



LEVERAGING TOSCA ORCHESTRATION TO ENABLE FULLY AUTOMATED CLOUD-BASED RESEARCH ENVIRONMENTS ON FEDERATED HETEROGENEOUS E-INFRASTRUCTURES

Marica Antonacci (INFN) on
behalf of the INFN Cloud team

International Symposium on Grids & Clouds (ISGC) 2023



CLOUD COMPUTING: A GAME-CHANGER IN SCIENTIFIC RESEARCH

Cloud computing has brought about a significant shift in scientific research, providing researchers with new tools and capabilities to accelerate their work.



OBSTACLES AND CHALLENGES

While cloud computing offers many benefits for scientific research, there are several obstacles to its broader adoption.

Some of the key obstacles include:

INTEGRATION CHALLENGES

The lack of integration of existing infrastructures can make it difficult to connect and manage resources across different cloud environments.

TECHNICAL COMPLEXITY


Cloud computing requires expertise in areas such as networking, security, and data management, which can be challenging for researchers to master.

CULTURAL BARRIERS

Some researchers may be resistant to adopting cloud computing because they are comfortable with existing infrastructures and methods.

INFN'S CONTRIBUTION TO COMPUTING TECHNOLOGIES IN SCIENTIFIC RESEARCH

“preparing the GRID”



“preparing the Cloud”



“expanding beyond HEP”




A long tradition in state-of-the-art distributed IT technologies and solutions, from the first small clusters to Grid and Cloud-based computing.

The main focus of INFN is not on computing in itself, but rather on utilizing computing as a crucial means to support its research and overall mission.

INFN CLOUD

Cloud
Resources
for research



<https://www.cloud.infn.it>

- In production since March 2021
- Based on scalable and production grade solutions
- A customisable portfolio of services accessible through various interfaces (web, terminal, API)

ARCHITECTURAL FOUNDATIONS




NO VENDOR LOCK-IN

Open-source,
vendor-neutral
architecture



FEDERATION

of existing Cloud
infrastructures for
both compute
and data



DYNAMIC ORCHESTRATION

of resources via
the INDIGO PaaS
Orchestrator



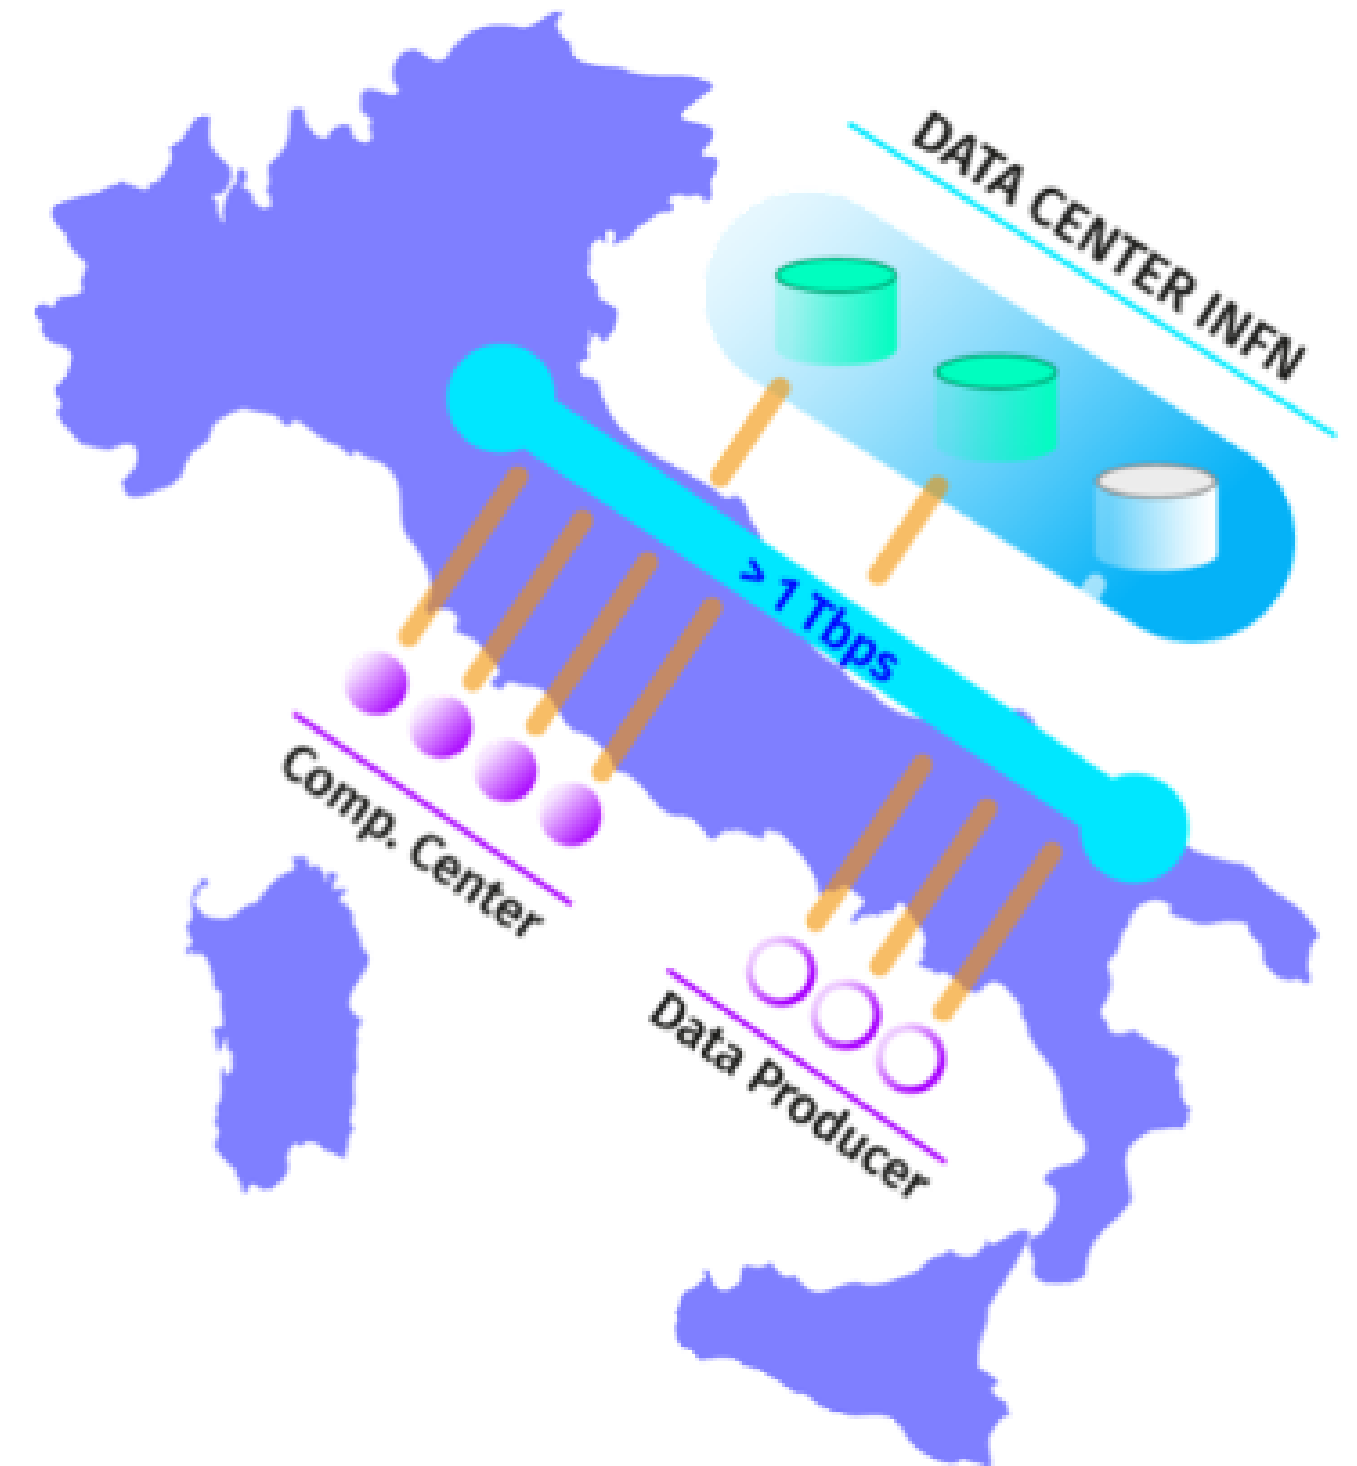
CONSISTENT AUTHN/AUTHZ

at all cloud levels
via OpenID-
Connect/OAuth2

INFN CLOUD IS DESIGNED AS A FEDERATION OF PRE-EXISTING INFRASTRUCTURES

- The Backbone of the INFN Cloud is made up of two closely linked federated sites, BARI and CNAF.
- A scalable set of satellite sites, geographically distributed across Italy and loosely coupled, expand the resources offered by the backbone.

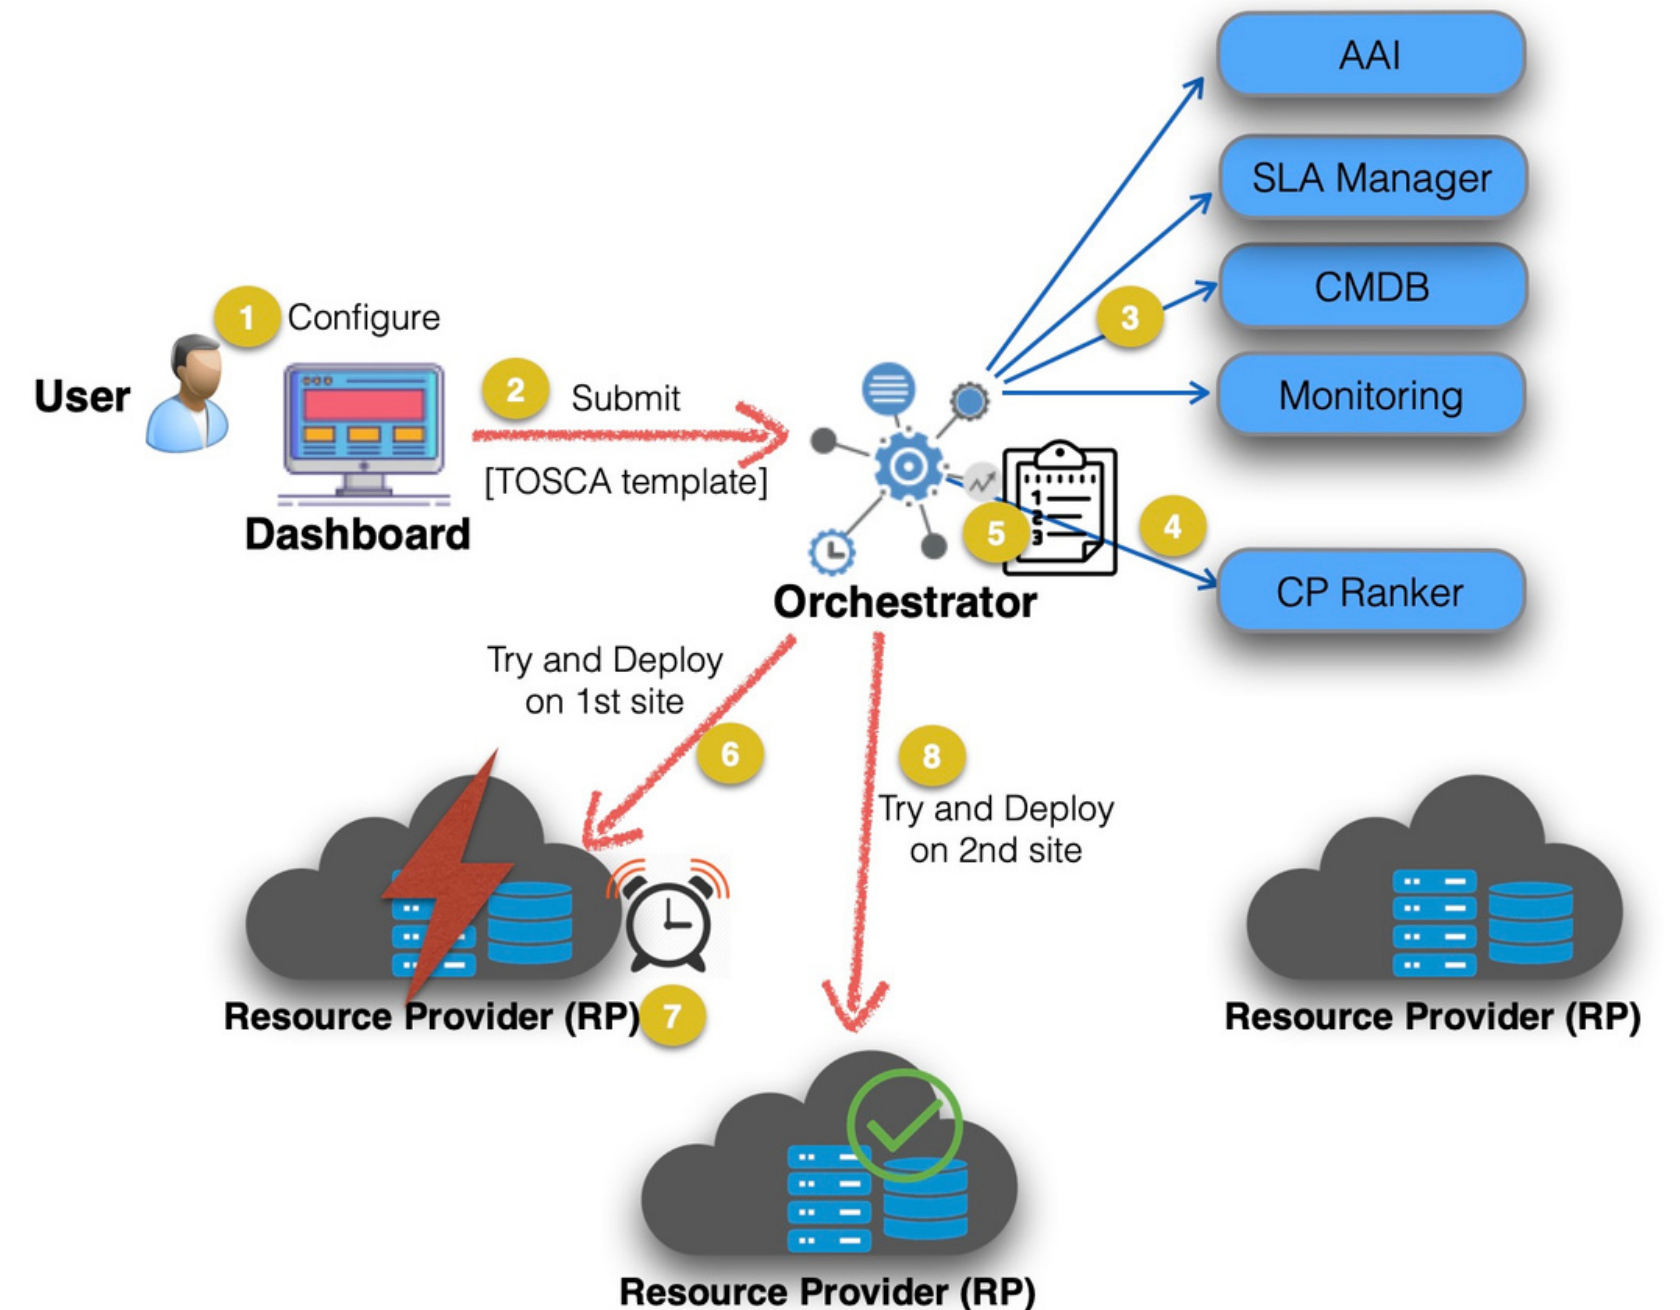
INFN Cloud core services and some centralised, fully managed, high-level services are hosted on the Backbone. This allows us to leverage high-availability and disaster recovery capabilities to ensure that these critical services are always available and operating at peak efficiency.



THE FEDERATION MIDDLEWARE

The INDIGO PaaS Orchestrator enables the federation of distributed and heterogeneous compute environments: clouds, docker orchestration platforms, HPC systems.

- Smart scheduling → Automatic selection of the best provider
 - based on compute/storage requirements vs provider capabilities including the following criteria:
 - Resource quotas (SLA)
 - Monitoring data
 - Support for specialized hardware (GPU, Infiniband)
 - Data location
- Support for hybrid deployments and network orchestration
- Client interfaces for advanced users (REST APIs, CLI, python bindings) and end-users (web dashboard - no skills required)

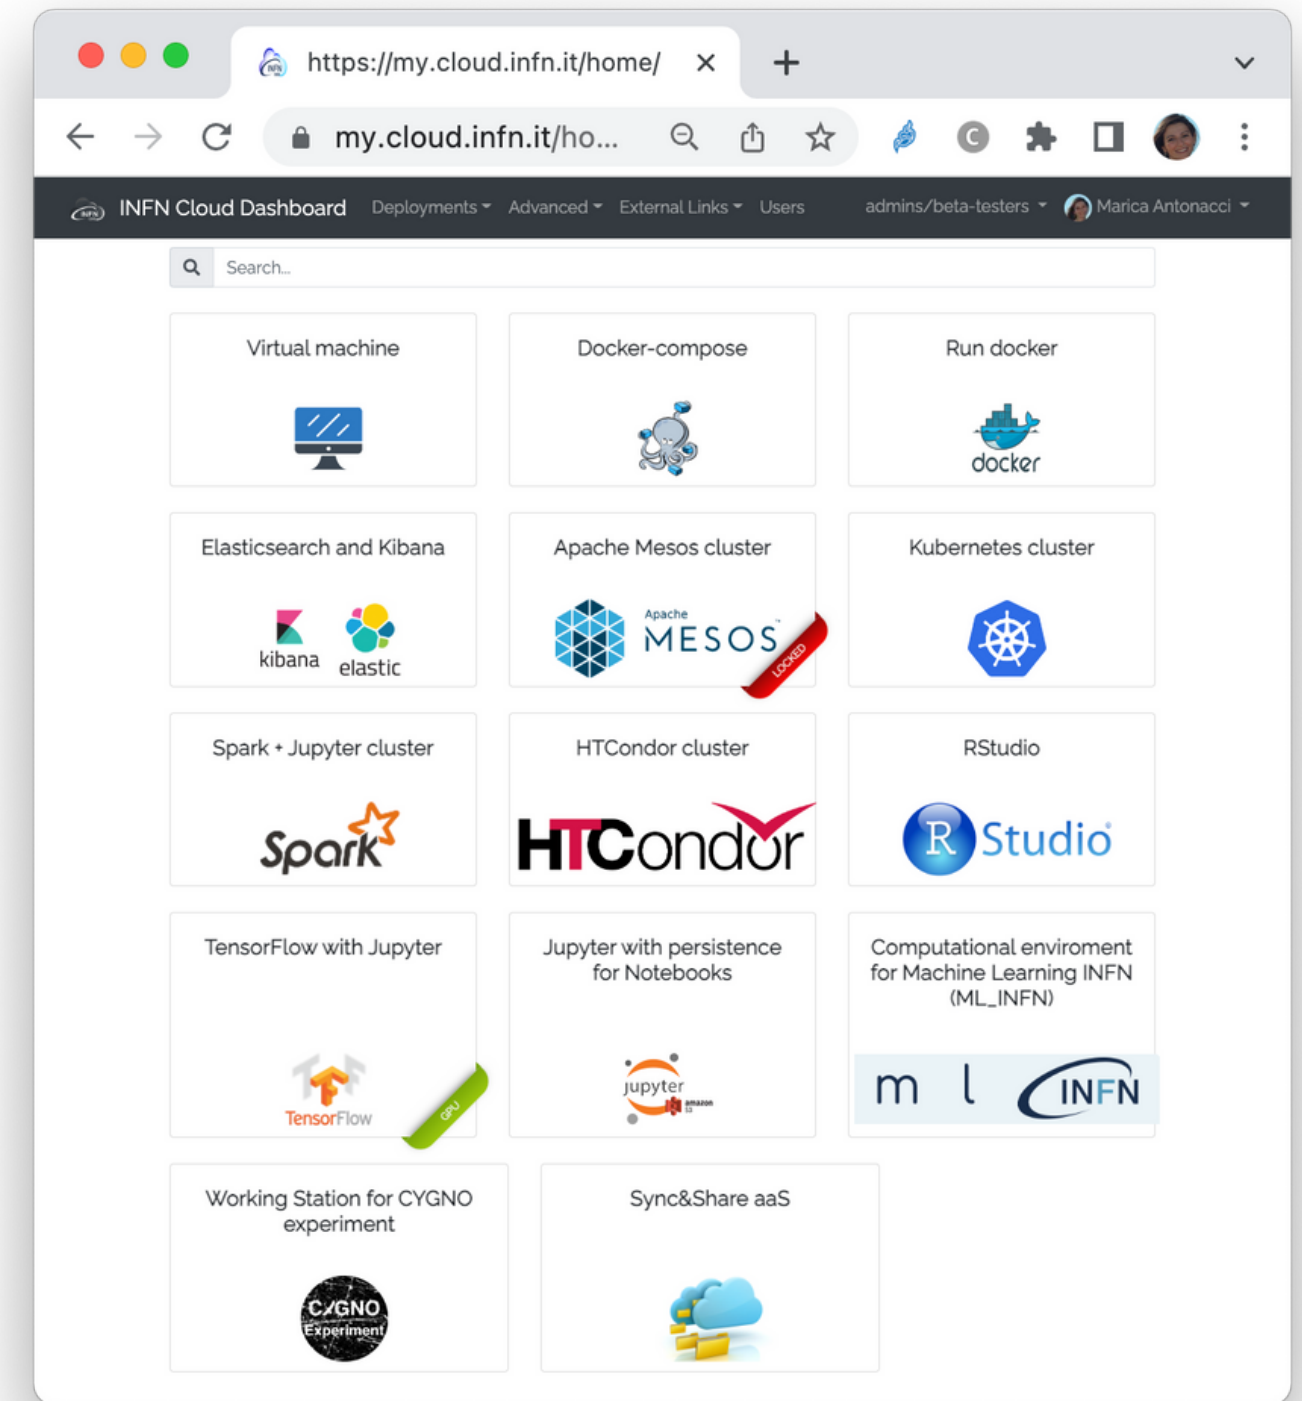


THE PAAS DASHBOARD

The INDIGO PaaS Dashboard is a web-based user interface that enables users to manage and monitor their deployments without requiring any TOSCA knowledge.

The dashboard hides all technical details and provides an intuitive interface for managing service deployments.

- OpenID-Connect Authentication
- Multi-tenancy
- Secrets management (via Vault integration)
- Dynamic view of service catalog (depending on the user group membership)



REQUEST SERVICES WITH JUST A FEW CLICKS

Kubernetes cluster

Description: Deploy a single master Kubernetes 1.23.8 cluster

Deployment description

Configuration **Advanced**

admin_token
.....

Password token for accessing K8s dashboard

number_of_nodes
3

Number of K8s node VMs

ports

Ports to open on the K8s master VM

master_flavor
--Select--

Number of vCPUs and memory size of the K8s master VM

node_flavor
--Select--

Number of vCPUs and Memory Size of each K8s node VM

Customize your deployment through the deployment input parameters

Choose the Scheduling strategy

- automatic: let the Orchestrator select the best provider
- manual: choose the provider from the drop down menu automatically created by the Dashboard with the list of providers returned by the SLA Manager service

Kubernetes cluster

Description: Deploy a single master Kubernetes 1.23.8 cluster

Deployment description

Configuration **Advanced**

Configure scheduling:
 Auto Manual

Select a provider:
RECAS-BARI: org.openstack.nova

Set deployment creation timeout (minutes) 720

Do not delete the deployment in case of failure

Send a confirmation email when complete

THE SERVICE IMPLEMENTATION STRATEGY

The employed strategy is based on the Infrastructure as Code paradigm.

Users describe "What" is needed rather than "How" a specific service or functionality should be implemented.

The adopted technologies enable a Lego-like approach: services can be composed and modules reused to create the desired infrastructure.

The logo for OASIS TOSCA, featuring the word "OASIS" in purple, a small blue icon of a person, and the word "TOSCA" in yellow.

TOSCA is used to model the topology of the whole application stack

The logo for Ansible, featuring a black circle with a white letter 'A' inside, followed by the word "ANSIBLE" in black capital letters.

Ansible is used to automate the configuration of the virtual environments

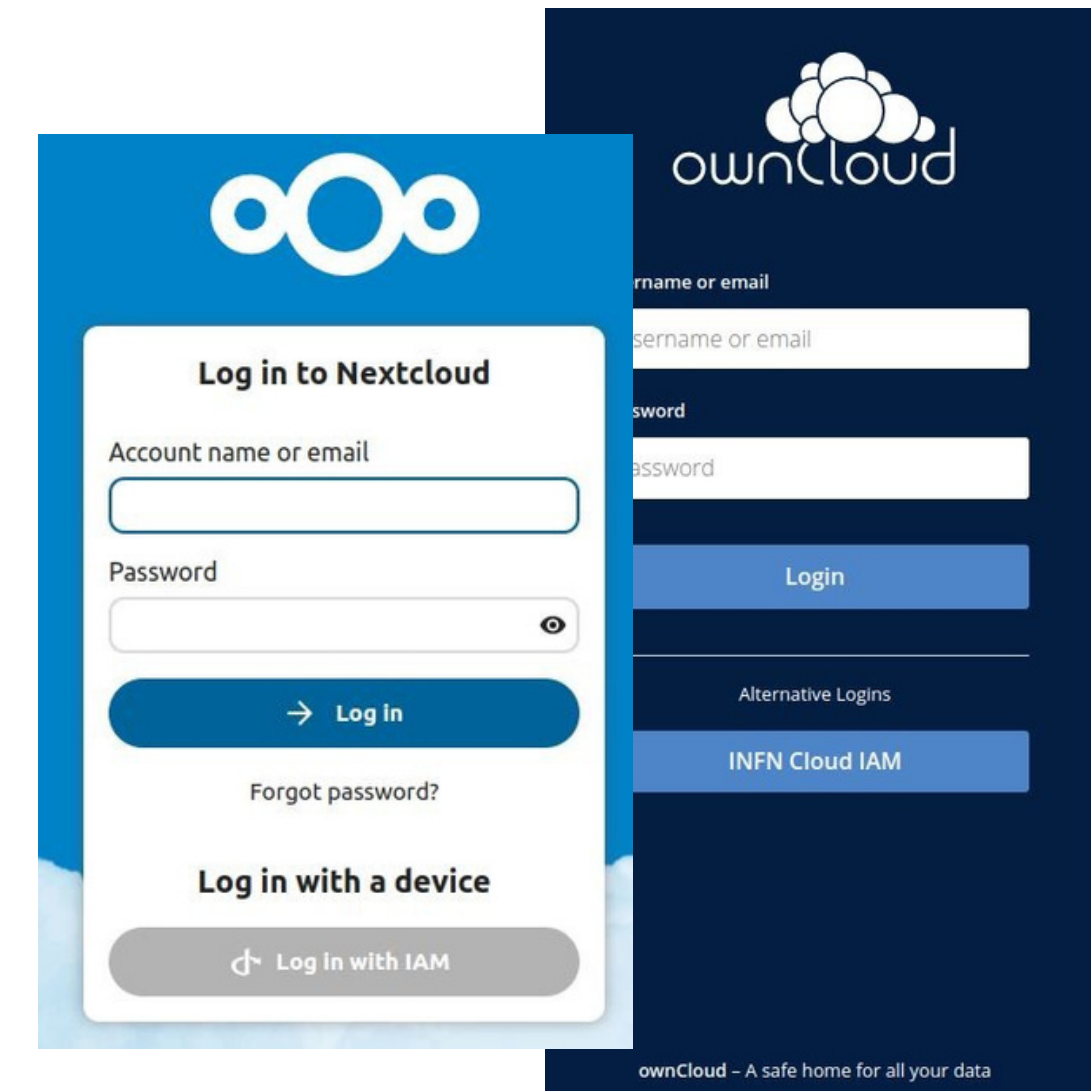
The logo for Docker, featuring a blue whale icon with a stack of blue containers on its back, followed by the word "docker" in a lowercase, dark green font.

Docker is used to encapsulate the high-level application software and runtime

THE SERVICE CATALOGUE

GENERAL PURPOSE SERVICES

- Virtual Machine with or without external block storage, eventually equipped with docker engine and docker-compose, on top of which dockerized services can be automatically started;
- Data analytics and visualization environments based on Elasticsearch and Kibana
- File sync & share solution based on OwnCloud/NextCloud with 1) replicated backend storage on the S3-compliant Object Storage provided by the INFN Cloud infrastructure; 2) automatic configuration for enabling INDIGO IAM OpenID Connect authentication; 3) pre-installed and configured backup cron jobs for safely storing configuration and data on the Object Storage for future restore in case of disaster; 4) integrated application and backup monitoring based on Nagios

The image shows a Nagios monitoring dashboard. At the top, it displays 'Current Network Status' (Last Updated: Sun Feb 26 22:09:27 CET 2023), 'Host Status Totals' (Up: 0, Down: 0, Unreachable: 0, Pending: 7), and 'Service Status Totals' (Ok: 16, Warning: 0, Unknown: 0, Critical: 0, Pending: 0). Below this is a table titled 'Service Status Details For All Hosts'.

Host	Service	Status	Last Check	Duration	Attempt	Status Information
backup	Check last backup	OK	02-26-2023 22:09:21	0d 1h 40m 11s	1/4	OK - Last backup was successful
	Ping	OK	02-26-2023 22:04:21	0d 1h 40m 11s	1/4	PING OK - Packet loss = 0%, RTA = 0.09 ms
db	MySQL DB	OK	02-26-2023 22:09:21	0d 1h 40m 11s	1/4	Uptime: 6610 Threads: 4 Questions: 108019 Slow queries: 0 Opens: 180 Open tables
	Ping	OK	02-26-2023 22:04:21	0d 1h 40m 11s	1/4	PING OK - Packet loss = 0%, RTA = 0.11 ms
dbbackup	Check DB backup	OK	02-26-2023 22:09:21	0d 1h 40m 11s	1/4	OK - 208 files found
	Ping	OK	02-26-2023 22:04:21	0d 1h 40m 11s	1/4	PING OK - Packet loss = 0%, RTA = 0.11 ms
nagios	Current Load	OK	02-26-2023 22:09:21	0d 1h 40m 11s	1/4	OK - load average: 1.17, 1.14, 1.19
	Ping	OK	02-26-2023 22:04:21	0d 1h 40m 11s	1/4	PING OK - Packet loss = 0%, RTA = 0.08 ms
root	Root Partition	OK	02-26-2023 22:09:21	0d 1h 40m 11s	1/4	DISK OK - free space: / 22934 MB (86.22% used=91%)
	Ping	OK	02-26-2023 22:04:21	0d 1h 40m 11s	1/4	PING OK - Packet loss = 0%, RTA = 0.08 ms
owncloud	Owncloud application	OK	02-26-2023 22:09:21	0d 1h 40m 11s	1/4	HTTP OK: HTTP/1.1 302 Found - 1046 bytes in 0.087 second response time
	Ping	OK	02-26-2023 22:04:21	0d 1h 40m 11s	1/4	PING OK - Packet loss = 0%, RTA = 0.09 ms
proxy	S3 bucket usage	OK	02-26-2023 21:03:38	0d 1h 45m 11s+	1/4	S3 OK - 3756879a-b608-11ed-b605-0242ac110002-data: 0 objects, 0m
	Application frontend	OK	02-26-2023 22:09:21	0d 1h 40m 11s	1/4	SSL OK - data.212.189.205.212.mysql.cloud.inf.it - certificate expires in 89 days
redis	Ping	OK	02-26-2023 22:04:21	0d 1h 40m 11s	1/4	PING OK - Packet loss = 0%, RTA = 0.08 ms
	Redis Service	OK	02-26-2023 22:09:21	0d 1h 40m 11s	1/4	PING OK - Packet loss = 0%, RTA = 0.11 ms
						TCP OK - 0.001 second response time on redis port 6379

ON-DEMAND INTERACTIVE DATA ANALYSIS ENVIRONMENTS

Web-based multi-user interactive development environment for notebooks, code and data built on JupyterLab and enhanced with:

- persistent storage areas for storing results and notebooks for future re-use;
- a monitoring system based on Prometheus and Grafana for collecting relevant metrics;
- experiment-specific customizations (pre-installed libraries, drivers, configurations, etc.)

Advanced use-cases:

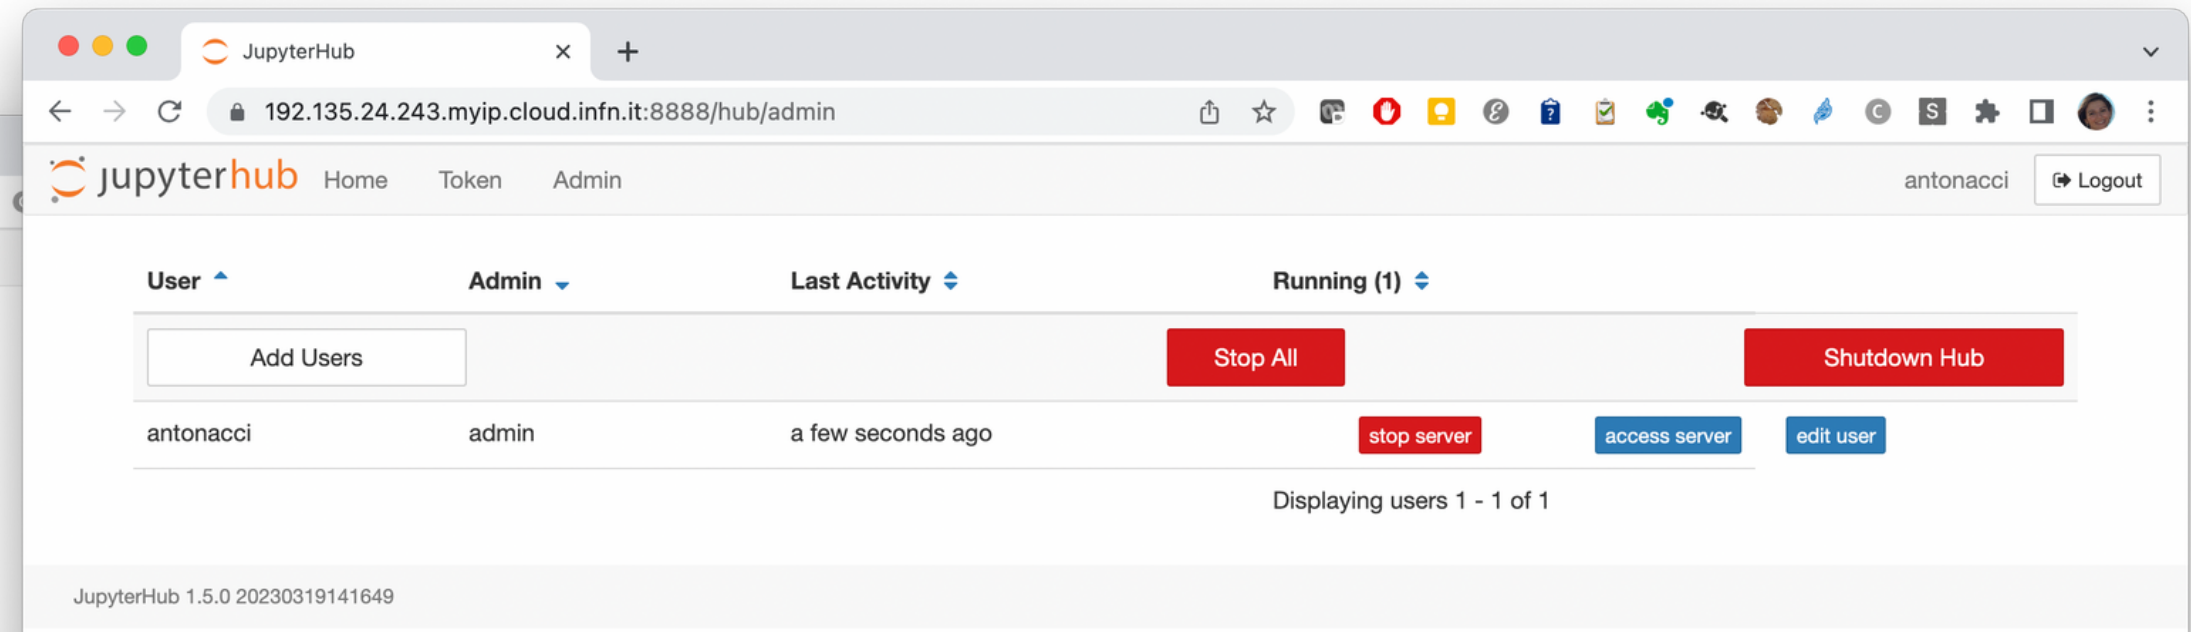
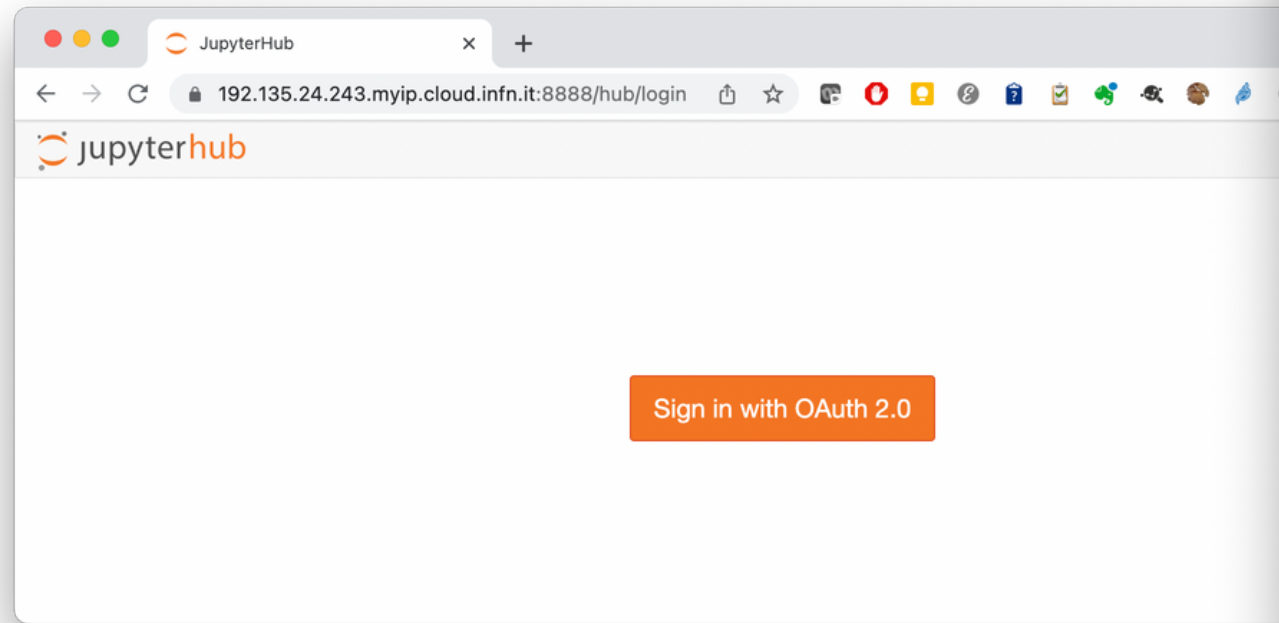
- Cygno experiment for Dark Matter direct detection: Python/ROOT kernels, pre-installed libraries for event reconstruction, data analysis and simulation (based on GEANT4 and Garfield++ software), CVMFS mounts
- ML-INFN Project, an INFN-funded project aiming at lowering the potential barriers for accessing specialized hardware for the exploitation of Machine Learning techniques: in this case the JupyterLab instances are able to access one or more GPUs as the needed drivers and configurations are automatically managed. Moreover, GPU partitioning (based on nvidia MIG feature) is also supported for optimal utilization.



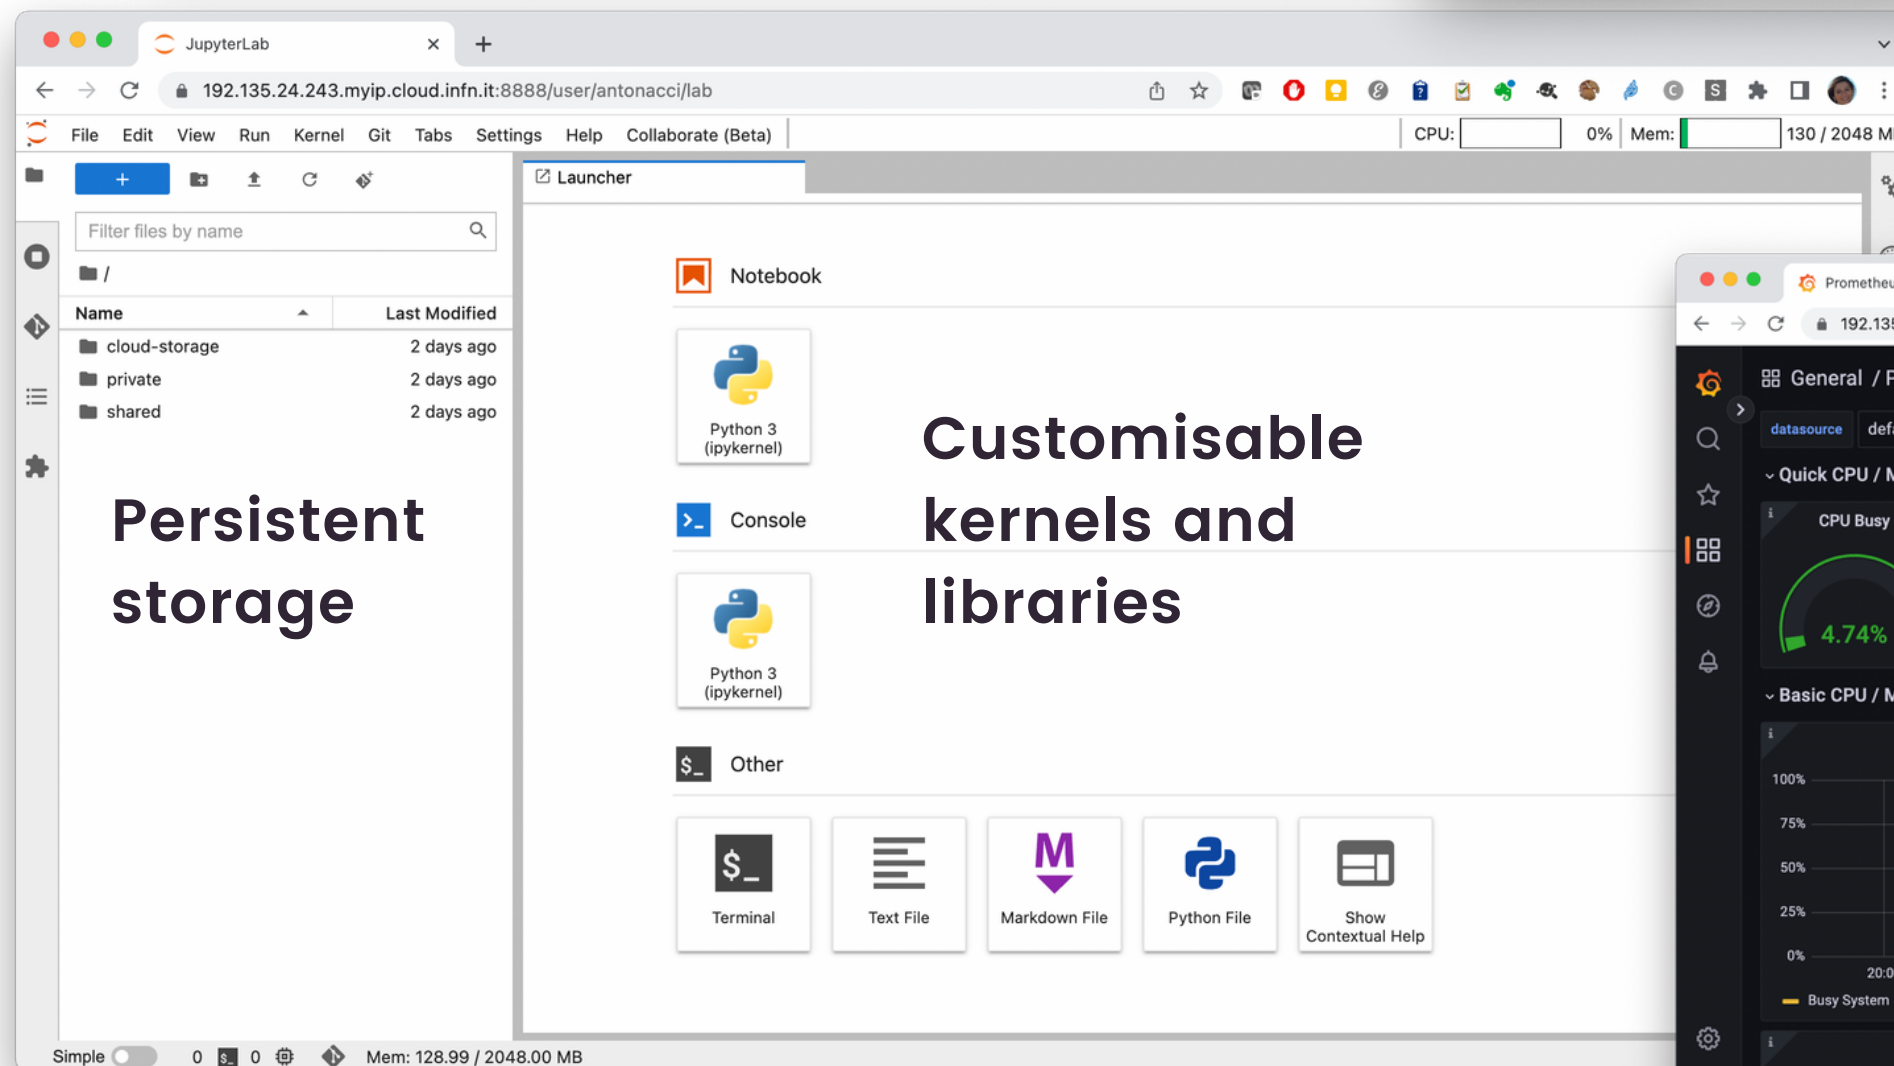
Automatic discovery of special hardware resources

Fully automated installation and configuration of software and drivers

OpenID-Connect AuthN/AuthZ



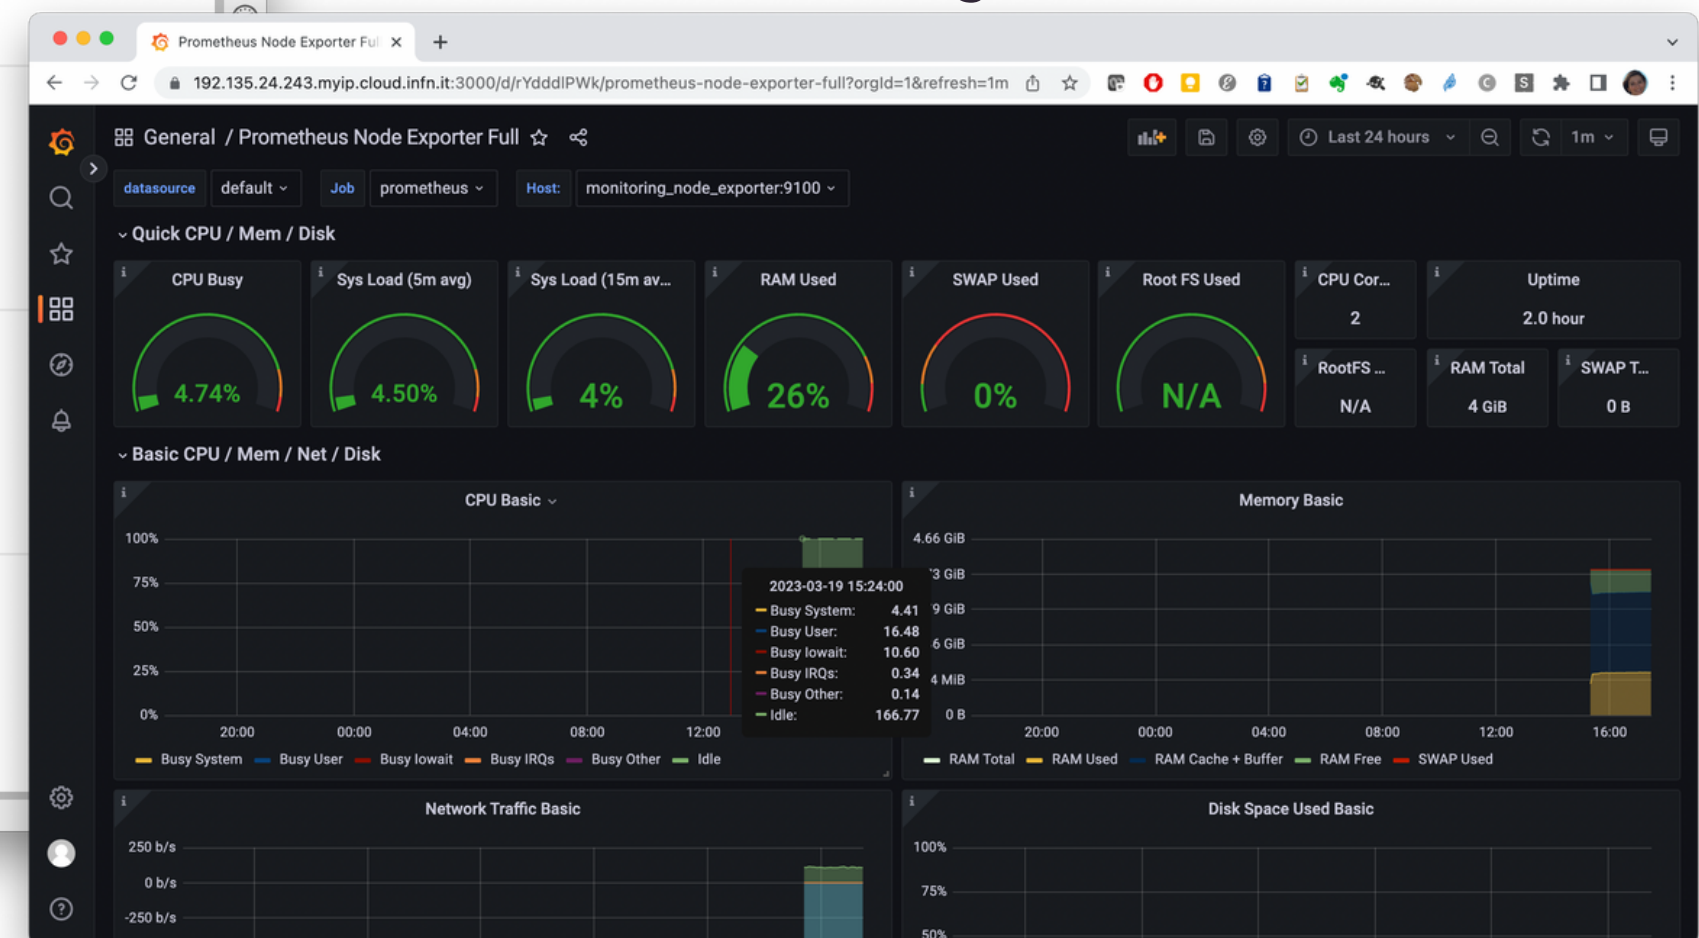
Role based access (admin vs user)



Persistent storage

Customisable kernels and libraries

Monitoring and visualization




THE SERVICE CATALOGUE

KUBERNETES-BASED ADVANCED SERVICES

Running services on Kubernetes provides evident benefits like:

- **Scalability:** Kubernetes can scale applications horizontally by adding or removing containers in response to demand.
- **Resilience:** Kubernetes provides self-healing mechanisms that detect and recover from failures automatically.
- **Declarative management:** Kubernetes allows users to define the desired state of their application using declarative syntax, instead of having to manage individual resources manually.
- **Automation:** Kubernetes automates many common tasks, such as load balancing, networking, and service discovery



kubernetes

OASIS TOSCA



We are extending the TOSCA-based templating and encapsulating logic by adopting HELM for managing applications

CENTRALISED FULLY MANAGED SERVICES

The following are SaaS managed services available for every INFN Cloud user.

- INFN Cloud Object Storage implemented by a multi-region Openstack Swift (Backbone BARI and CNAF); S3 APIs and web interface via MinIO Gateway are provided as well. Open Policy Agent (OPA) is used to define access control rules based on IAM token claims.
- Notebook as a Service (NaaS) is a JupyterHub backed by kubernetes clusters hosted on the INFN Cloud Backbone (HA/failover on the two sites), connected with the Object Storage per notebook and data persistence
- INFN Cloud Registry provides a centrally managed registry, based on Harbor software.

INFN Cloud object storage



Notebooks as a Service
(NaaS)



INFN Cloud Registry



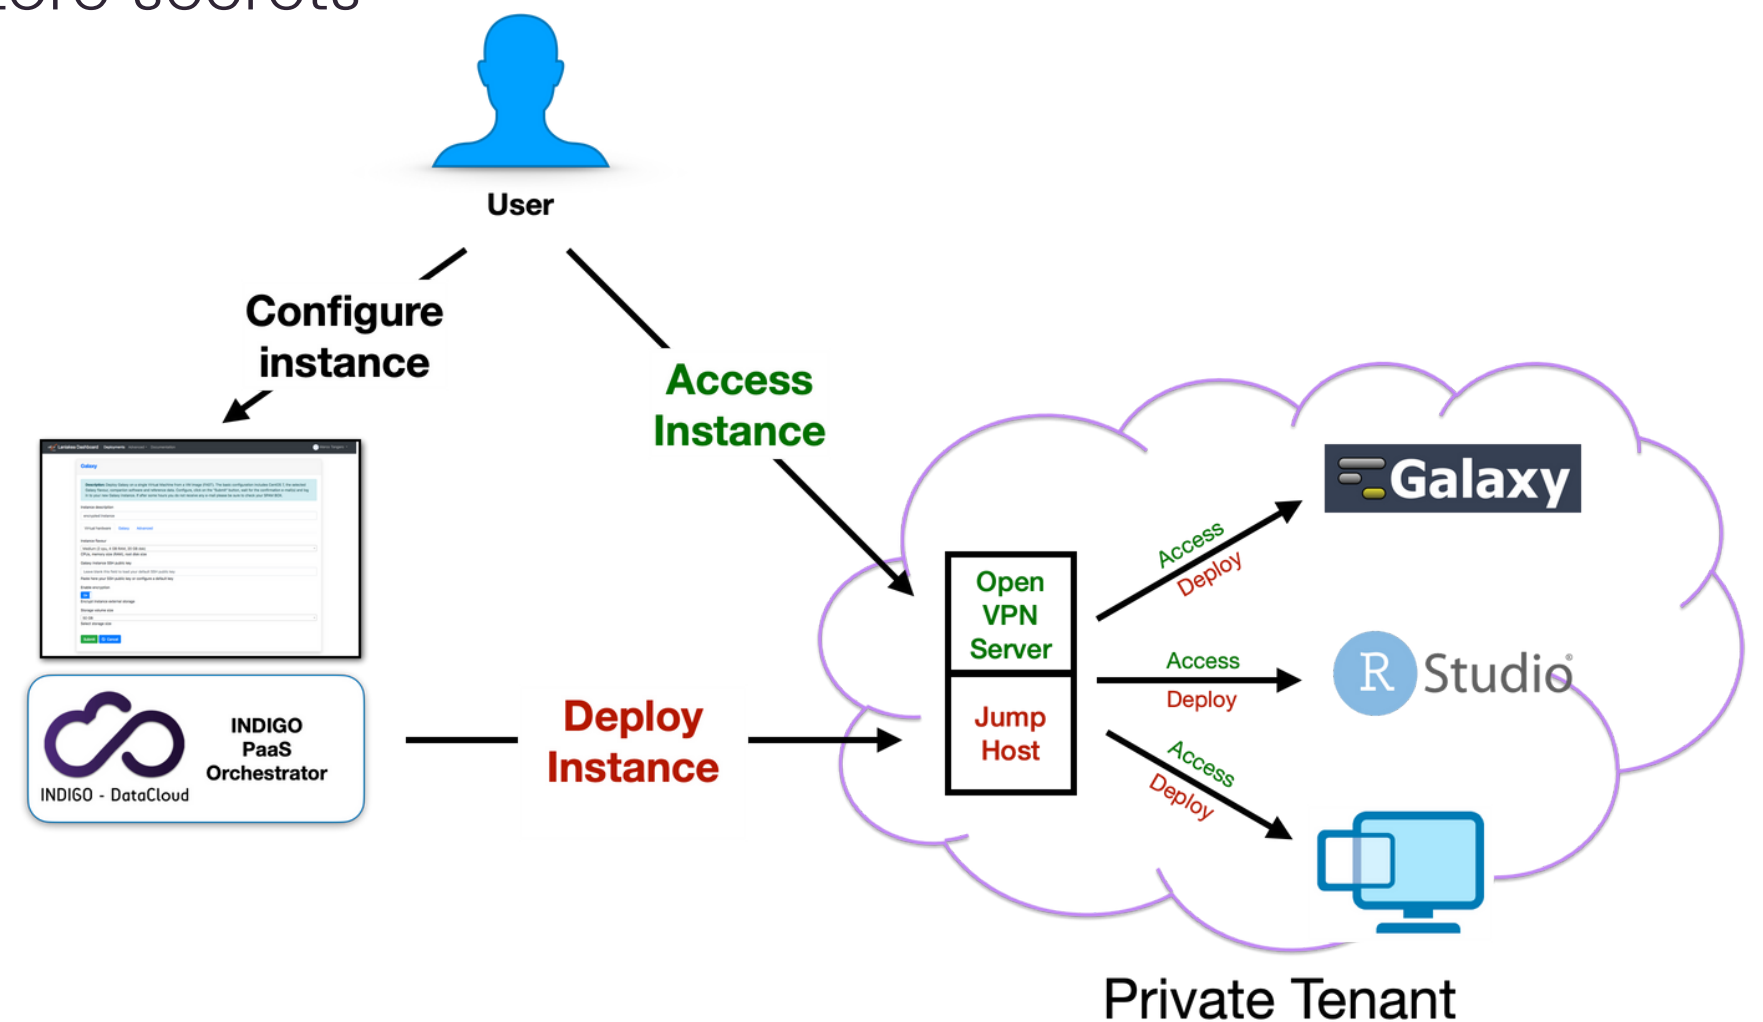
SECURE ENVIRONMENTS

Data security, legal, and ethical requirements for genetic and medical data are becoming increasingly stringent. Therefore we have been working to implement the necessary measures to improve the security of the deployed services.

- The PaaS Orchestration system has been extended to allow for the deployment of VMs on private networks (the contextualization is done via a "jump host")
- Moreover, a Vault instance is made available to store secrets

The deployment can be accessed by the user through VPN.

- To authenticate users, a PAM module has been developed that uses IAM by taking advantage of the device code authorization flow. The implementation is based on OpenVPN.




See also the contribution:

[Secure deployments of Galaxy Servers for analyzing personal and Health Data leveraging the Laniakea service](#)

CONTINUOUS TESTING & IMPROVEMENT

Average stage times:
(Average full run time: ~7min 23s)

	Declarative: Checkout SCM	Create test deployment	Scan endpoints	Declarative: Post Actions
#49 Oct 13 11:14 3 commits	6s	6min 31s	6s	11s
#48 Oct 13 09:53 No Changes	6s	6min 24s	5s	11s



Providing service templates that are frequently updated and include security patches is of crucial importance.

We have implemented an automatic testing system based on Jenkins that allows to test the TOSCA templates with pre-defined, fully automated job pipelines.

These pipelines perform automatic checks for each of the services available in the catalogue, including scans related to security vulnerabilities.

CONCLUSIONS

- INFN Cloud aims to address the new challenges of computing for scientific research.
 - It provides researchers with easy access to distributed resources in a transparent way, allowing them to perform complex analyses without requiring specialized IT knowledge.
 - Additionally, the platform enables the deployment of complex services with just a few clicks, streamlining the research process and enabling researchers to focus on their work.
-
- Future work will focus on improving the scheduling algorithm to enhance the platform's performance.
 - Integrated solutions for data management across distributed sites will be provided (based for example on Rucio + FTS)





THANK YOU

Marica Antonacci
antonacci@inf.n.it

www.cloud.infn.it

