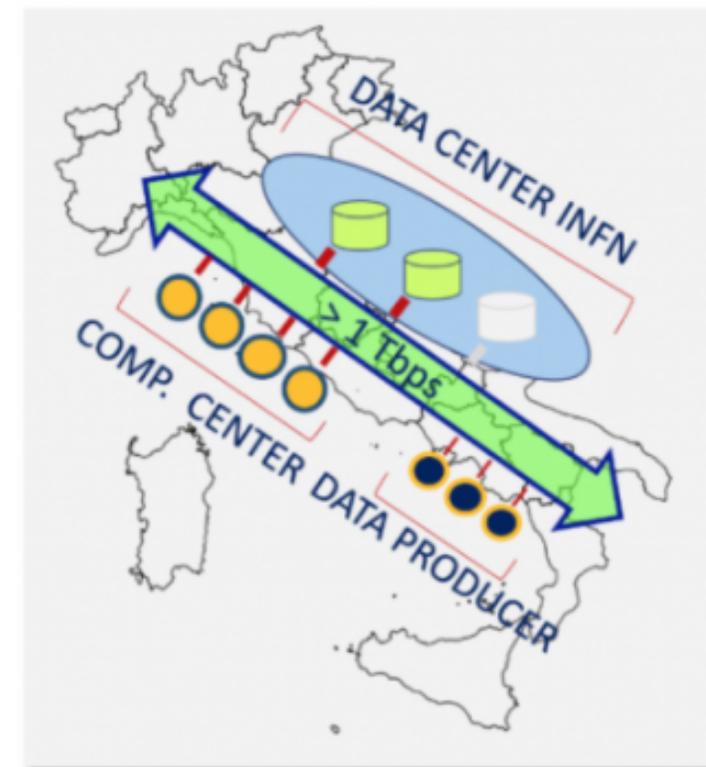
Architectural foundations

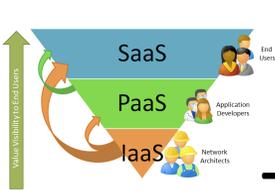
- 1 • **Open source, vendor neutral architecture**, leveraging many years of INFN leadership, investments and know-how in e-infrastructures and distributed computing projects, implementing **extensible service composition**.
- 2 • **Federation of existing Cloud infrastructures** for both compute and data.
- 3 • **Consistent authentication and authorization technologies and policies** at all Cloud levels (IaaS, PaaS, SaaS) via OAuth and OpenID-Connect, supporting also legacy AAI solutions, via **INDIGO-IAM**.
- 4 • **Dynamic orchestration of resources via the INDIGO PaaS Orchestrator** across all participating Cloud infrastructures, according to agreed SLAs and Rules of Participation.



The INFN Cloud architecture

- **An INFN Cloud backbone** spanning the two main INFN computing sites (CNAF and Bari).
 - In each of these two sites there is an “INFN Cloud backbone infrastructure”, connected at high speed with the other.
 - The backbone is used to host the INFN Cloud core services, such as the PaaS core, the internal DNS, the logging and monitoring services, as well as user services that leverage backbone features, such as automated replication of object storage data across the two sites.
- **A set of distributed, federated cloud infrastructures** connecting to the backbone. Currently, the cloud infrastructures at CNAF and Bari (which are *not* the corresponding backbone infrastructures) are already connected to the INFN Cloud backbone, with several other INFN sites in the pipeline.





The INFN Cloud services

- The INFN Cloud services are based on **modular components and span the IaaS, PaaS and SaaS models** for both computing and data.
- All services are described by [TOSCA templates](#) (which can refer internally to other components such as Ansible playbooks, HELM charts, etc.).
- The services can be **deployed** via the INFN Cloud Dashboard or via a command line interface:
 - **Automatically** by the INFN Cloud Orchestrator on one of the federated Cloud infrastructures, depending on resource availability and policies.
 - **Manually** by a user on a specific federated Cloud infrastructure.



Which services?

- They **include**, for example:
 - Creation of VMs with different flavors and sizes.
 - Creation of containers (specify the container name) or of applications via docker-compose files.
 - Building blocks “as a service” for example for container orchestration (e.g. creation of a Mesos cluster or of a Kubernetes cluster as a service).
 - Pre-configured environments for example for data analytics (e.g. using ElasticSearch and Kibana or Spark).
 - Nonvolatile, object storage and Posix-compliant virtual file system solutions transparently connected to higher-layer services (e.g Jupyter notebooks as a service with permanent, replicated storage).
 - Dynamic clusters tailored to specific experiments (e.g. an automated full HTCondor installation realized on a k8s cluster, or a GPU-based Machine Learning-optimized environment).
 - Services leveraging transparent user-level encryption of disk volumes.
- **The service portfolio can be easily extended** with the simple addition / modification of TOSCA templates.



The INFN Cloud Dashboard



Welcome to **inf-n-cc**

Sign in with



Username

Password

Sign in

[Forgot your password?](#)

Or sign in with



Not a member?

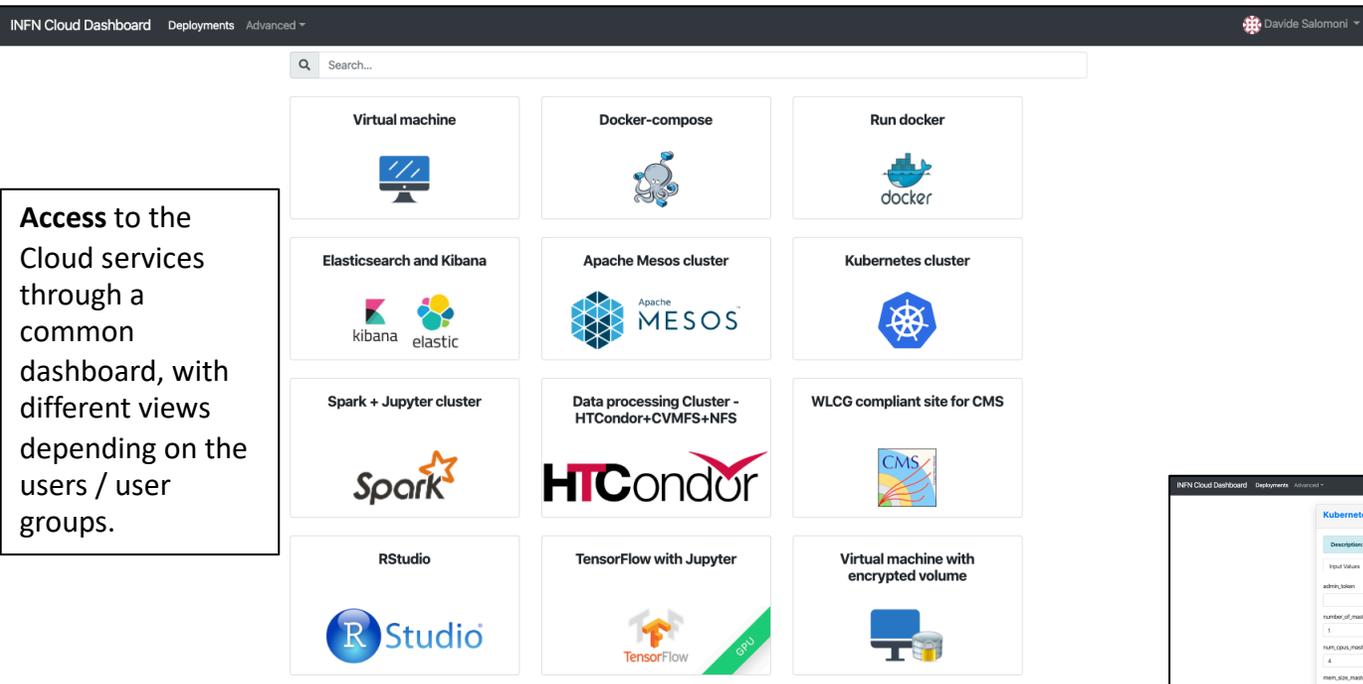
[Register a new account](#)

Authentication can be enabled for:

- Local username/password
- Google accounts
- EduGAIN (e.g. University, research centers, etc.)
- Other OIDC providers

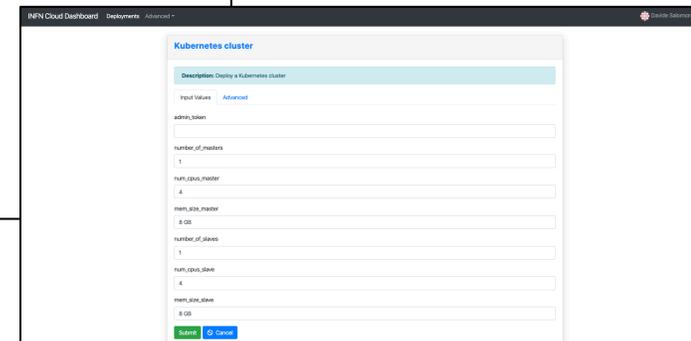
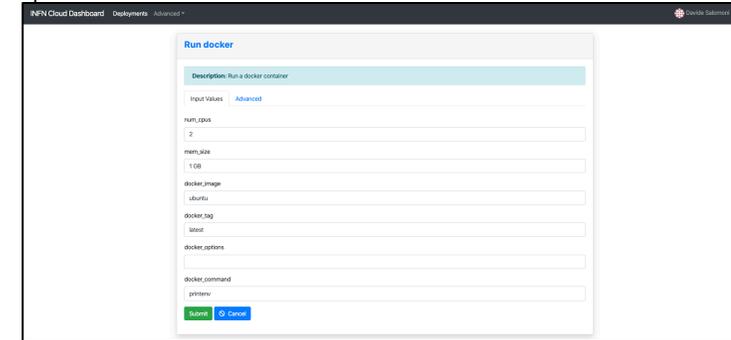


Transparent, multi-site **federation** for users of Cloud resources belonging to INFN and/or to other Cloud providers (private or public)



Access to the Cloud services through a common dashboard, with different views depending on the users / user groups.

Composed, high-level services easily customizable and configurable directly by users





The INFN Cloud Organization

- INFN Cloud is internally organized into **5 Work Packages**, run by people belonging to several INFN sites in a fully distributed way:
 - WP1: Architecture, Operations and Service Portfolio
 - WP2: Documentation, User Support, Communication and Training
 - WP3: Monitoring and Accounting
 - WP4: Security, Policies and Rules of Participation
 - WP5: Service Evolution and New Developments
- The **current operational state** of INFN Cloud is pre-production but already serving several INFN experiments and collaborations, with full production state and general availability planned by beginning 2021, in sync with already planned resource expansion and connection of additional federated Clouds.



The INFN Cloud Rules of Participation

- The INFN Cloud **Rules of Participation (RoP)** document describes the general requirements that must be satisfied by resource providers in order to join the INFN Cloud federation. It is complemented by the more specific “Resource Center Operation Level Agreement (OLA)” and “User Community Operation Level Agreement (OLA)” documents.
- The RoP document includes the following topics:
 - Why joining the INFN Cloud federation?
 - Compliance for resource access
 - Authentication and Authorization
 - Resource allocation
 - Resource configuration
 - Supported end users' services
 - Networking
 - Support
 - Service level targets
 - Security
 - Management of security incidents
 - Traceability and logging
 - Accounting
 - Certification
 - Withdrawals
 - Violations



Next steps

- INFN Cloud will announce its general availability for INFN users at the beginning of 2021.
- In 2021, we aim at opening INFN Cloud with interested academic and research infrastructures, as well as the connection with some commercial providers.
- For any further information, expressions of interest or other queries, send an email to:

cloud-info@lists.infn.it

INFN Cloud

- Integration, not replacement
- Start from the needs, not from the tools
- Focus on high-level added value services, not on “infrastructures”
- Create an expandable service model
- Develop a sustainable organizational model
- Final remarks:
 - The goal is to provide end-users with compute and storage services. Possibly as part of a “complex” computing environment, by offering
 - **a portfolio of technical solutions** already developed but extensible – continuously evolving following a user driven development approach
 - **technical support for the end user applications migration** to a cloud-based environment
 - **transparent solutions hiding the resources allocation complexity** in a federation of distributed clouds

INFN Cloud – during #EGI2020

- Monday, 2 Nov.
 - **Clinic: Authentication-Authorisation services**

INDIGO IAM and VOMS

Andrea Ceccanti

Room: <http://go.egi.eu/zoom2>

13:25 - 13:35

- Tuesday, 3 Nov.
 - **Data Management Solutions**

Integrated, heterogeneous data access in INFN-Cloud and beyond

Stefano Stalio

Room: <http://go.egi.eu/zoom2>

11:55 - 12:20

- **Cloud Computing - Part I**

The INFN-Cloud PaaS Dashboard

Marica Antonacci

Room: <http://go.egi.eu/zoom4>

13:30 - 13:45

- Wednesday, 4. Nov.
 - **Cloud computing - Part 2**

Monitoring and accounting in the INFN-Cloud infrastructure

Vincenzo Spinoso

Room: <http://go.egi.eu/zoom4>

14:00 - 14:15



Acknowledgments

- This presentation, like INFN Cloud in general, has **many explicit and implicit authors**.
- INFN Cloud currently involves about 30 people, located in 10 INFN sites. We **thank** all of them, the INFN users that throughout these months have contributed to refine the INFN Cloud portfolio, as well as INFN management for their continued support.

And thank you for your attention!