



Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani

PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



Centro Nazionale di Ricerca in HPC,  
Big Data and Quantum Computing



# Deploy services over CLOUD

A brief overview of the INFN CLOUD and the offered services

- Michele Delli Veneri – [delliven@na.infn.it](mailto:delliven@na.infn.it)



Istituto Nazionale di Fisica Nucleare

Workshop Jennifer 2, CERN, 20/02/2024



Istituto Nazionale di Fisica Nucleare  
Sezione di Napoli



Finanziato dall'Unione europea  
NextGenerationEU



Ministero dell'Università e della Ricerca



Italiadomani

PIANO NAZIONALE DI RIPRESA E RESILIENZA



Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

# Who am I and most importantly what I do?



Tecnologist at INFN Section of Naples



Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing



openstack.®

IBISCO CLOUD NAPOLI





# What is CLOUD Computing?

Cloud computing refers to the **delivery of computing services** over the internet ("the cloud") to offer faster innovation, flexible resources, and economies of scale. In essence, cloud computing allows users to access and utilize computing resources, such as servers, storage, databases, networking, software, and analytics, without the need for owning or managing the physical infrastructure themselves.



**Infrastructure as a Service (IaaS):** Provides virtualized computing resources. Users can provision and manage virtual machines, storage, and networking infrastructure.

**Platform as a Service (PaaS):** Offers a platform allowing customers to develop, run, and manage applications without worrying about the underlying infrastructure. PaaS providers handle the hardware and software infrastructure, including middleware, development tools, database management systems, and runtime environments.

**Software as a Service (SaaS):** Delivers software applications. Users can access and use the software without needing to install or maintain it.



Finanziato  
dall'Unione europea  
NextGenerationEU







Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



# IBISCO Cloud NA – OpenStack and Overview



Openstack Services are built on VM generated through Proxmox. This ensure recoverability and HA



Finanziato dall'Unione europea  
NextGenerationEU



Ministero dell'Università e della Ricerca



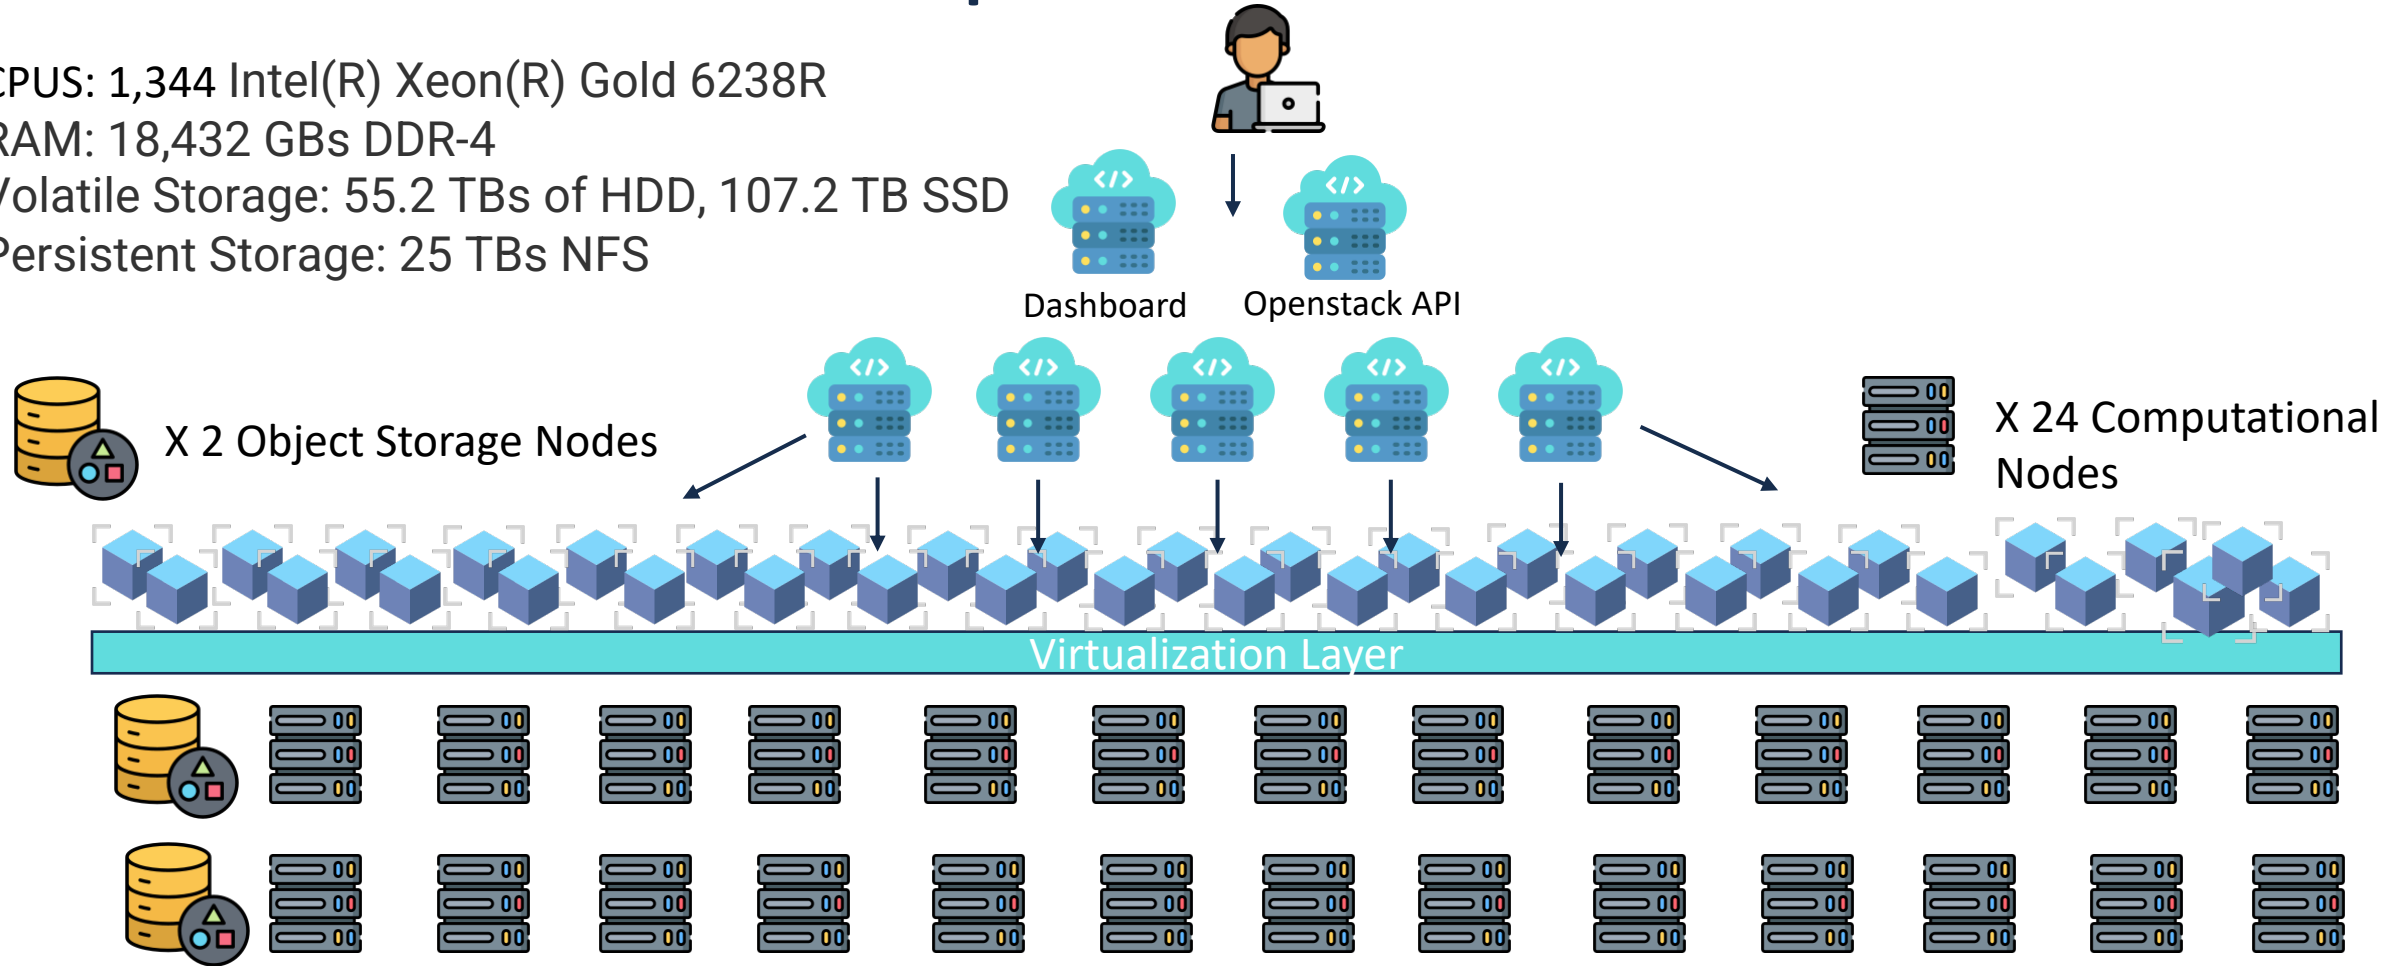
Italiadomani  
PIANO NAZIONALE DI RIPRESA E RESILIENZA



Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

# IBISCO Cloud NA – OpenStack and Overview

CPUS: 1,344 Intel(R) Xeon(R) Gold 6238R  
RAM: 18,432 GBs DDR-4  
Volatile Storage: 55.2 TBs of HDD, 107.2 TB SSD  
Persistent Storage: 25 TBs NFS





Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



# What is CLOUD Federation?

Cloud federation, also known as cloud interoperability or **federated cloud computing**, refers to the integration and collaboration of multiple cloud computing environments to enable seamless data and resource sharing across disparate cloud platforms. The goal of cloud federation is to create a unified computing environment that spans multiple cloud providers, enabling users to leverage resources and services from different providers as if they were part of a single, cohesive infrastructure.



## Lightly Coupled Cloud Federation:

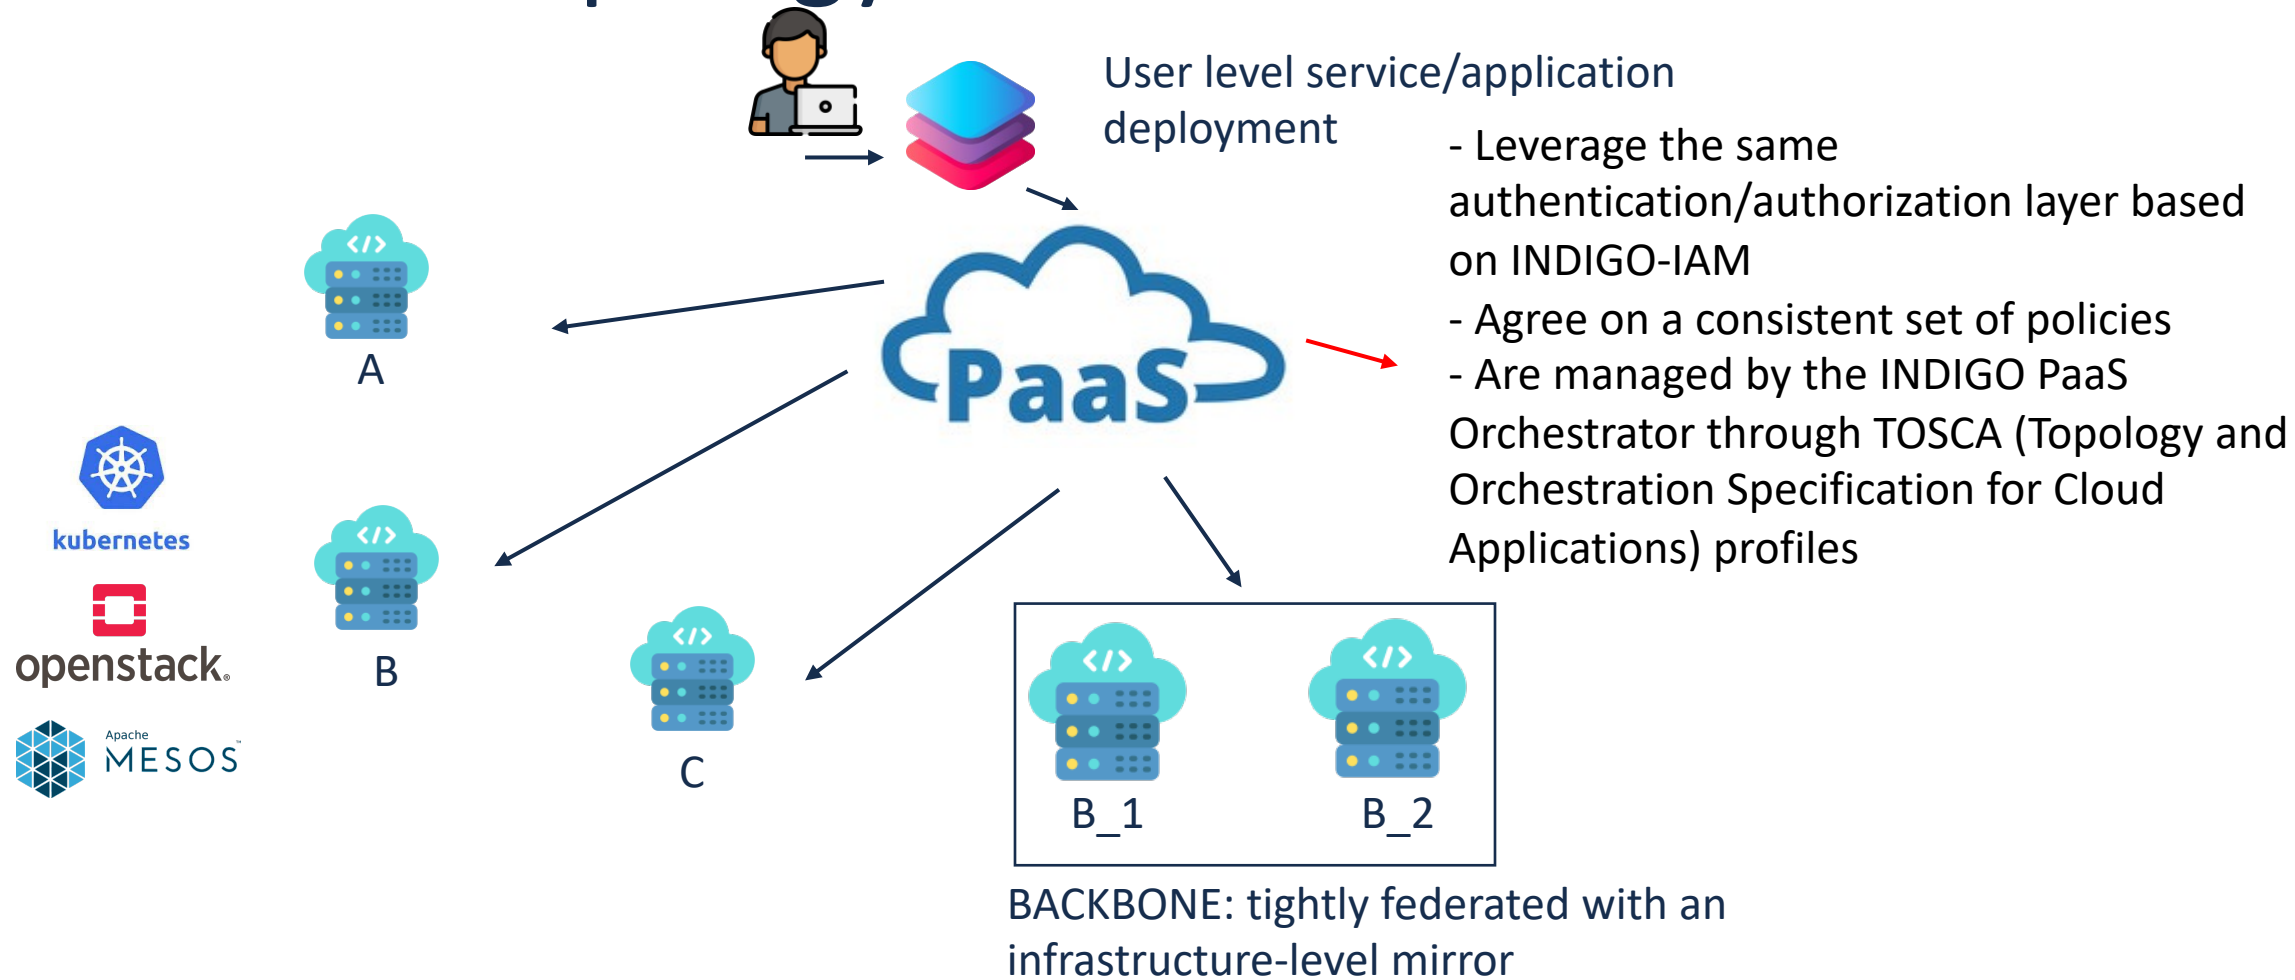
- the cloud environments maintain a degree of independence and autonomy;
- Each cloud provider retains control over its infrastructure, policies, and services, no resource sharing between the federated clouds

## Tightly Coupled Cloud Federation:

- the participating cloud environments are tightly integrated and interdependent, forming a unified and cohesive computing environment; advanced capabilities such as workload migration, federated identity management, unified billing, and cross-cloud service orchestration



# INFN CLOUD Topology





Finanziato dall'Unione europea  
NextGenerationEU



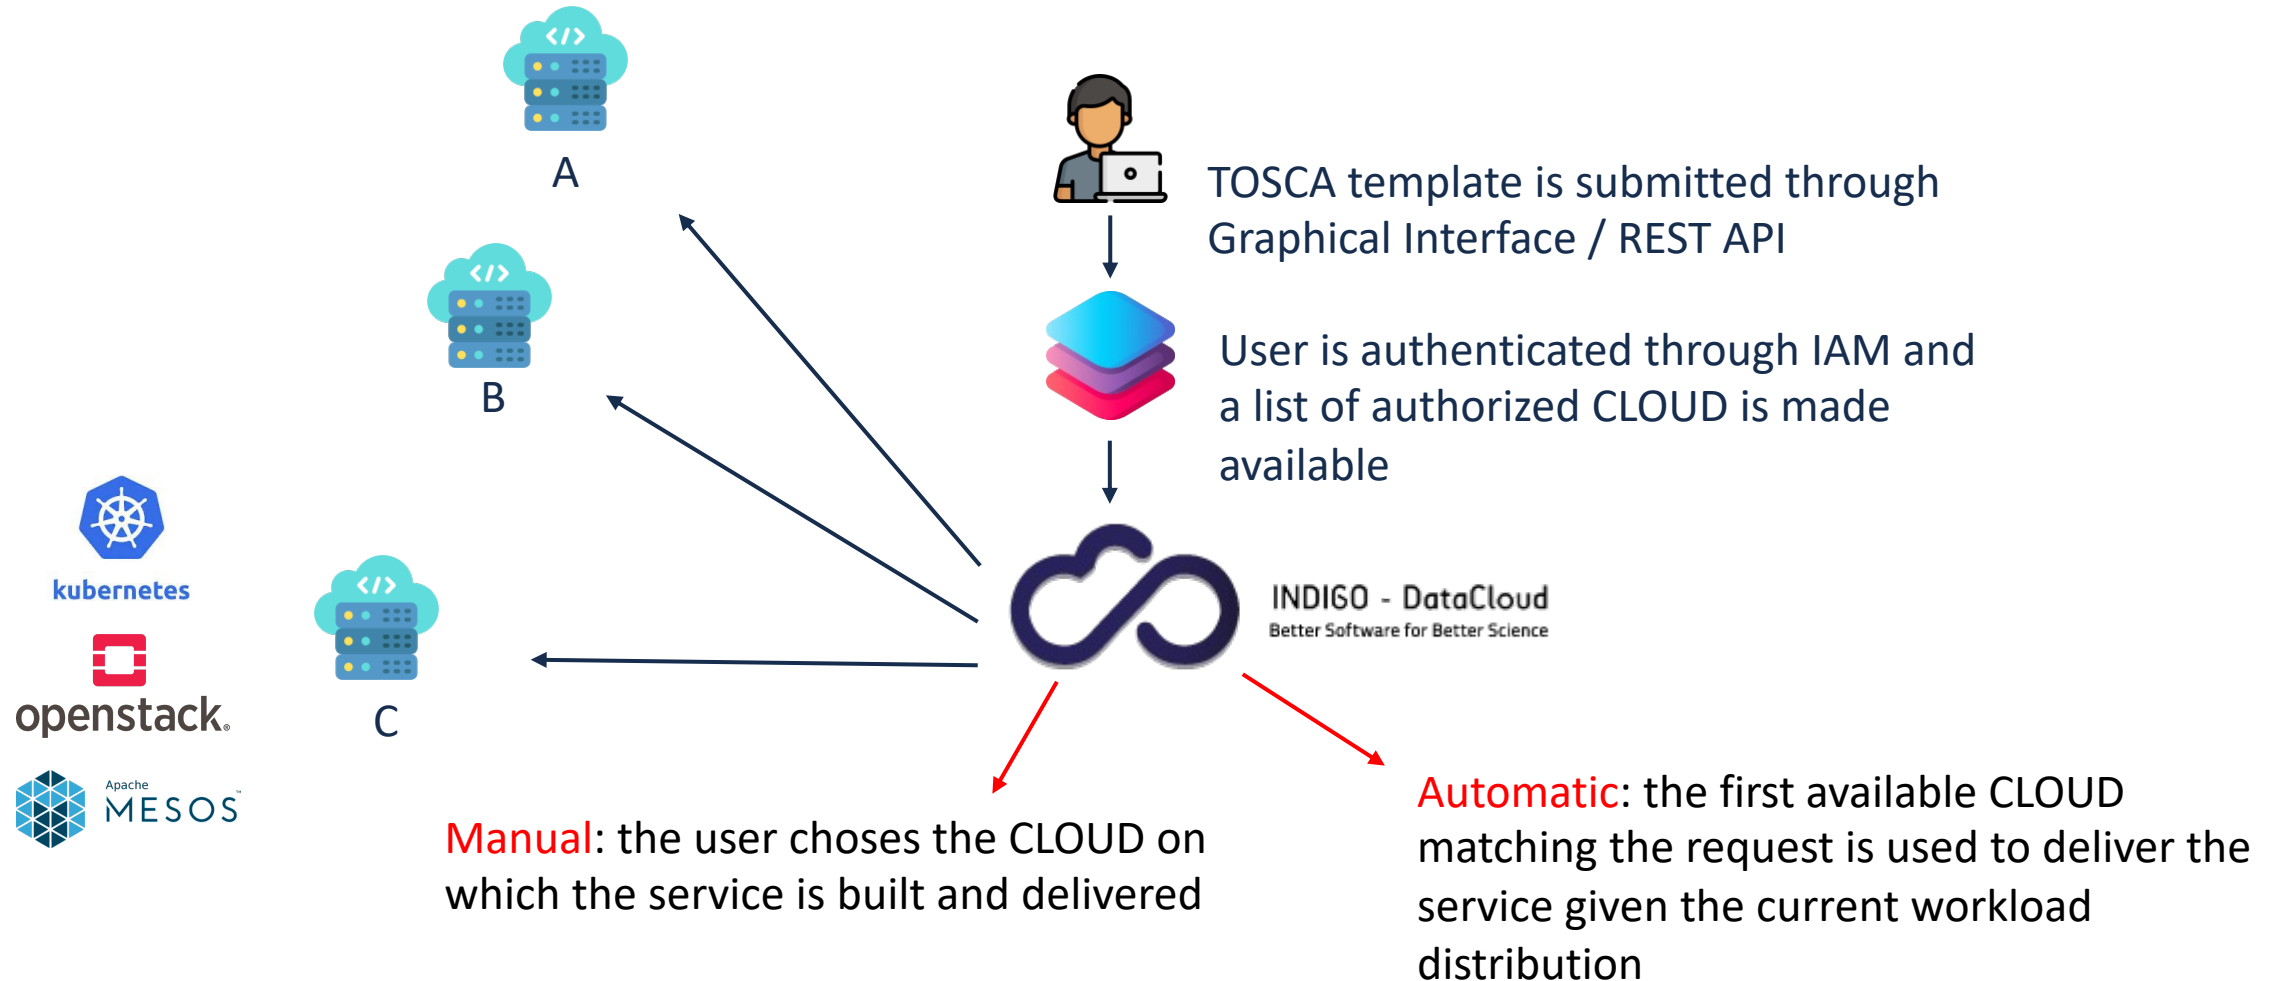
Ministero dell'Università e della Ricerca



Italiadomani  
PIANO NAZIONALE DI RIPRESA E RESILIENZA



# INDIGO PaaS Orchestrator

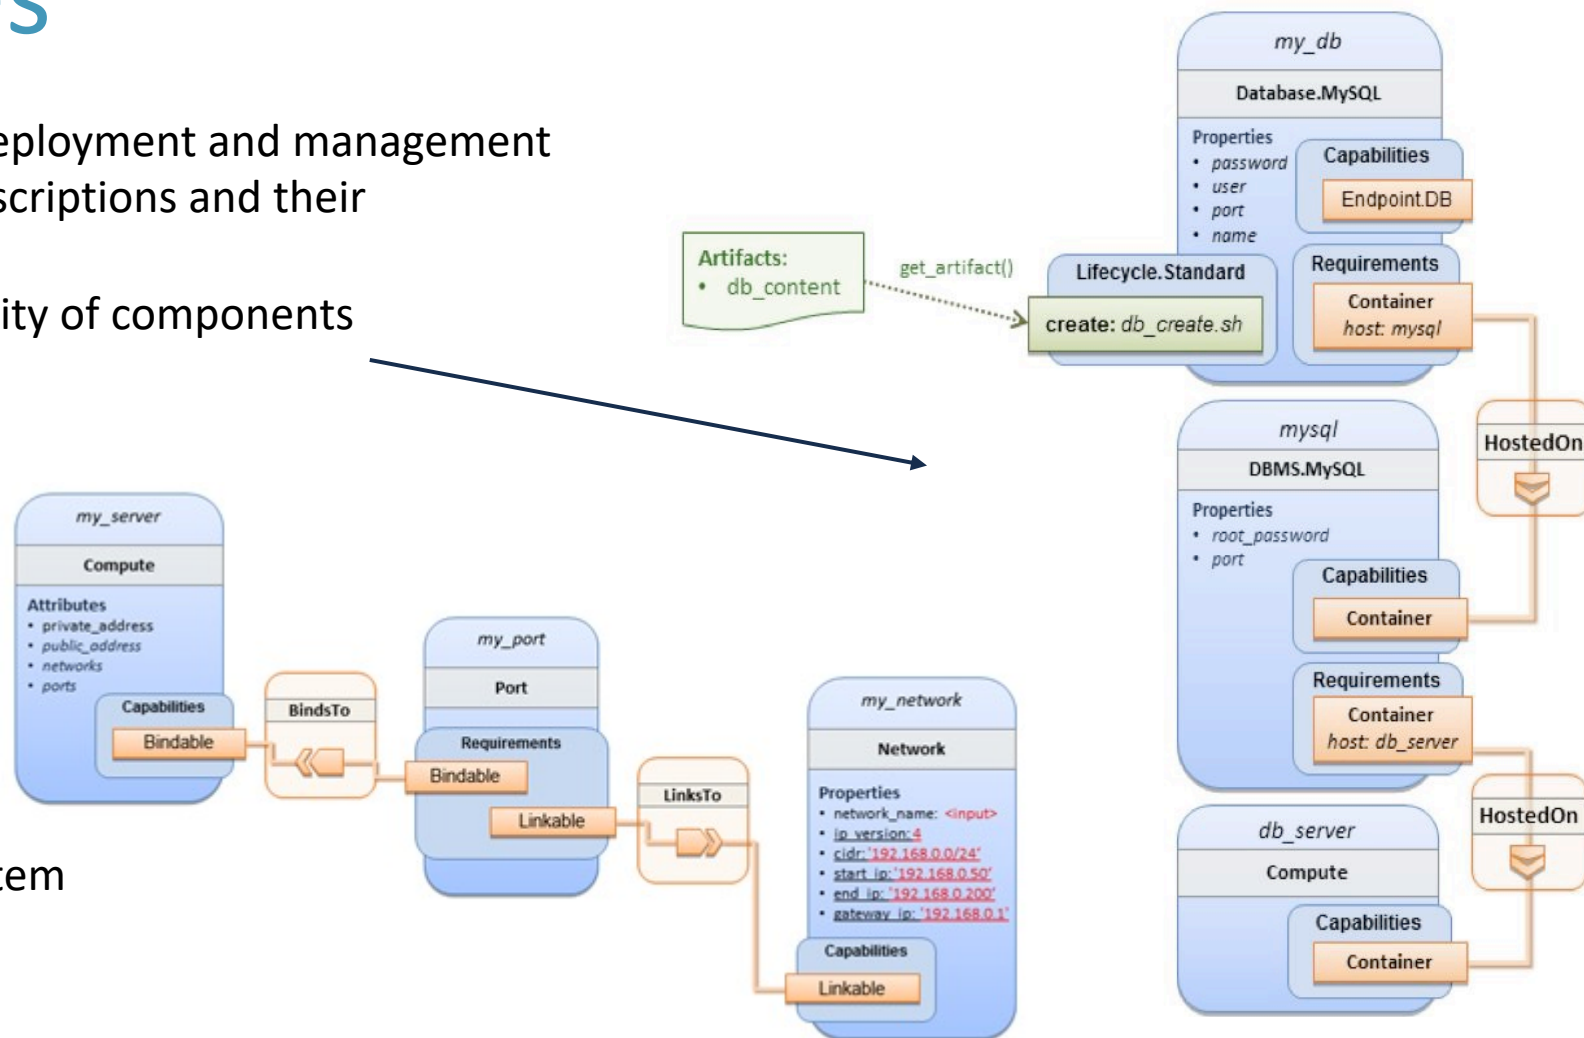






# TOSCA Templates

- Allow the automatic application deployment and management
- Allow portability of application descriptions and their management
- Allow interoperability and reusability of components




Allow the creation of complex system through a “LEGO-like” syntax.



# TOSCA Service Catalogue

The Catalogue is a graphical representation of TOSCA templates:

- Each card in the catalogue is associated to one or more templates
- Complex templates are built by combining simpler templates
- Once a template has been built and tested is made available to scientists



1. Allows to easily build services on top of a IaaS and PaaS infrastructure
2. Lowers the entry barrier for scientists which do not need any programming skill to use / create their applications

COMPLEXITY

- Creation of VMs with different flavors, images and sizes
- Creation of containers or services via docker-compose files
- Building blocks as a service (Kubernetes or Mesos clusters as a service)
- Pre-configured environments for data analytics (ElasticSearch, Spark)
- Non volatile, object storage and posix-compliant file systems connected to high-layer services (Jupyter Notebook as a service with permanent storage)
- Data Pipelines for specific experiments (MPI Simulations, ML-optimized environments)
- User-level encryption of disk volumes, Monitoring Services, Spark Clusters, Kubernetes Clusters



# PaaS Interfaces - user perspective



1. PaaS REST API
2. Orchestrator CLI
3. Orchestrator Bindings (python)



## Requirements:

- Basic TOSCA Knowledge
- Installation and configuration of CLI tools
- Authentication token managements



1. Orchestrator dashboard reachable through your browser



## Requirements:

- No TOSCA Knowledge is needed
- Accessible through web
- Automatic token management through user login / authorization
- Easy to use and intuitive



Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



# PaaS Dashboard – Authentication and Authorization

INFN Cloud Dashboard

Welcome to the  
INFN Cloud

**INFN**  
CLOUD  
Istituto Nazionale di Fisica Nucleare

Welcome to **infn-cc**

Sign in with

**INFN**  
CCR - AAI

Not a member?

Apply for an account

INDIGO IAM manages the authentication/authorization through the whole stack (from PaaS to IaaS)

- Users are organized into different IAM groups which correspond to Virtual Organizations (VO) roles
- Each IAM group is mapped to a specific set of resource quotas spread over different Cloud infrastructures
- A specific set of services is made available to the user on the basis of the resource quotas
- Once a service has been selected, a TOSCA template is used to build and deliver the service to the user



# PaaS Dashboard – Service Request Customization

**Virtual machine**

**Description:** Launch a compute node getting the IP and SSH credentials to access via ssh

Deployment description

description

Configuration **Advanced**

hostname

ports

Add rule

Ports to open on the host

flavor

--Select--

Number of vCPUs and memory size of the Virtual Machine

operating\_system

--Select--

Operating System for the Virtual Machine

Submit Cancel

```

topology_template:
  inputs:
    num_cpus:
      type: integer
      description: Number of virtual cpus for the VM
      required: true
    mem_size:
      type: scalar-unit.size
      description: Amount of memory for the VM
      required: true
    os_distribution:
      type: string
      required: true
      description: Operating System distro
      constraints:
        - valid_values: [ "ubuntu", "centos" ]
    os_version:
      type: version
      required: true
      description: Operating System distribution version
      constraints:
        - valid_values: [ "16.04", "18.04", "7" ]
    service_ports:
      type: map
      required: false
      constraints:
        - min_length: 0
      entry_schema:
        type: toska.datatypes.network.PortSpec
        description: Ports to open on the host
  
```

The configuration form allows the user to specify their requirements for the deployment

- It hides all the complexity of TOSCA
- Complex TOSCA types are managed with dedicated Javascript functions

Protocol	Port Range	Source	
TCP	e.g. [8080,8082] or 80	0.0.0.0/0	Remove

Add rule

Ports to open on the host



# PaaS Dashboard – Advance Customization

**Virtual machine**

**Description:** Launch a compute node getting the IP and SSH credentials to access via ssh

Deployment description  
tester

Configuration **Advanced**

Configure scheduling:  
 Auto  Manual

Set deployment creation timeout (minutes) 720

Do not delete the deployment in case of failure

Send a confirmation email when complete

Submit Cancel

**Virtual machine**

**Description:** Launch a compute node getting the IP and SSH credentials to access via ssh

Deployment description  
tester

Configuration **Advanced**

Configure scheduling:  
 Auto  Manual

Select a provider:

- CLOUD-VENETO: org.openstack.nova
- CLOUD-IBISCO-NAPOLI: org.openstack.nova
- CLOUD-CNAF-T1: org.openstack.nova

Submit Cancel

Dashboard interacts with the SLA Manager Service to get the list of available Cloud providers for the user



Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



Centro Nazionale di Ricerca in HPC,  
Big Data and Quantum Computing

# PaaS Dashboard – Deployment Outputs

INFN Cloud Dashboard Deployments ▾ Advanced ▾ External Links ▾ priv-admins/catchall ▾ Michele Delli Veneri ▾

**My deployments** Refresh + New deployment

Show  entries Search:

Description	Deployment identifier	Status	Creation time	Deployed at	Actions
test deployment	<a href="#">11eecbf8-6c78-11aa-8be4-56fce75e0bfa</a>	CREATE_COMPLETE	2024-02-15 11:50:00	CLOUD-IBISCO-NAPOLI	<span>Details</span>

Showing 1 to 1 of 1 entries

**Actions**

- Details
- Edit
- Show template
- Log
- Request Ports
- Manage VMs
- Lock
- Delete

[11eecbf8-6c78-11aa-8be4-56fce75e0bfa](#) Back

Description: test deployment

Overview Input values Output values

```

vpn_server: ibisco-endpoint.na.infn.it:42120
node_ip: 192.168.204.88
ssh_account: delliven
vpn_client_conf_url: https://baltig.infn.it/infn-cloud/vpnconfiles/-/raw/main/vpn-ibisco-endpoint.na.infn.it:42120-client.ovpn?inline=false

```



# PaaS Dashboard – Secret Management

## SSH keys management

SSH keys allow you to establish a secure connection between your computer and your virtual server(s).

### Upload SSH public key

Paste your public SSH key, which is usually contained in the file `~/.ssh/id_ed25519.pub` or `~/.ssh/id_rsa.pub` and begins with `ssh-ed25519` or `ssh-rsa`. Don't use your private SSH key.

Upload

### Create new key pair

SSH key pair will be created from scratch. The private key will be safely stored in the Vault, while the public key will be stored in the Dashboard database.

Create new SSH key pair

The Dashboard is integrated with Hashcorp Vault to support:

- ssh key pair management
- Service credential stores (e.g. AWS)

The Vault has been integrated with the Cloud IAM and policies grant read/write permissions to specific Vault paths depending on the user claims

## Service Credentials

Site	Endpoint	Manage your credentials
AWS-us-east-1	https://ec2.us-east-1.amazonaws.com	GET SET DELETE
AWS-us-east-2	https://ec2.us-east-2.amazonaws.com	GET SET DELETE





Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA




Centro Nazionale di Ricerca in HPC,  
Big Data and Quantum Computing

# INFN Cloud – List of Services


CENTRALISED SERVICES:

**INFN Cloud object storage**  
the centrally managed service based on  
MinIO-Gateway




[Go to service](#)

**Notebooks as a Service (NaaS)**




**INFN-Cloud monitoring**




- The list of Services depends on the User IAM group.
- It grows as users and admins create and configure TOSCA templates for services

ON-DEMAND SERVICES:

**Virtual machine**



**INDIGO IAM as a Service**



**Spark + Jupyter cluster**




**Jupyter + Matlab (with persistence for Notebooks)**

**Docker compose**  
Deploy a virtual machine with docker engine and docker compose pre-installed. Optionally run a docker compose file [...]

[Read More](#) [Configure](#)

**Elasticsearch and Kibana**




**HTCondor mini**



**Sync&Share aaS**

**Run docker**



**Kubernetes cluster**



**Jupyter with persistence for Notebooks**





# INFN Cloud – Jupyter Notebook with Persistence

Jupyterhub is automatically installed and configured in a Virtual Machine

- The integration with the INFN Cloud IAM allows authorized users to login and spawn their containerized (docker) service
- It can be built from a default jupyter image or customized docker image (containing, for example, a scientific pipeline)
- The integration with the INFN Cloud Storage provide persistence to the Notebooks

Jupyter with persistence for Notebooks

**Description:** Run Jupyter on a single VM enabling Notebooks persistence

Deployment description  
description

General Authorizations Advanced

num\_cpus  
2  
Number of virtual cpus for the VM

mem\_size  
4 GB  
Amount of memory for the VM

enable\_monitoring  
false  
Enable/disable monitoring

jupyter\_images  
dodasts/snj-base-lab-persistence.v1.1.1-snj  
Default image

jupyterlab\_collaborative  
false  
enable the jupyter collaborative service

jupyterlab\_collaborative\_image  
dodasts/snj-base-jlab.v1.1.1-snj  
Default image for jupyter collaborative service

contact\_email  
Email address of certificate management administrator

ports  
Add rule  
Ports to open on the VM

certificate\_type  
letsencrypt-prod

Submit Cancel

Jupyter with persistence for Notebooks

**Description:** Run Jupyter on a single VM enabling Notebooks persistence

Deployment description  
description

General Authorizations Advanced

iam\_groups  
IAM groups for authorization management (comma delimited list of strings)

iam\_admin\_groups  
IAM groups for JupyterHub ADMIN authorization management (comma delimited list of strings)

**Warning!** You have not filled some mandatory fields. Please check the red boxes in each tab.

Submit Cancel



Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



# INFN Cloud – Community Customizations

Computational environment for  
Machine Learning INFN  
(ML\_INFNO)



Working Station for CYGNO  
experiment



Community customized environments are built upon the Jupyter with persistent storage VM use-case

New Environments can be constructed through interactions with the Community



Finanziato dall'Unione europea  
NextGenerationEU



Ministero dell'Università e della Ricerca

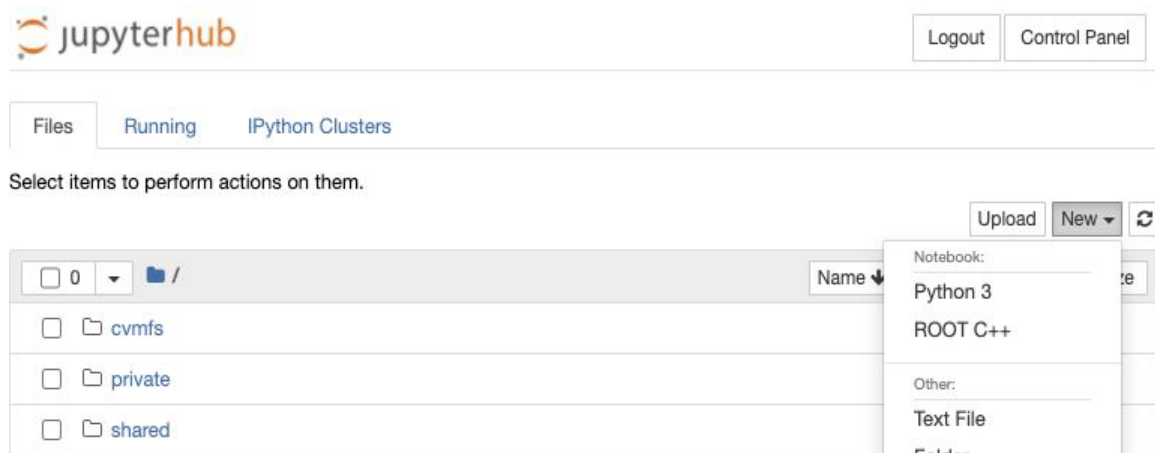


Italiadomani  
PIANO NAZIONALE DI RIPRESA E RESILIENZA



Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

# INFN Cloud – Community Customizations



**CernVM-FS**  
CernVM File System

Virtual File System mounted to allow data access and storage

Specific Libraries are pre-configured

GPUs are made available through GPU passthrough





Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca

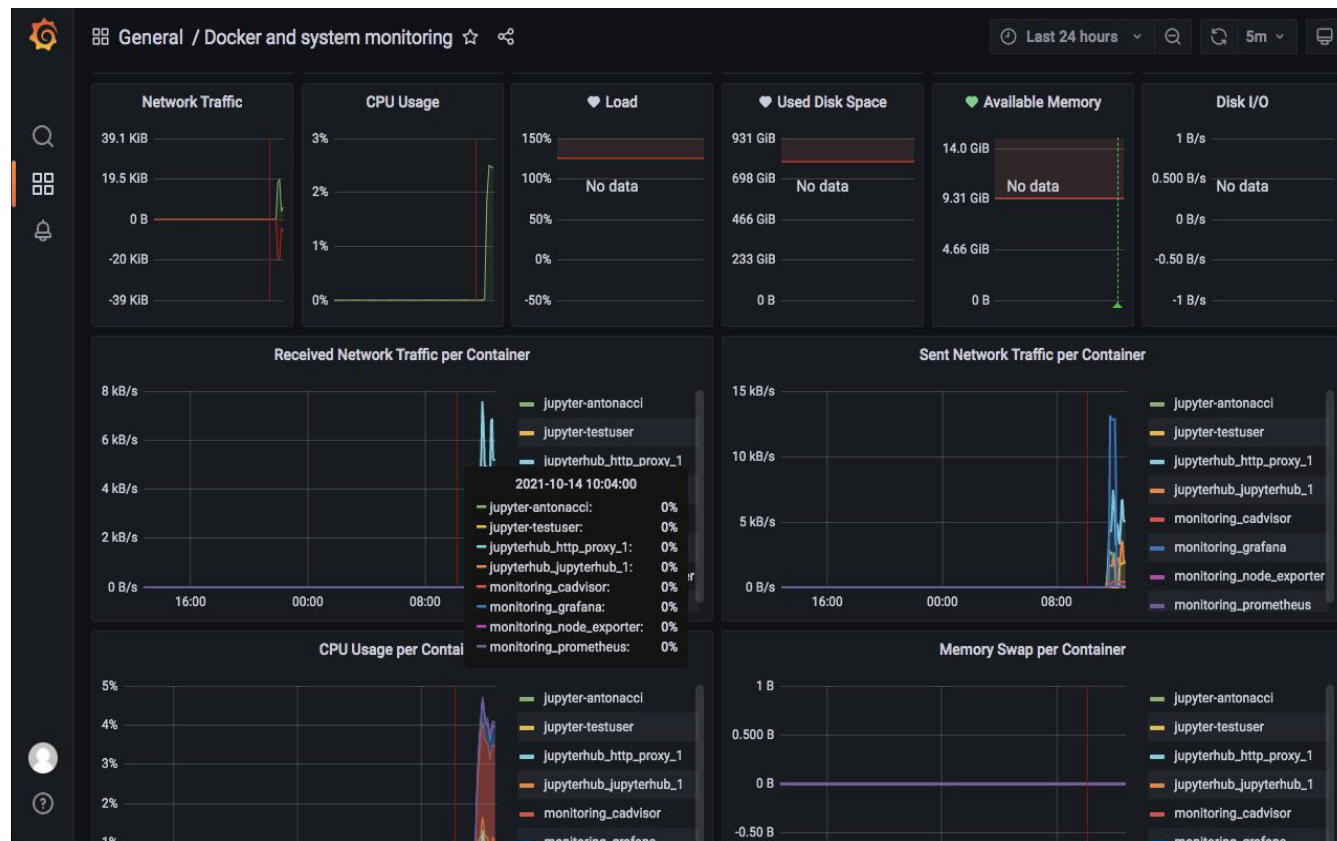
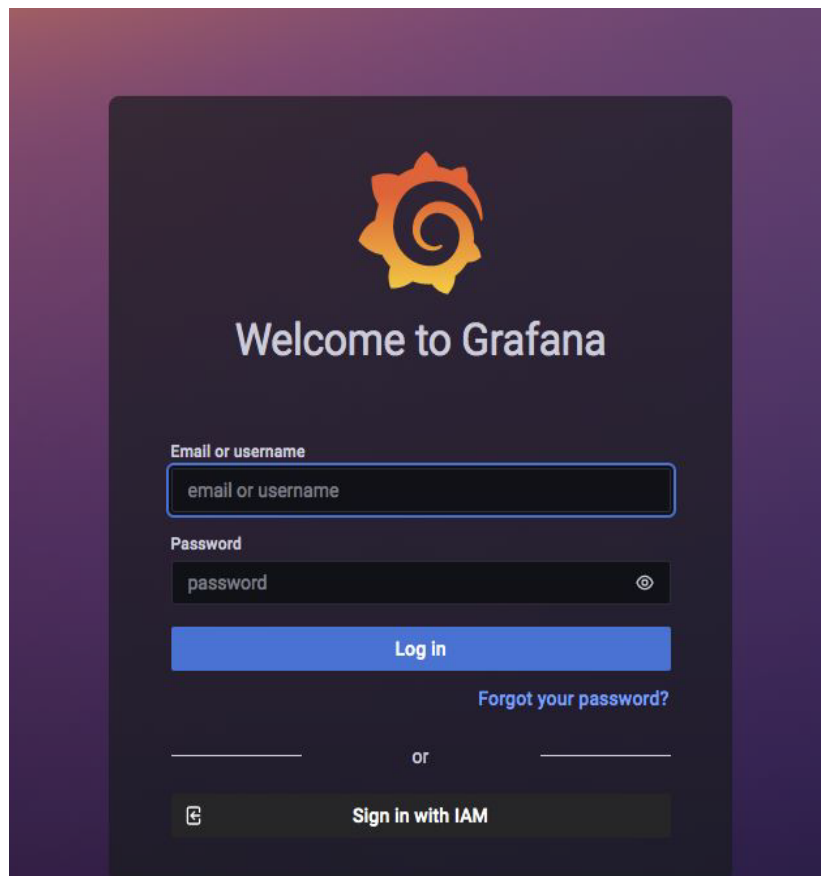


Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



Centro Nazionale di Ricerca in HPC,  
Big Data and Quantum Computing

# INFN Cloud – Monitoring through Grafana



APEL + CASO + STUNNEL on the CLIENT Side

Thank you  
for the  
attention

