Introduction to INFN Cloud services

Marica Antonacci (INFN BA)

Tutorial Days CCR :
Amministrazione di risorse erogate attraverso INFN Cloud
3-11 November 2021

The INFN Cloud

- **INFN Cloud aims to offer a full set of high-level cloud services to INFN user communities**
  - the service catalogue is not static: new applications are included through a defined “on-boarding” process for new use-cases
- **Architecturally INFN Cloud is a federation of existing infrastructures**
  - the *INFN Cloud backbone*, consists of two tightly coupled federated sites: BARI and CNAF
  - a scalable set of satellite sites, geographically distributed across Italy, and loosely coupled.
- **Key enabling factors for the federation**
  - leverage the same authentication/authorization layer based on **INDIGO-IAM**
  - agree on a consistent set of policies and participation rules (user management, SLA, security, etc.)
  - transparent and dynamic orchestration of the resources across all the federated infrastructures through the **INDIGO PaaS Orchestrator**

PaaS Orchestration System (from 10Km)

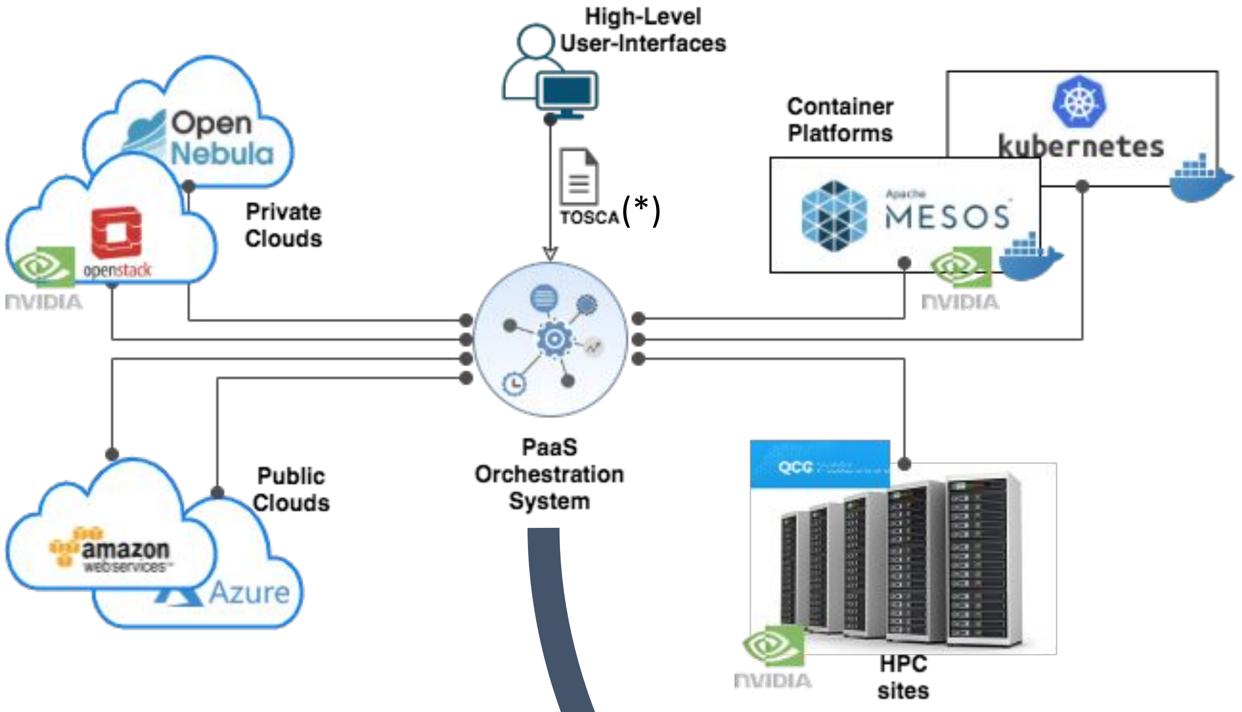

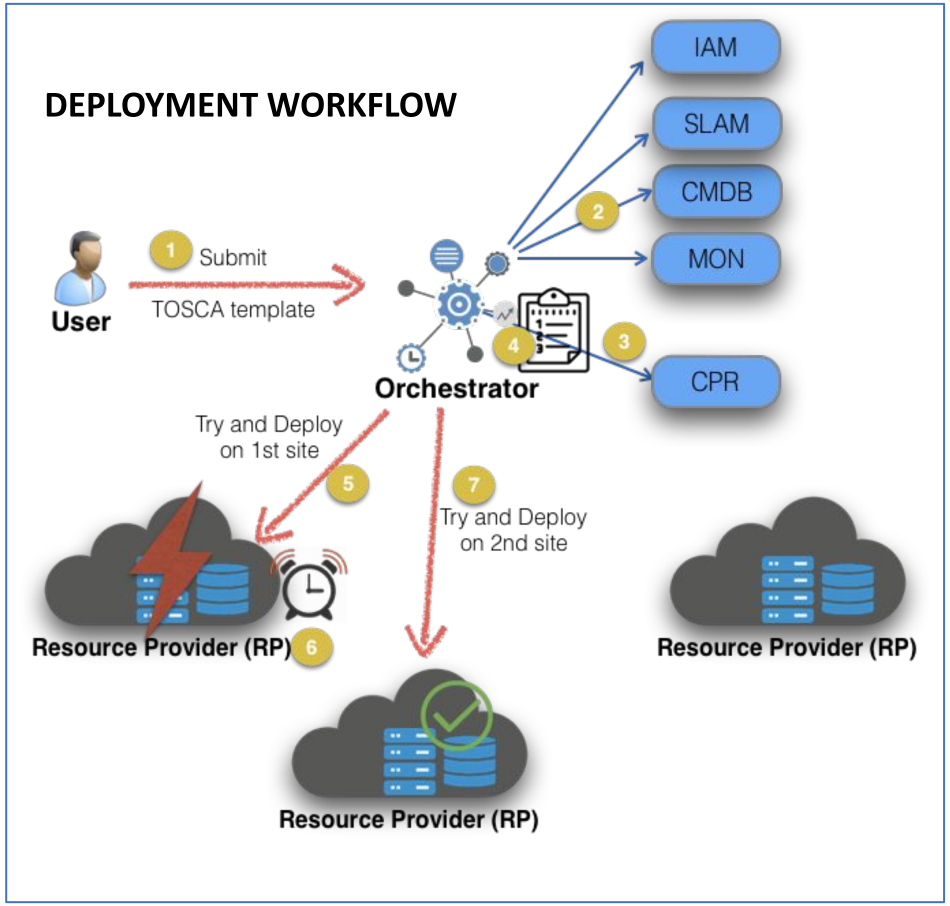

(*) **T**opology and **O**rchestration **S**pecification for **C**loud **A**pplications

Ref: TOSCA Simple Profile in YAML Version 1.1

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.

The service catalogue

The catalogue is a graphical representation of the TOSCA templates repository that we have been developing extending the INDIGO-DC custom types

- Each card in the catalogue is associated to one or more templates
- We are following a **lego-like** approach, building on top of reusable components and exploiting the TOSCA service composition pattern

Main objectives:

#1 - build added value services on top of IaaS and PaaS infrastructures
#2 - lower the entry barrier for non-skilled scientists

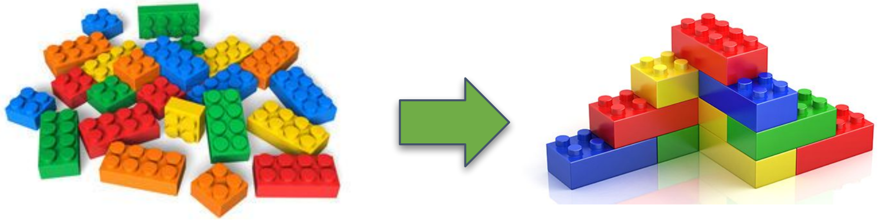

Which services are available?

SIMPLE

- Creation of VMs with different flavors and sizes.
- Creation of containers or of services 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 data analytics (e.g. using ElasticSearch and Kibana or Spark).
- Non volatile, 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.

COMPLEX

The service catalogue can be easily extended with the simple addition/customization of TOSCA templates.

The INFN Cloud Dashboard

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

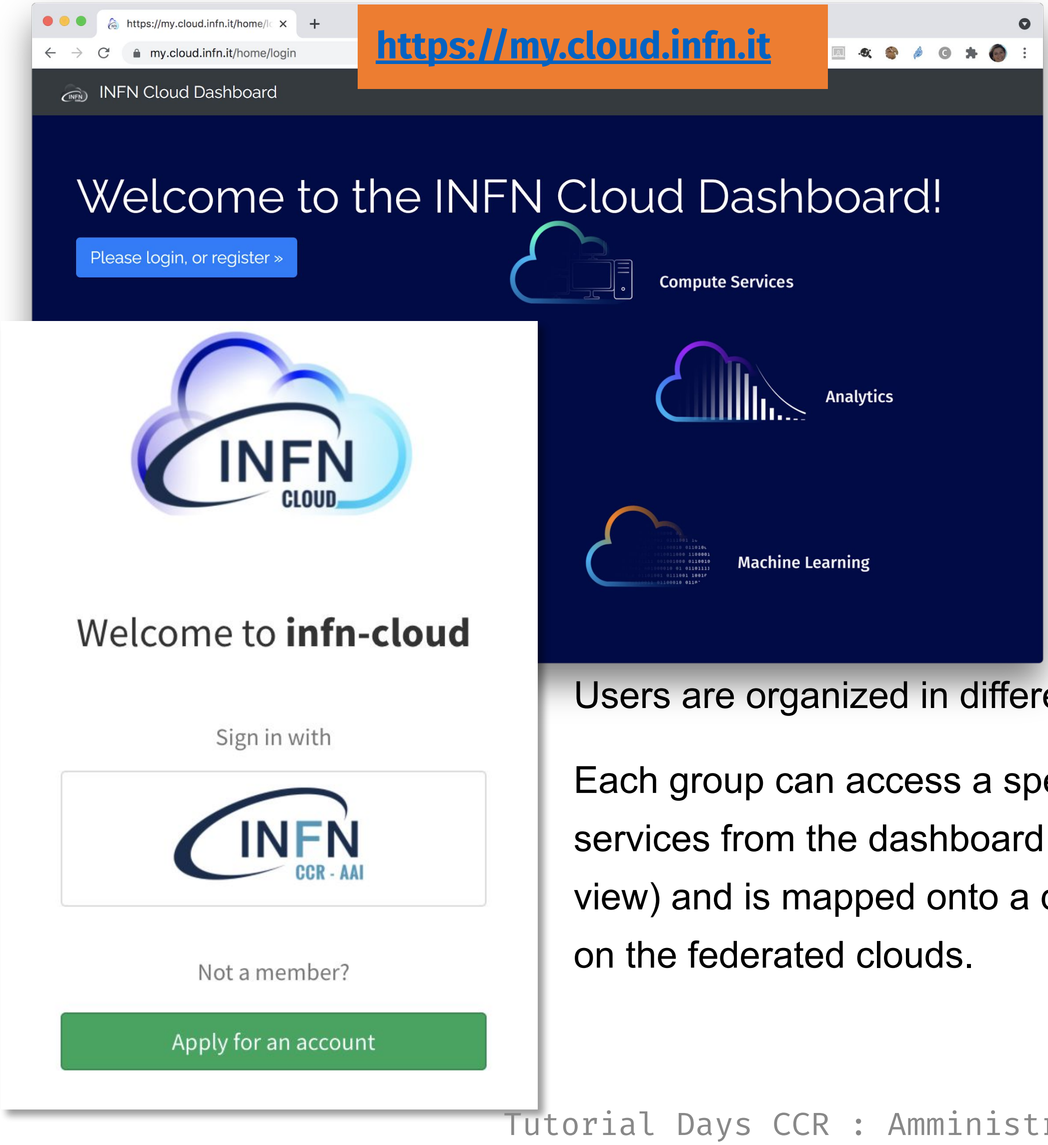

https://my.cloud.infn.it

my.cloud.infn.it/home/login

INFN Cloud Dashboard

Welcome to the INFN Cloud Dashboard!

Please login, or register »

Welcome to **infn-cloud**

Sign in with

Not a member?

Apply for an account

Users are organized in different IAM **groups**.

Each group can access a specific set of services from the dashboard (personalized view) and is mapped onto a dedicated tenant on the federated clouds.

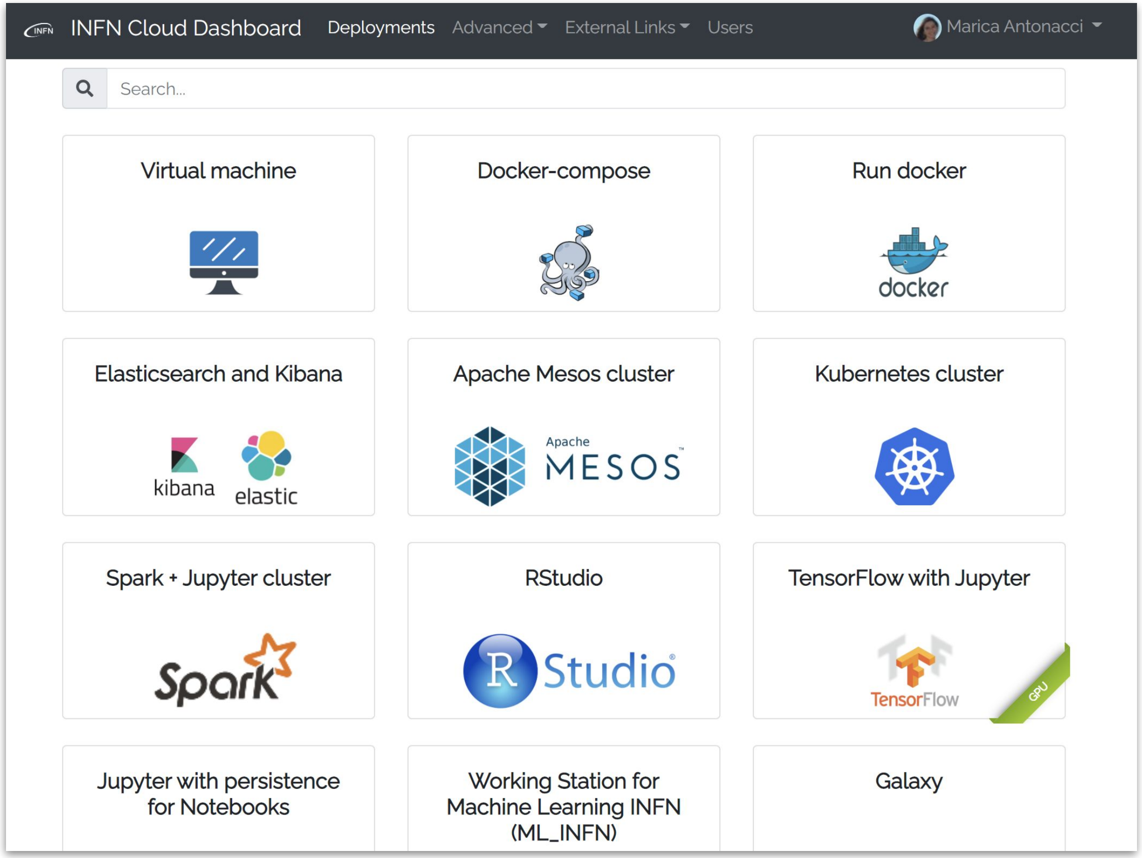

INFN Cloud Dashboard | Deployments | Advanced ▾ | External Links ▾ | Users | Marica Antonacci ▾

Search...

- Virtual machine
- Docker-compose
- Run docker
- Elasticsearch and Kibana
- Apache Mesos cluster
- Kubernetes cluster
- Spark + Jupyter cluster
- RStudio
- TensorFlow with Jupyter
- Jupyter with persistence for Notebooks
- Working Station for Machine Learning INFN (ML_INFN)
- Galaxy

The INFN Cloud Dashboard

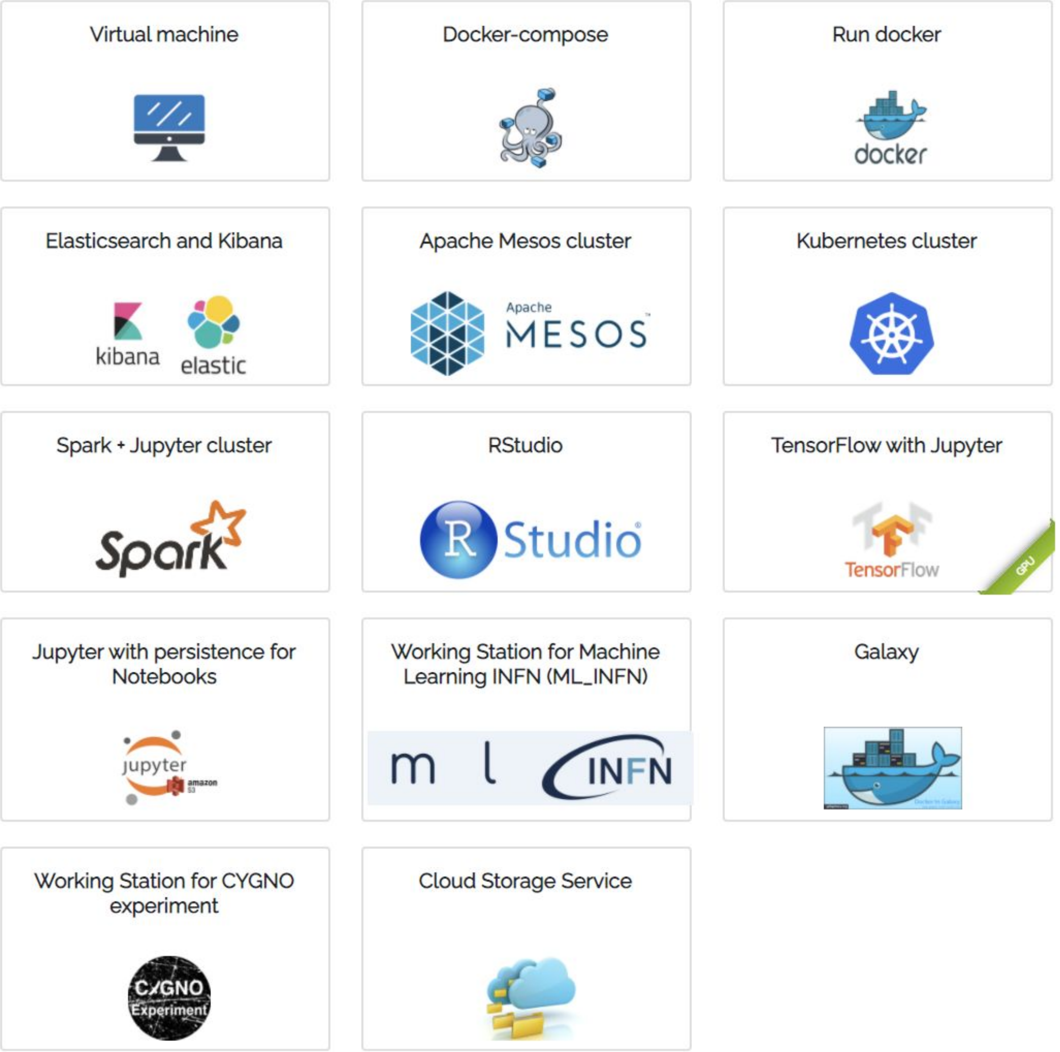

The services are **easily customizable** and configurable directly by users

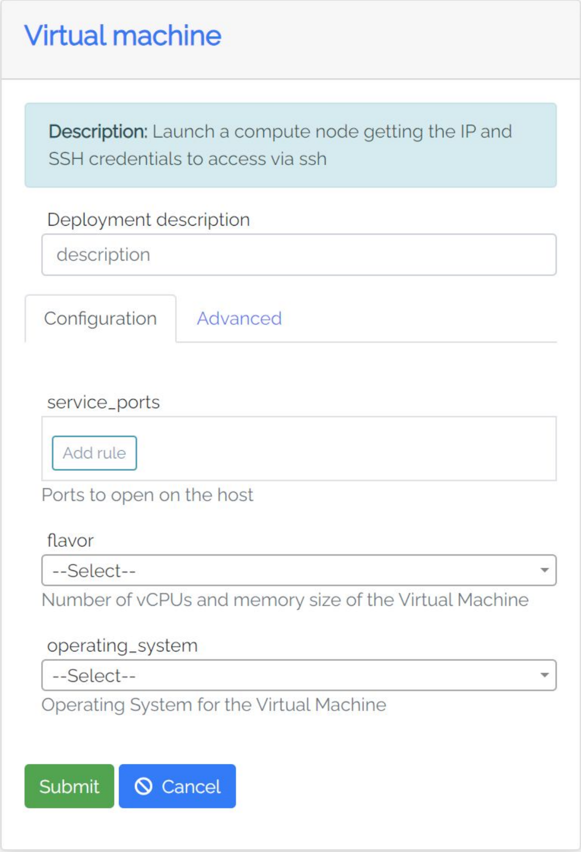

Virtual machine

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

Deployment description

description

Configuration | Advanced

service_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

Transparent, multi-site **federation or site selection** made manually by the user

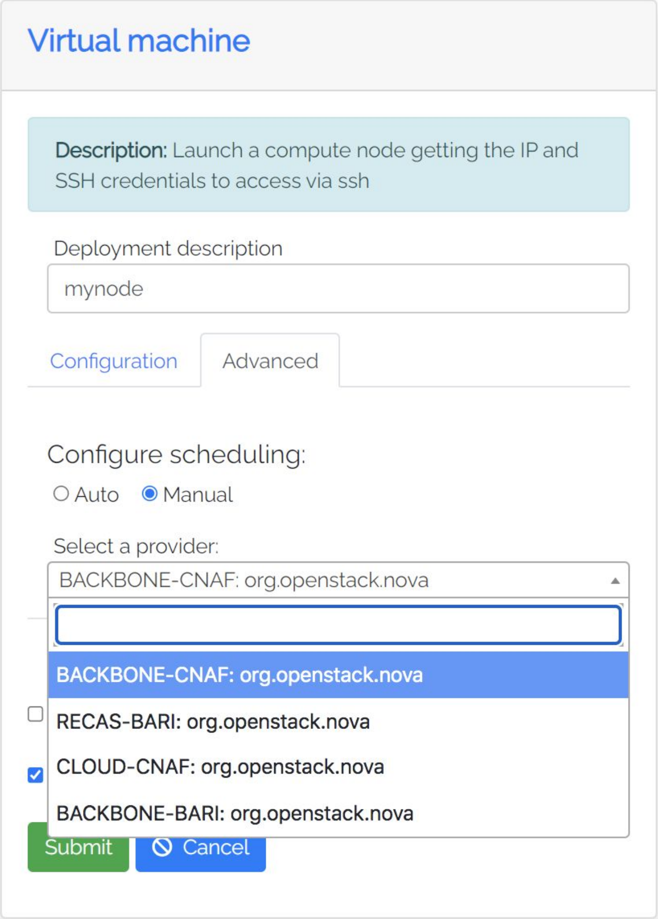

Virtual machine

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

Deployment description

mynode

Configuration | Advanced

Configure scheduling:

○ Auto ● Manual

Select a provider:

BACKBONE-CNAF: org.openstack.nova

- BACKBONE-CNAF: org.openstack.nova
- RECAS-BARI: org.openstack.nova
- CLOUD-CNAF: org.openstack.nova
- BACKBONE-BARI: org.openstack.nova

Submit | Cancel

Virtual machine use-case

How to start a virtual server on INFN Cloud

Configure your VM

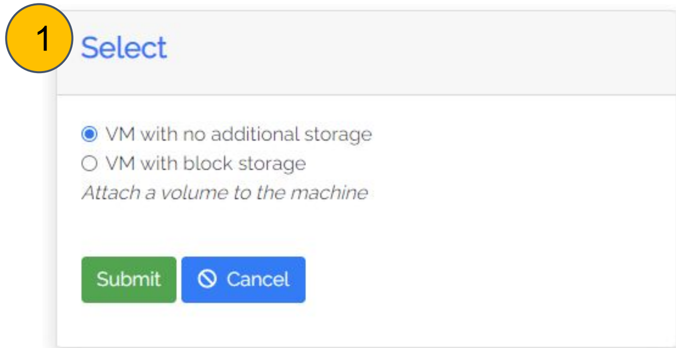

1

Select

● VM with no additional storage
○ VM with block storage
Attach a volume to the machine

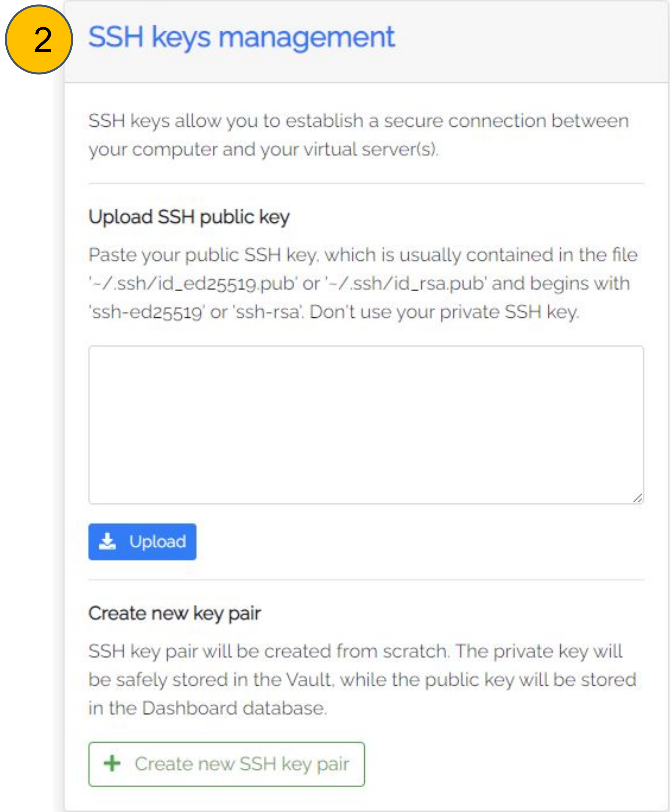

2

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.

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.

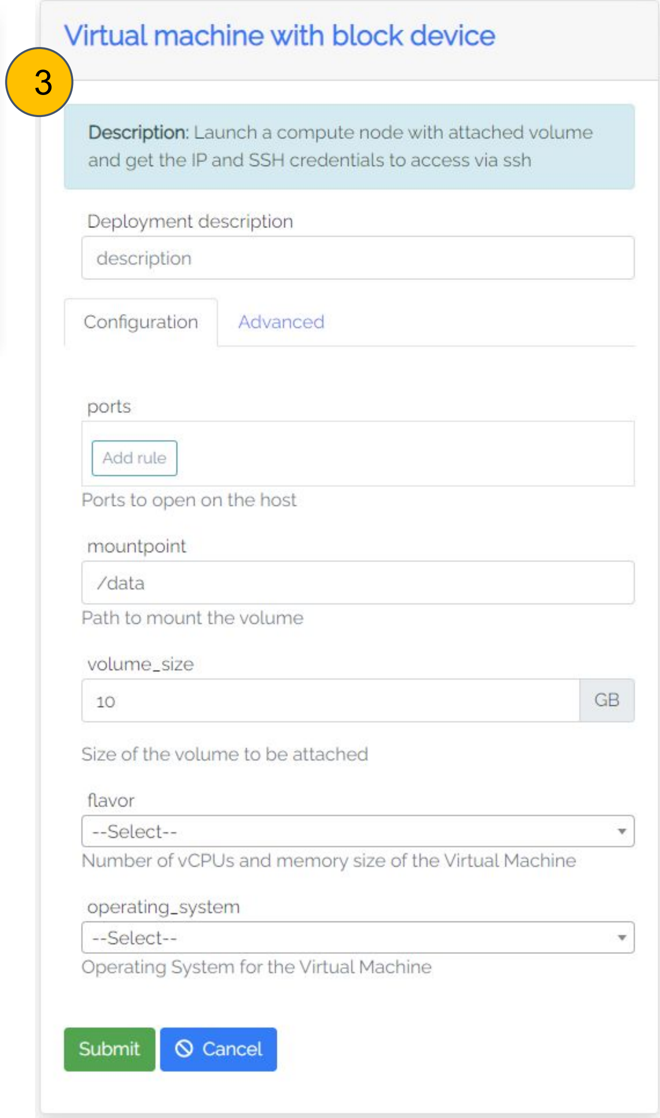

3

Virtual machine with block device

Description: Launch a compute node with attached volume and get the IP and SSH credentials to access via ssh

Deployment description: description

ports: Add rule
Ports to open on the host

mountpoint: /data
Path to mount the volume

volume_size: 10 GB
Size of the volume to be attached

flavor: --Select--
Number of vCPUs and memory size of the Virtual Machine

operating_system: --Select--
Operating System for the Virtual Machine

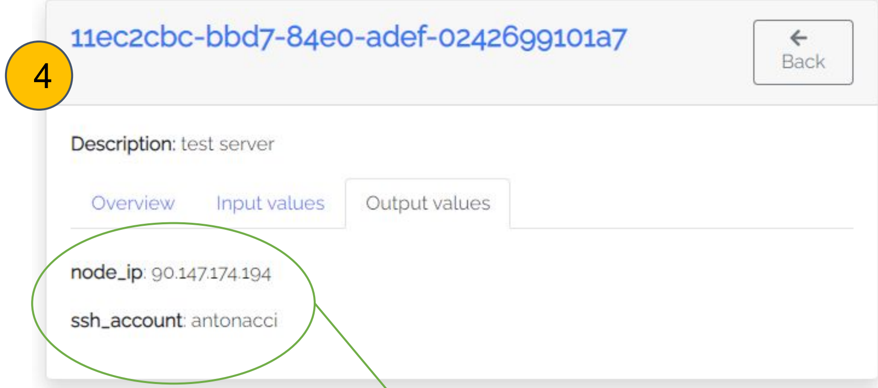

4

11ec2cbc-bbd7-84e0-adef-0242699101a7

Description: test server

node_ip: 90.147.174.194

ssh_account: antonacci

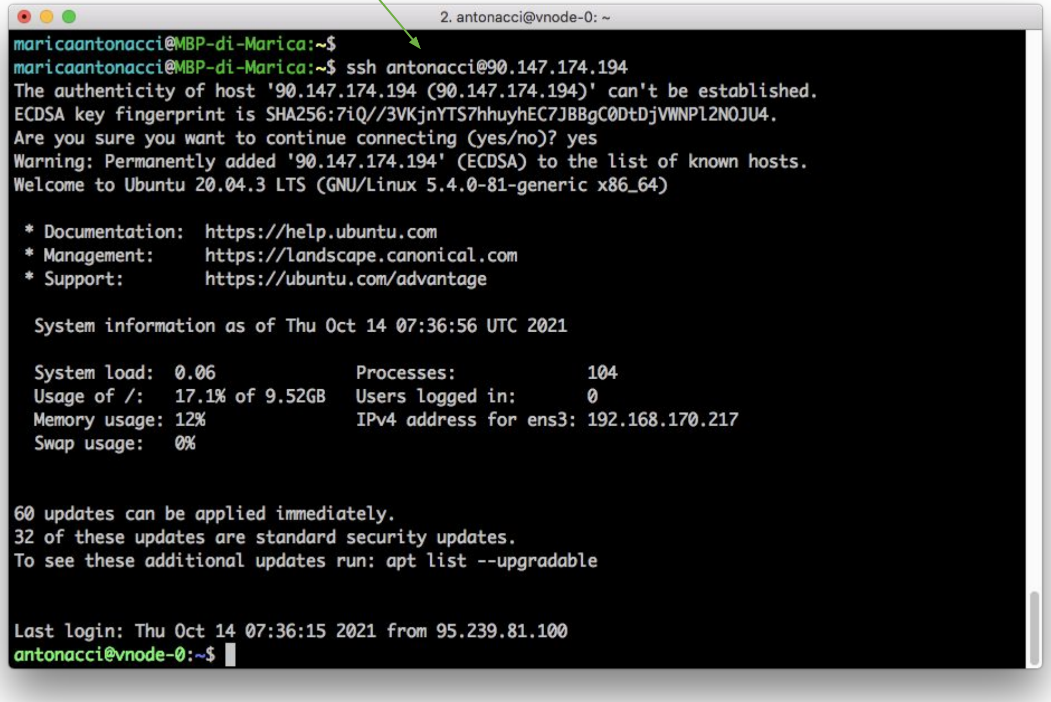

2. antonacci@vnode-0: ~

```
maricaantonacci@MBP-di-Marica:~$
maricaantonacci@MBP-di-Marica:~$ ssh antonacci@90.147.174.194
The authenticity of host '90.147.174.194 (90.147.174.194)' can't be established.
ECDSA key fingerprint is SHA256:7iQ//3VKjnYTS7hhuyhEC7JBBgC0DtDjVWNPl2NOJU4.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '90.147.174.194' (ECDSA) to the list of known hosts.
Welcome to Ubuntu 20.04.3 LTS (GNU/Linux 5.4.0-81-generic x86_64)

 * Documentation:  https://help.ubuntu.com
 * Management:     https://landscape.canonical.com
 * Support:        https://ubuntu.com/advantage

  System information as of Thu Oct 14 07:36:56 UTC 2021

  System load:  0.06              Processes:             104
  Usage of /:   17.1% of 9.52GB   Users logged in:       0
  Memory usage: 12%               IPv4 address for ens3: 192.168.170.217
  Swap usage:   0%


60 updates can be applied immediately.
32 of these updates are standard security updates.
To see these additional updates run: apt list --upgradable


Last login: Thu Oct 14 07:36:15 2021 from 95.239.81.100
antonacci@vnode-0:~$
```

Docker run use-case

How to run a container on INFN Cloud

Configure your dockerized service

The configuration form allows you to customize your deployment.

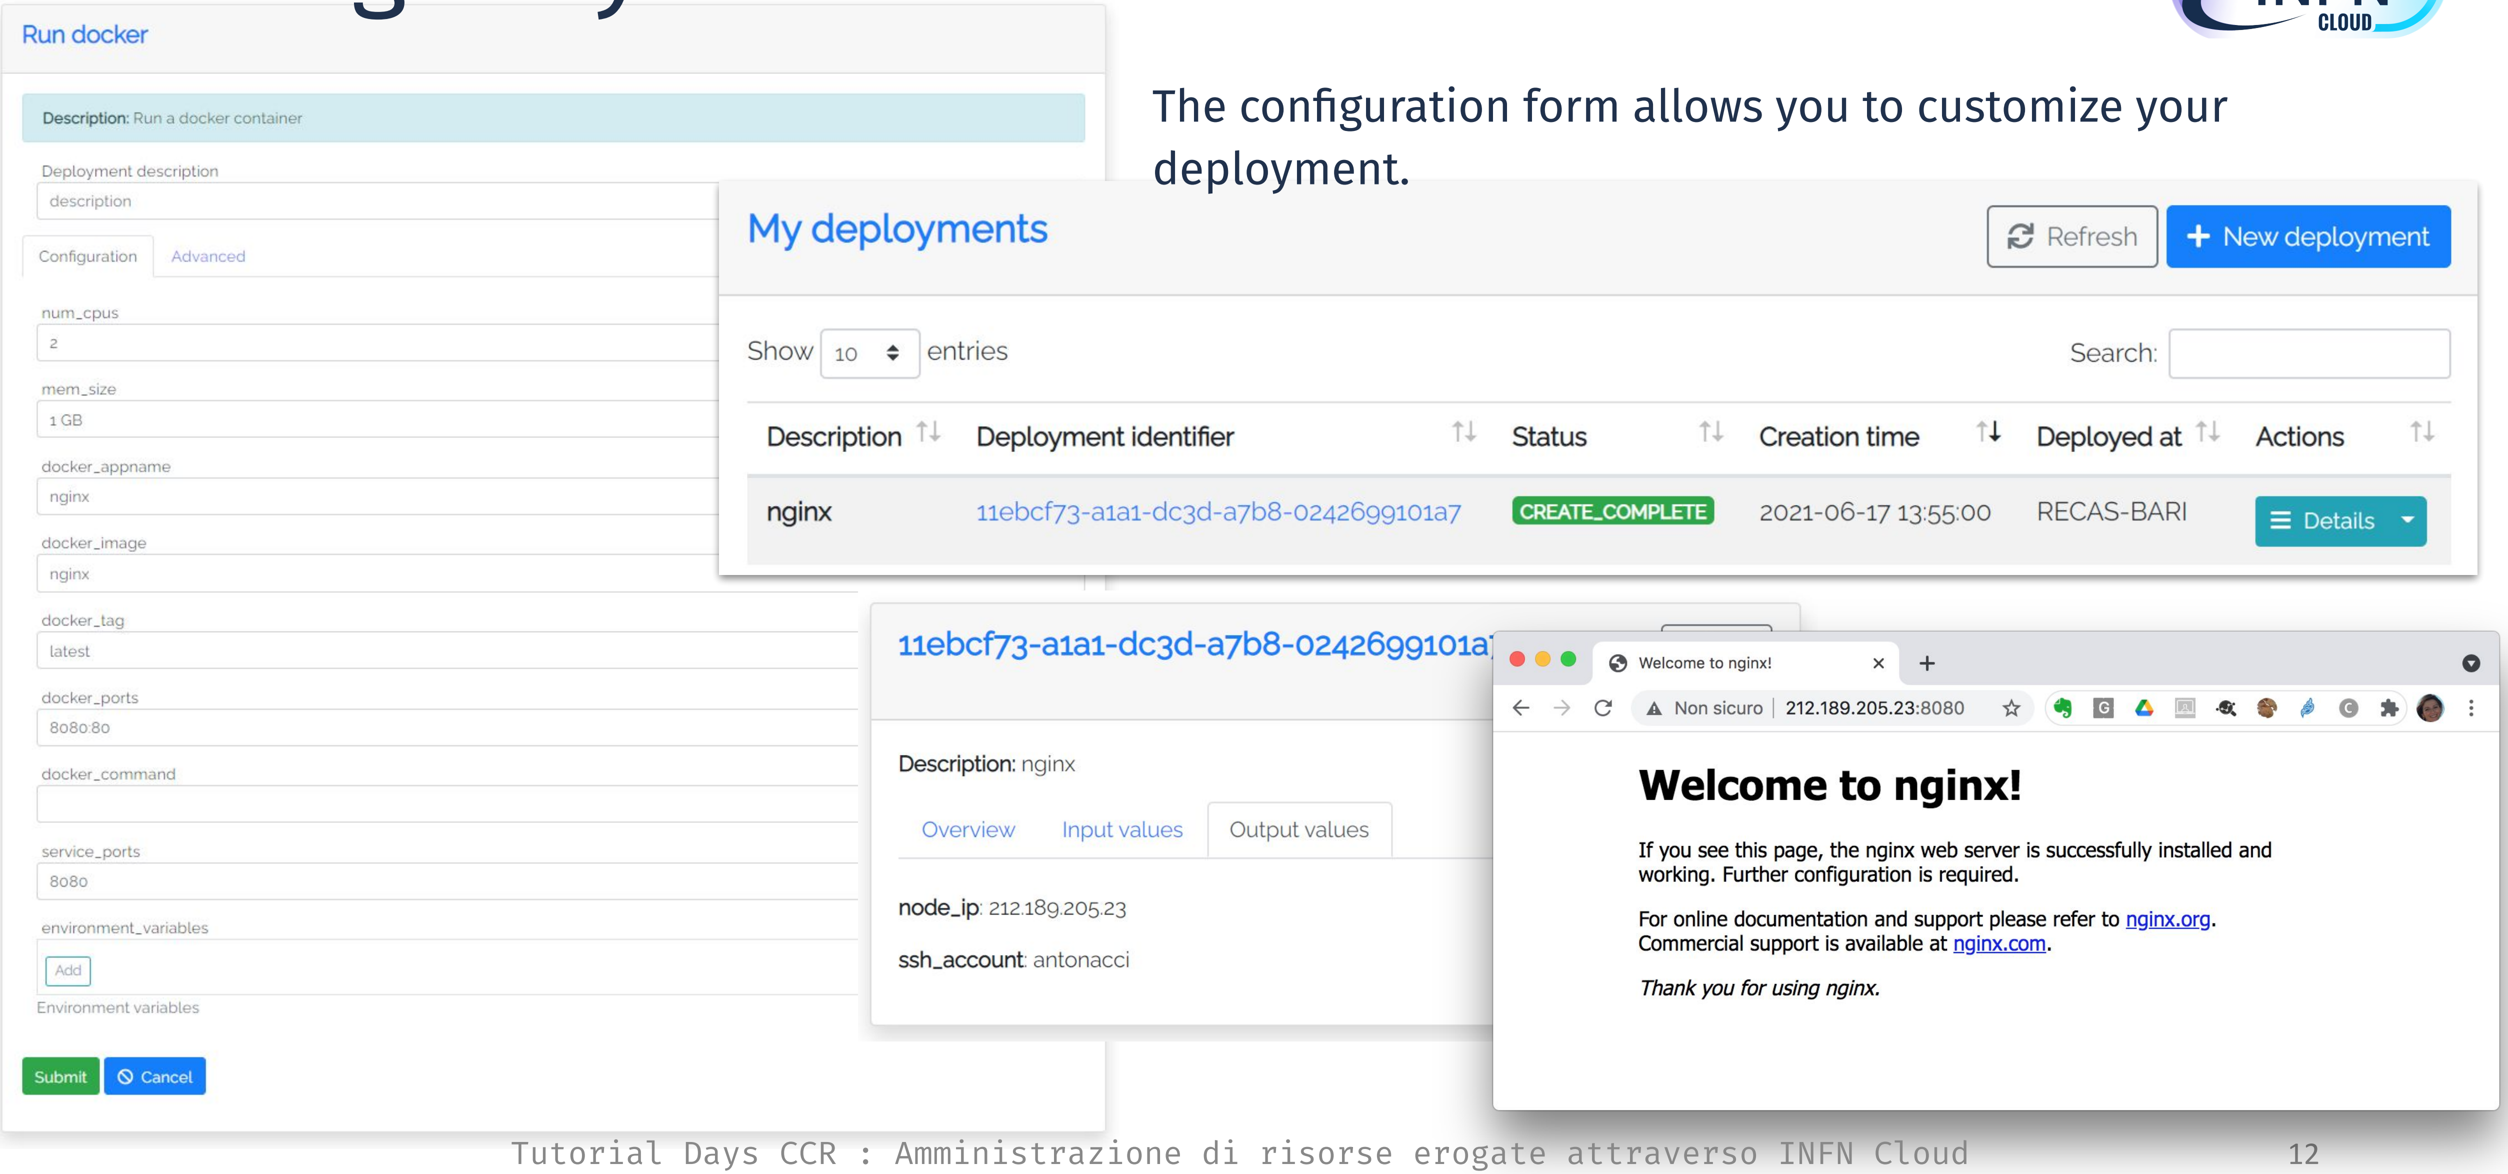

Run docker

Description: Run a docker container

Deployment description
description

Configuration | Advanced

num_cpus
2

mem_size
1 GB

docker_appname
nginx

docker_image
nginx

docker_tag
latest

docker_ports
8080:80

docker_command

service_ports
8080

environment_variables
Add
Environment variables

Submit | Cancel

My deployments

Refresh | + New deployment

Show 10 entries

Search:

Description	Deployment identifier	Status	Creation time	Deployed at	Actions
nginx	11ebcf73-a1a1-dc3d-a7b8-0242699101a7	CREATE_COMPLETE	2021-06-17 13:55:00	RECAS-BARI	Details

11ebcf73-a1a1-dc3d-a7b8-0242699101a7

Description: nginx

Overview | Input values | Output values

node_ip: 212.189.205.23

ssh_account: antonacci

Welcome to nginx!

Non sicuro | 212.189.205.23:8080

Welcome to nginx!

If you see this page, the nginx web server is successfully installed and working. Further configuration is required.

For online documentation and support please refer to nginx.org.
Commercial support is available at nginx.com.

Thank you for using nginx.

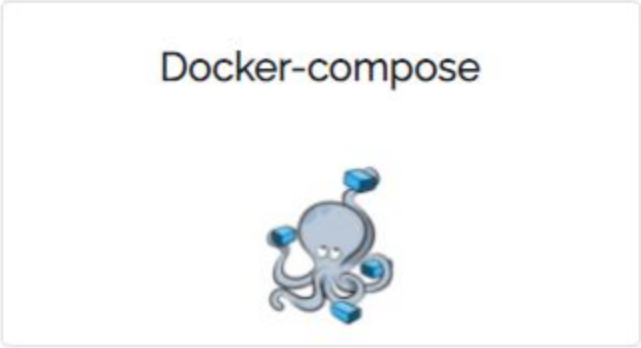

Docker-compose use-case

How to run a docker compose file fetched from a given URL

Configure your service

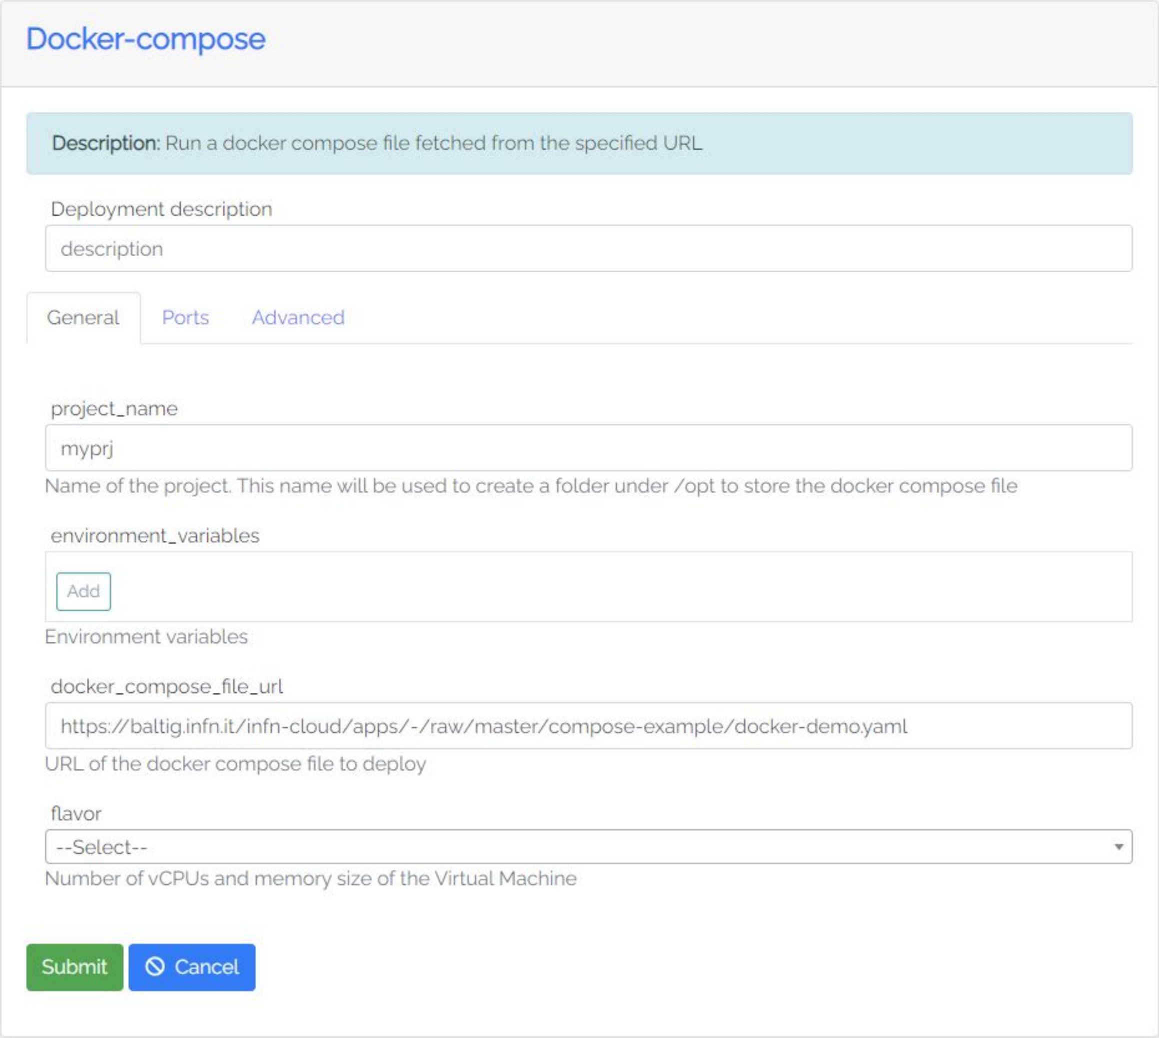

Docker-compose

Description: Run a docker compose file fetched from the specified URL

Deployment description

description

project_name

myprj

Name of the project. This name will be used to create a folder under /opt to store the docker compose file

environment_variables

Environment variables

docker_compose_file_url

https://baltig.infn.it/infn-cloud/apps/-/raw/master/compose-example/docker-demo.yaml

URL of the docker compose file to deploy

flavor

--Select--

Number of vCPUs and memory size of the Virtual Machine

Environment variables management

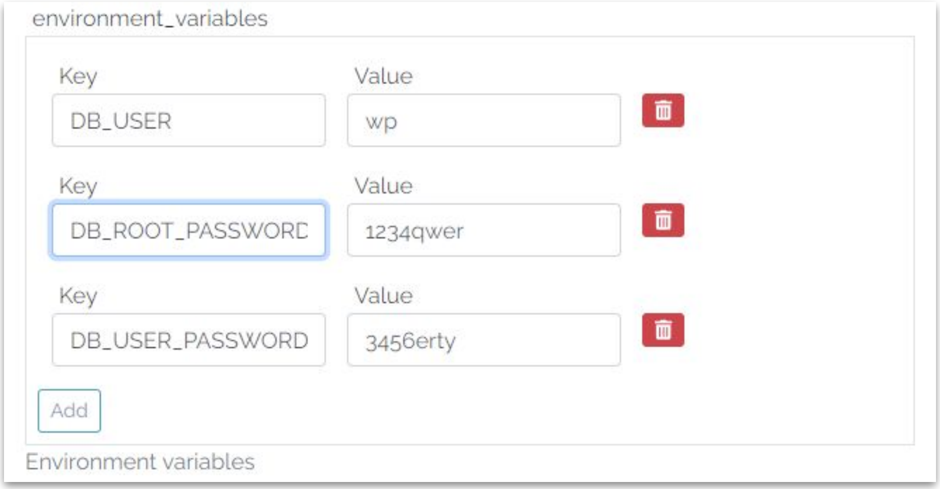

environment_variables

Key	Value
DB_USER	wp
DB_ROOT_PASSWORD	1234qwer
DB_USER_PASSWORD	3456erty

Environment variables

```
services:
 ........
 app:
  depends_on:
   - db
  image: wordpress
  container_name: app
  volumes:
   - wp-content:/var/www/html/wp-content
  environment:
   - WORDPRESS_DB_HOST=db:3306
   - WORDPRESS_DB_USER=${DB_USER}
   - WORDPRESS_DB_PASSWORD=${DB_USER_PASSWORD}
   - VIRTUAL_HOST=wp.${HOST_PUBLIC_IP}.myip.cloud.infn.it
  expose:
   - 80
```

- The special variable *HOST_PUBLIC_IP* is made available by the PaaS system and contains the public IP assigned to the VM
- This env variable can be used as a normal env variable inside the user docker compose file

Ports management

You can define the set of ports that must be automatically opened on the server in order to access your services

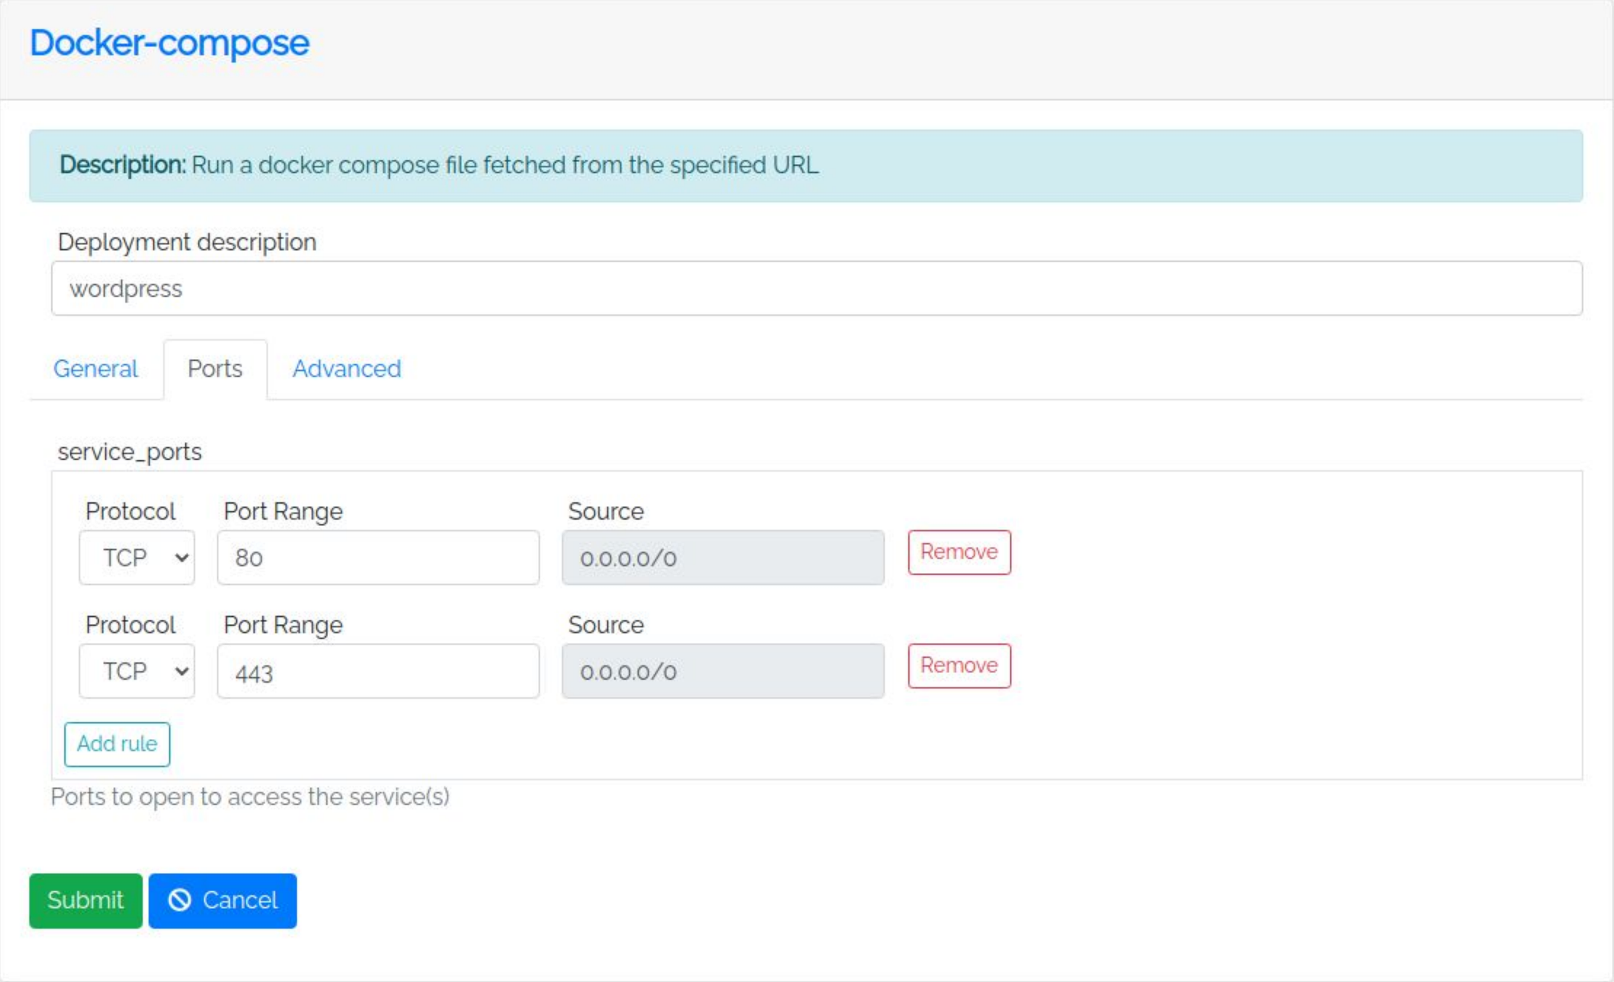

Docker-compose

Description: Run a docker compose file fetched from the specified URL

Deployment description: wordpress

service_ports

Protocol	Port Range	Source	
TCP	80	0.0.0.0/0	Remove
TCP	443	0.0.0.0/0	Remove

Ports to open to access the service(s)

DNS @INFN Cloud

INFN Cloud provides a DNSaaS mechanism that associates a DNS name to each VM public IP

```
$ host wp.90.147.174.132.myip.cloud.infn.it
wp.90.147.174.132.myip.cloud.infn.it has address
90.147.174.132
```

This mechanism is based on xip.io (wildcard DNS) and is exploited for the automatic generation of ssl certificates (e.g. with letsencrypt)

```
services:
  db:
    image: mariadb
    container_name: db
    volumes:
      - db:/var/lib/mysql
    environment:
      - MYSQL_ROOT_PASSWORD=${DB_ROOT_PASSWORD}
      - MYSQL_DATABASE=wordpress
      - MYSQL_USER=${DB_USER}
      - MYSQL_PASSWORD=${DB_USER_PASSWORD}
    expose:
      - 3306
  app:
    depends_on:
      - db
    image: wordpress
    container_name: app
    volumes:
      - wp-content:/var/www/html/wp-content
    environment:
      - WORDPRESS_DB_HOST=db:3306
      - WORDPRESS_DB_USER=${DB_USER}
      - WORDPRESS_DB_PASSWORD=${DB_USER_PASSWORD}
      - VIRTUAL_HOST=wp.${HOST_PUBLIC_IP}.myip.cloud.infn.it
    expose:
      - 80
```

SSL Terminator & Load-balancer

- You can use Traefik as load balancer and SSL terminator. https://traefik.io/traefik/

- Traefik is able to renew letsencrypt certificates

```
services:
 load_balancer:
  image: traefik
  container_name: traefik
  volumes:
   - letsencrypt:/letsencrypt
   - /var/run/docker.sock:/var/run/docker.sock:ro
  ports:
   - "80:80"
   - "443:443"
  command:
   - "--api.insecure=true"
   - "--providers.docker=true"
   - "--providers.docker.exposedbydefault=false"
   - "--entrypoints.web.address=:80"
   - "--entrypoints.websecure.address=:443"
   - "--certificatesresolvers.myhttpchallenge.acme.httpchallenge=true"
   -
"--certificatesresolvers.myhttpchallenge.acme.httpchallenge.entrypoint=w
eb"
   -
"--certificatesresolvers.myhttpchallenge.acme.email=${CONTACT_EMAIL}"
   -
"--certificatesresolvers.myhttpchallenge.acme.storage=/letsencrypt/acme.
json"
```

Traefik configuration

Traefik is automatically configured through the labels* exposed by the containers

(*) “A label is a **key=value** *pair that applies metadata to a container.”*

```
services:
  app:
    depends_on:
      - db
    image: wordpress
    container_name: app
    volumes:
      - wp-content:/var/www/html/wp-content
    environment:
      - WORDPRESS_DB_HOST=db:3306
      - WORDPRESS_DB_USER=${DB_USER}
      - WORDPRESS_DB_PASSWORD=${DB_USER_PASSWORD}
      - VIRTUAL_HOST=wp.${HOST_PUBLIC_IP}.myip.cloud.infn.it
    expose:
      - 80
    labels:
      - "traefik.enable=true"
      - "traefik.http.middlewares.app-redirect-ssl.redirectscheme.scheme=https"
      - "traefik.http.routers.app-nossl.middlewares=app-redirect-ssl"
      -
"traefik.http.routers.app-nossl.rule=Host(`wp.${HOST_PUBLIC_IP}.myip.cloud.infn.
it`)"
      - "traefik.http.routers.app-nossl.entrypoints=web"
      -
"traefik.http.routers.app.rule=Host(`wp.${HOST_PUBLIC_IP}.myip.cloud.infn.it`)"
      - "traefik.http.routers.app.entrypoints=websecure"
      - "traefik.http.routers.app.tls.certresolver=myhttpchallenge"
      - "traefik.http.routers.app.tls=true"
```

Docker compose example

https://baltig.infn.it/infn-cloud/apps/-/blob/master/compose-example/docker-demo.yaml

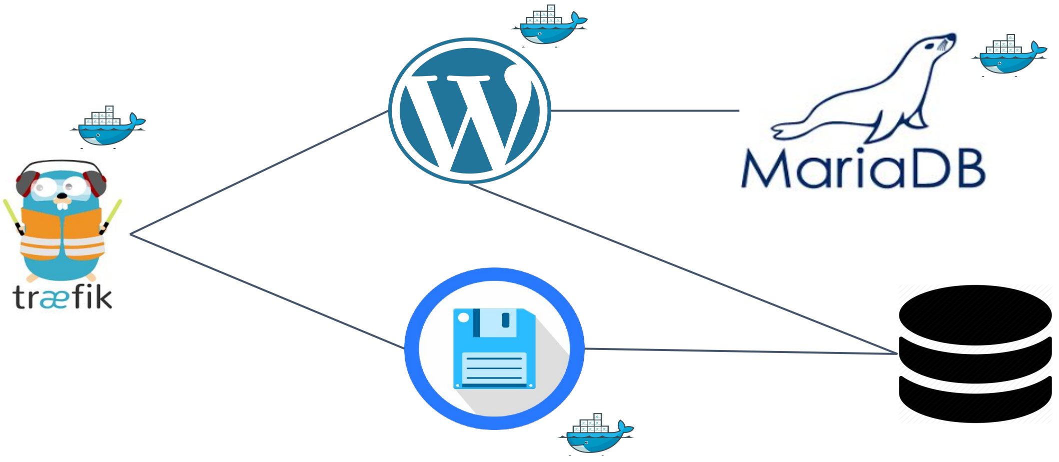

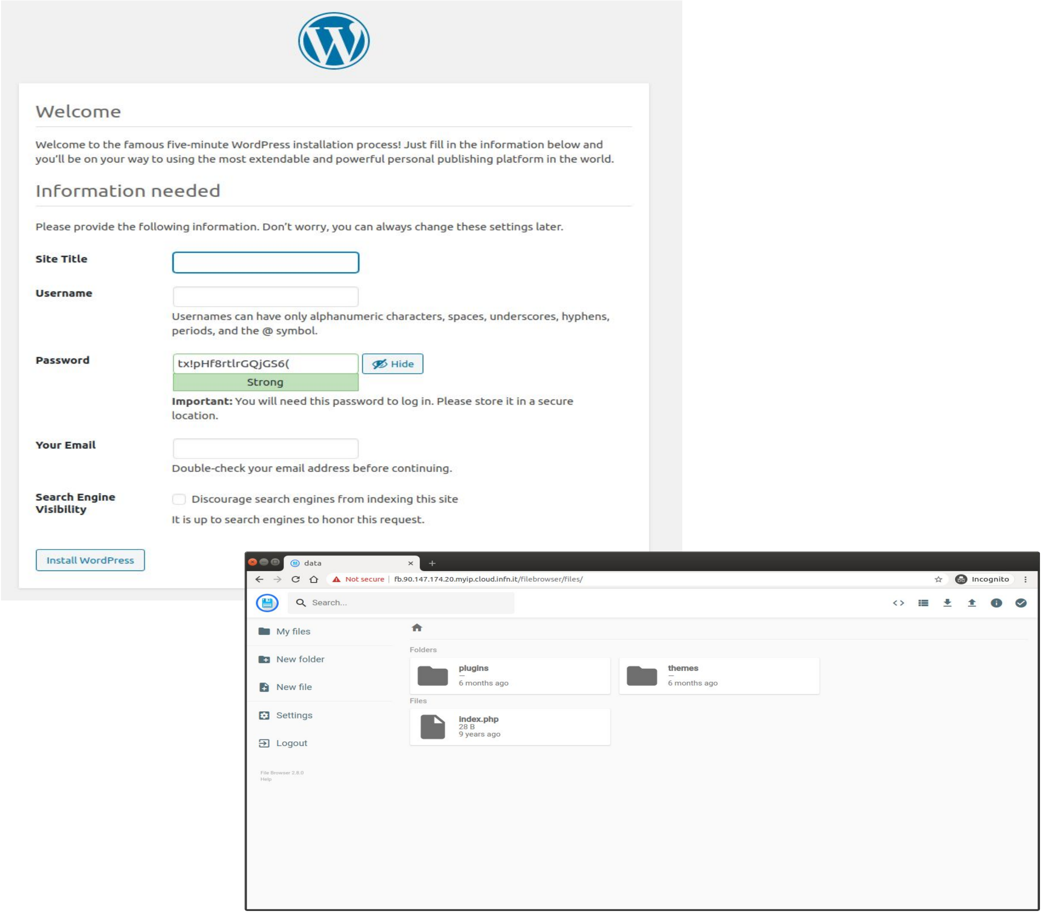

Welcome

Welcome to the famous five-minute WordPress installation process! Just fill in the information below and you'll be on your way to using the most extendable and powerful personal publishing platform in the world.

Information needed

Please provide the following information. Don't worry, you can always change these settings later.

Site Title

Username

Usernames can have only alphanumeric characters, spaces, underscores, hyphens, periods, and the @ symbol.

Password: tx!pHf8rtlrGQjGS6(

Strong

Important: You will need this password to log in. Please store it in a secure location.

Your Email

Double-check your email address before continuing.

Search Engine Visibility

□ Discourage search engines from indexing this site

It is up to search engines to honor this request.

Not secure | fb.90.147.174.20.myip.cloud.infn.it/filebrowser/files/

Folders

plugins — 6 months ago

themes — 6 months ago

Files

index.php 28 B 9 years ago

File Browser 2.8.0

Sync&Share aaS

How to deploy a cloud storage service based on Owncloud

Sync&Share aaS

- ownCloud is automatically configured to use the **INFN Cloud Object Storage** as main storage backend
- Rclone can be used for programmatic access to user data, including remote mount and folder sync
- embedded **automated DB and configuration backup (duplicati)**
- **embedded pre-configured monitoring system** with alert notifications (nagios)

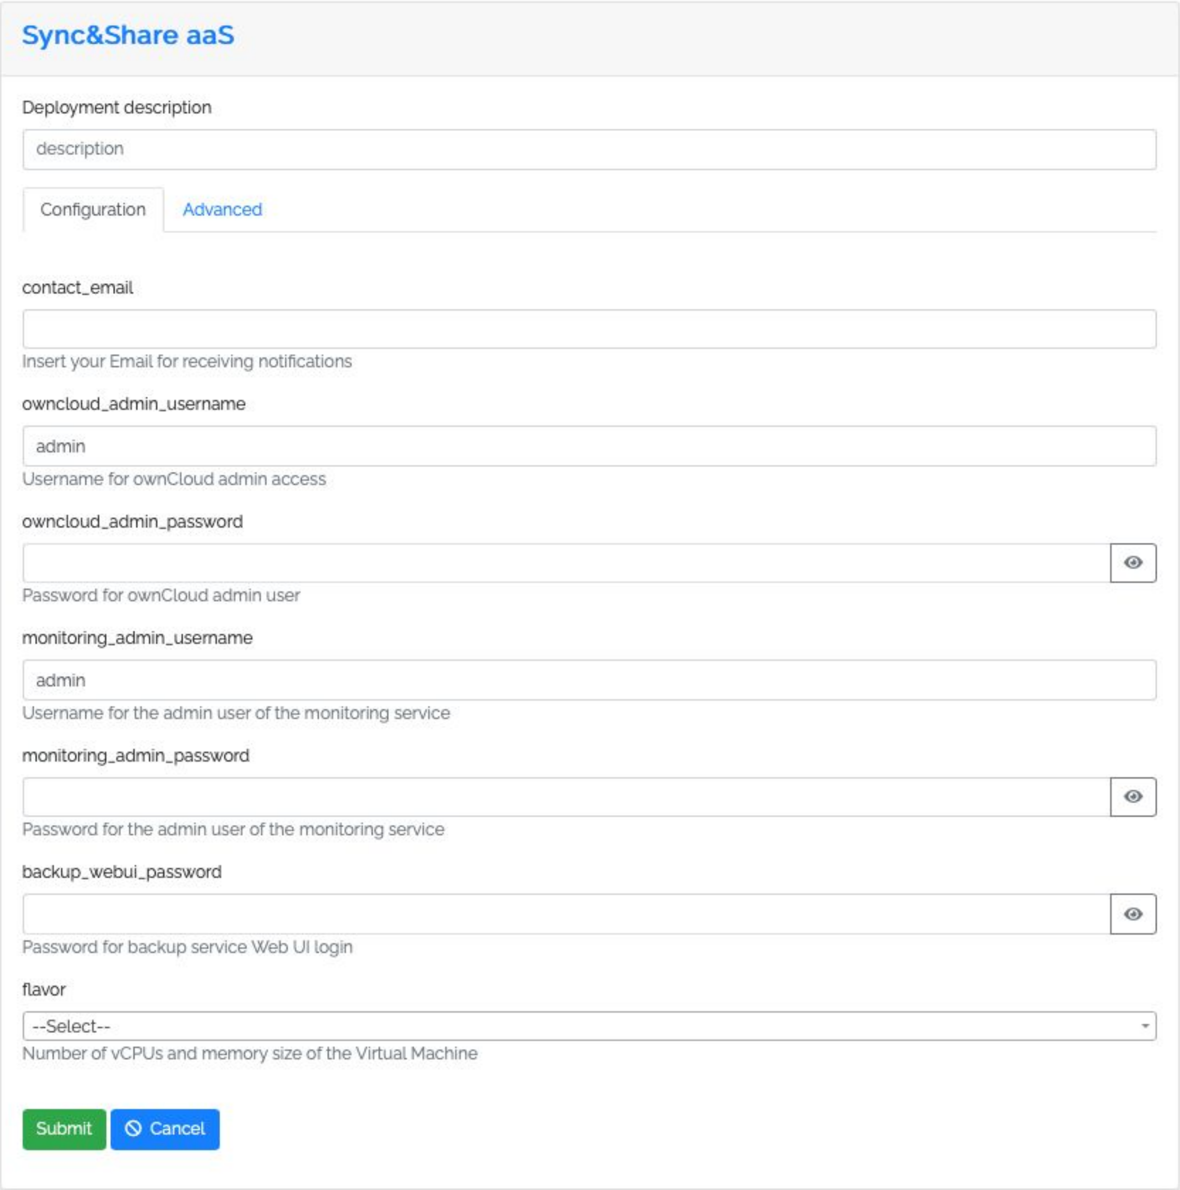

Sync&Share aaS

Deployment description

description

contact_email

Insert your Email for receiving notifications

owncloud_admin_username

admin

Username for ownCloud admin access

owncloud_admin_password

Password for ownCloud admin user

monitoring_admin_username

admin

Username for the admin user of the monitoring service

monitoring_admin_password

Password for the admin user of the monitoring service

backup_webui_password

Password for backup service Web UI login

flavor

--Select--

Number of vCPUs and memory size of the Virtual Machine

Deployment outputs

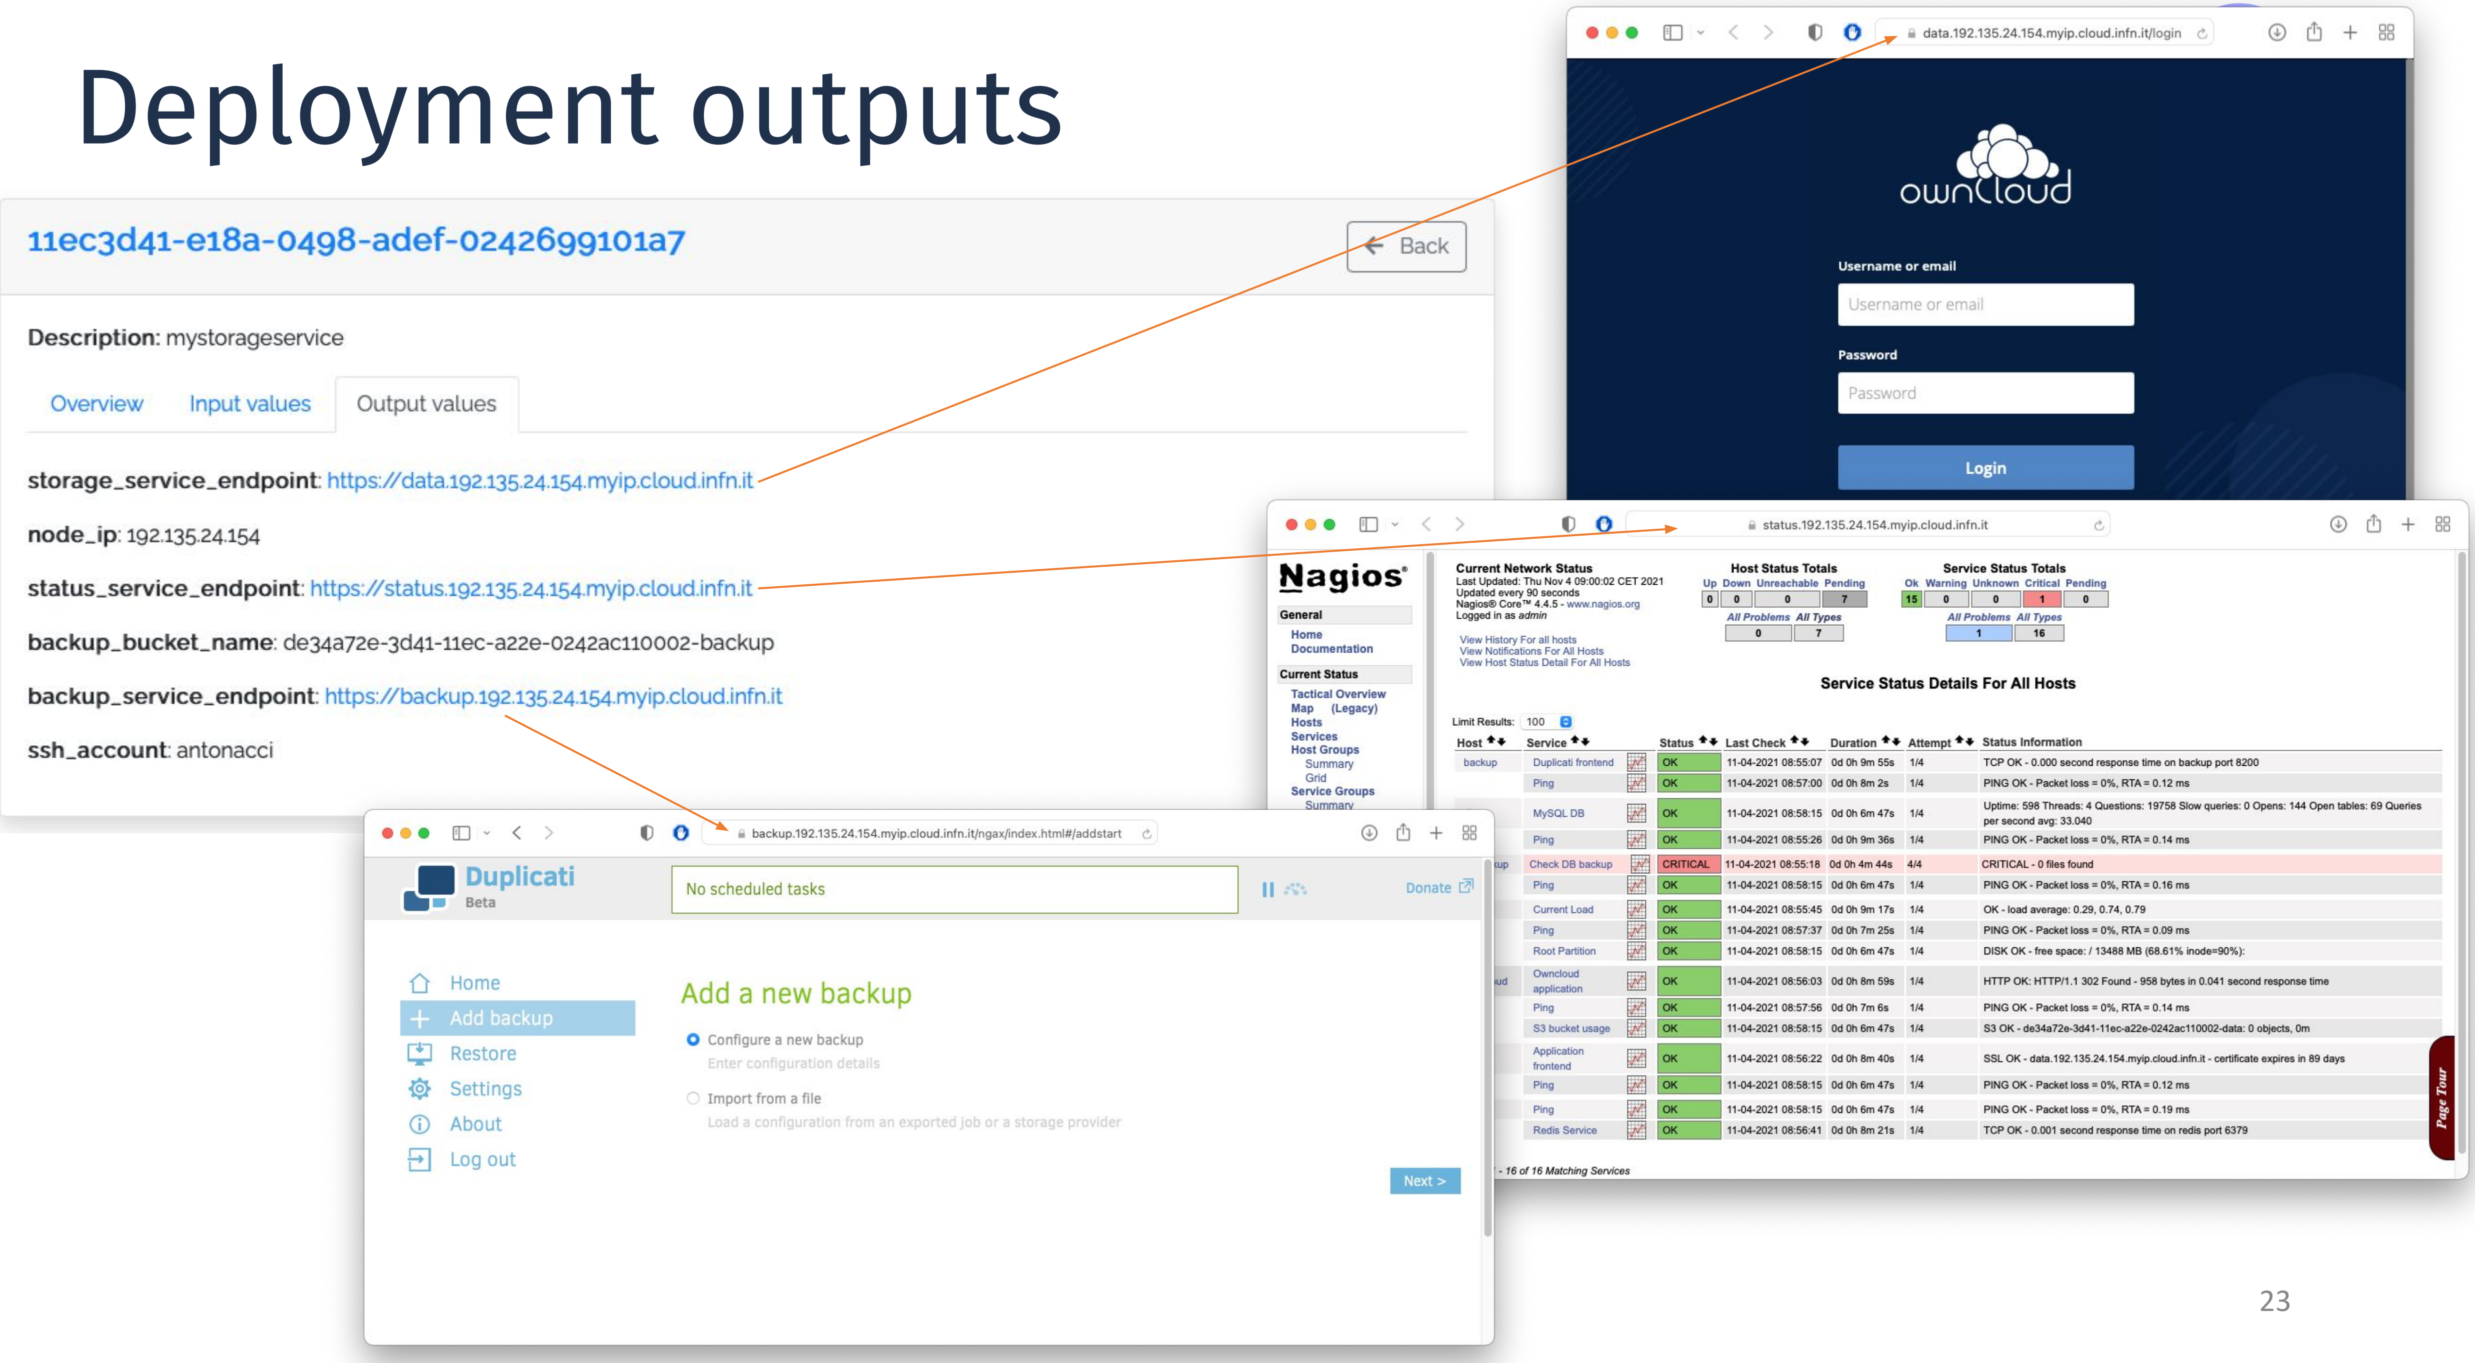

11ec3d41-e18a-0498-adef-0242699101a7

Description: mystorageservice

storage_service_endpoint: https://data.192.135.24.154.myip.cloud.infn.it

node_ip: 192.135.24.154

status_service_endpoint: https://status.192.135.24.154.myip.cloud.infn.it

backup_bucket_name: de34a72e-3d41-11ec-a22e-0242ac110002-backup

backup_service_endpoint: https://backup.192.135.24.154.myip.cloud.infn.it

ssh_account: antonacci

data.192.135.24.154.myip.cloud.infn.it/login

Username or email: Username or email

Password: Password

status.192.135.24.154.myip.cloud.infn.it

Current Network Status
Last Updated: Thu Nov 4 09:00:02 CET 2021
Updated every 90 seconds
Nagios® Core™ 4.4.5 - www.nagios.org
Logged in as *admin*

View History For all hosts
View Notifications For All Hosts
View Host Status Detail For All Hosts

Host Status Totals

Up	Down	Unreachable	Pending
0	0	0	7

All Problems	*All Types*
0	7

Service Status Totals

Ok	Warning	Unknown	Critical	Pending
15	0	0	1	0

All Problems	*All Types*
1	16

Service Status Details For All Hosts

Limit Results: 100

Host	Service	Status	Last Check	Duration	Attempt	Status Information
backup	Duplicati frontend	OK	11-04-2021 08:55:07	0d 0h 9m 55s	1/4	TCP OK - 0.000 second response time on backup port 8200
	Ping	OK	11-04-2021 08:57:00	0d 0h 8m 2s	1/4	PING OK - Packet loss = 0%, RTA = 0.12 ms
	MySQL DB	OK	11-04-2021 08:58:15	0d 0h 6m 47s	1/4	Uptime: 598 Threads: 4 Questions: 19758 Slow queries: 0 Opens: 144 Open tables: 69 Queries per second avg: 33.040
	Ping	OK	11-04-2021 08:55:26	0d 0h 9m 36s	1/4	PING OK - Packet loss = 0%, RTA = 0.14 ms
kup	Check DB backup	CRITICAL	11-04-2021 08:55:18	0d 0h 4m 44s	4/4	CRITICAL - 0 files found
	Ping	OK	11-04-2021 08:58:15	0d 0h 6m 47s	1/4	PING OK - Packet loss = 0%, RTA = 0.16 ms
	Current Load	OK	11-04-2021 08:55:45	0d 0h 9m 17s	1/4	OK - load average: 0.29, 0.74, 0.79
	Ping	OK	11-04-2021 08:57:37	0d 0h 7m 25s	1/4	PING OK - Packet loss = 0%, RTA = 0.09 ms
	Root Partition	OK	11-04-2021 08:58:15	0d 0h 6m 47s	1/4	DISK OK - free space: / 13488 MB (68.61% inode=90%):
ud	Owncloud application	OK	11-04-2021 08:56:03	0d 0h 8m 59s	1/4	HTTP OK: HTTP/1.1 302 Found - 958 bytes in 0.041 second response time
	Ping	OK	11-04-2021 08:57:56	0d 0h 7m 6s	1/4	PING OK - Packet loss = 0%, RTA = 0.14 ms
	S3 bucket usage	OK	11-04-2021 08:58:15	0d 0h 6m 47s	1/4	S3 OK - de34a72e-3d41-11ec-a22e-0242ac110002-data: 0 objects, 0m
	Application frontend	OK	11-04-2021 08:56:22	0d 0h 8m 40s	1/4	SSL OK - data.192.135.24.154.myip.cloud.infn.it - certificate expires in 89 days
	Ping	OK	11-04-2021 08:58:15	0d 0h 6m 47s	1/4	PING OK - Packet loss = 0%, RTA = 0.12 ms
	Ping	OK	11-04-2021 08:58:15	0d 0h 6m 47s	1/4	PING OK - Packet loss = 0%, RTA = 0.19 ms
	Redis Service	OK	11-04-2021 08:56:41	0d 0h 8m 21s	1/4	TCP OK - 0.001 second response time on redis port 6379

- 16 of 16 Matching Services

backup.192.135.24.154.myip.cloud.infn.it/ngax/index.html#/addstart

No scheduled tasks

Add a new backup

● Configure a new backup
Enter configuration details

○ Import from a file
Load a configuration from an exported job or a storage provider

INFN Cloud Object Storage

Based on Openstack Swift, the service is available for all INFN end users

- data replicated on two DCs (CNAF and Bari)
- Swift & S3 APIs
  - tested integration with several tools
    - Duplicati, OwnCloud, Rclone, Minio, etc.
- support for IAM-authentication (via keystone)
- High level user interface via Minio GW
  - extended ACLs via OPA (Open Policy Agent) + IAM

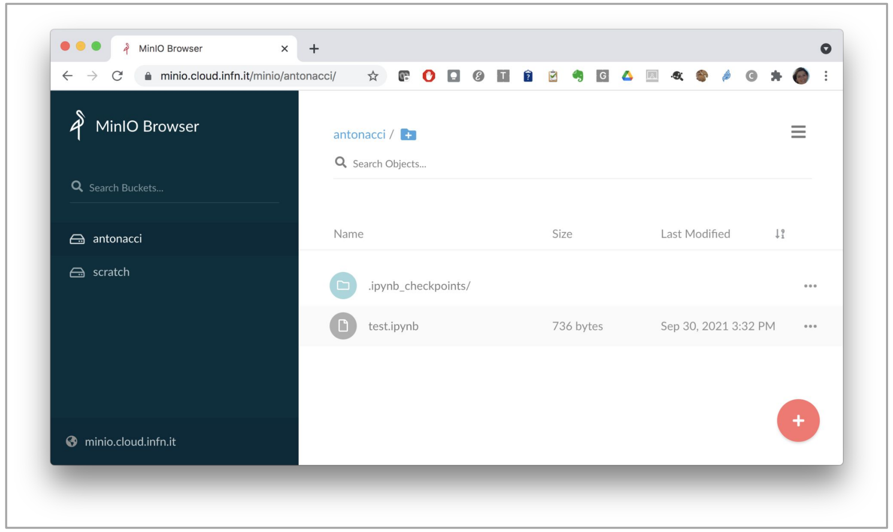

Name	Size	Last Modified
.ipynb_checkpoints/		
test.ipynb	736 bytes	Sep 30, 2021 3:32 PM

The high-level services offered via PaaS exploit this storage service in a transparent way for the user

Apache Mesos cluster

Mesos use-case

How to deploy a complete Mesos cluster

Cluster architecture

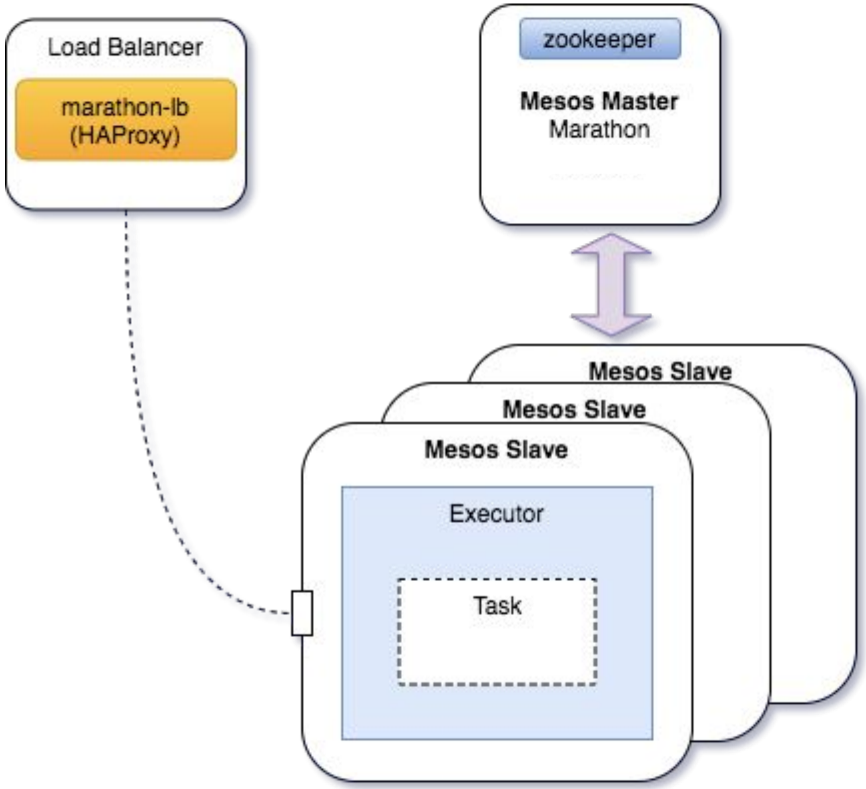

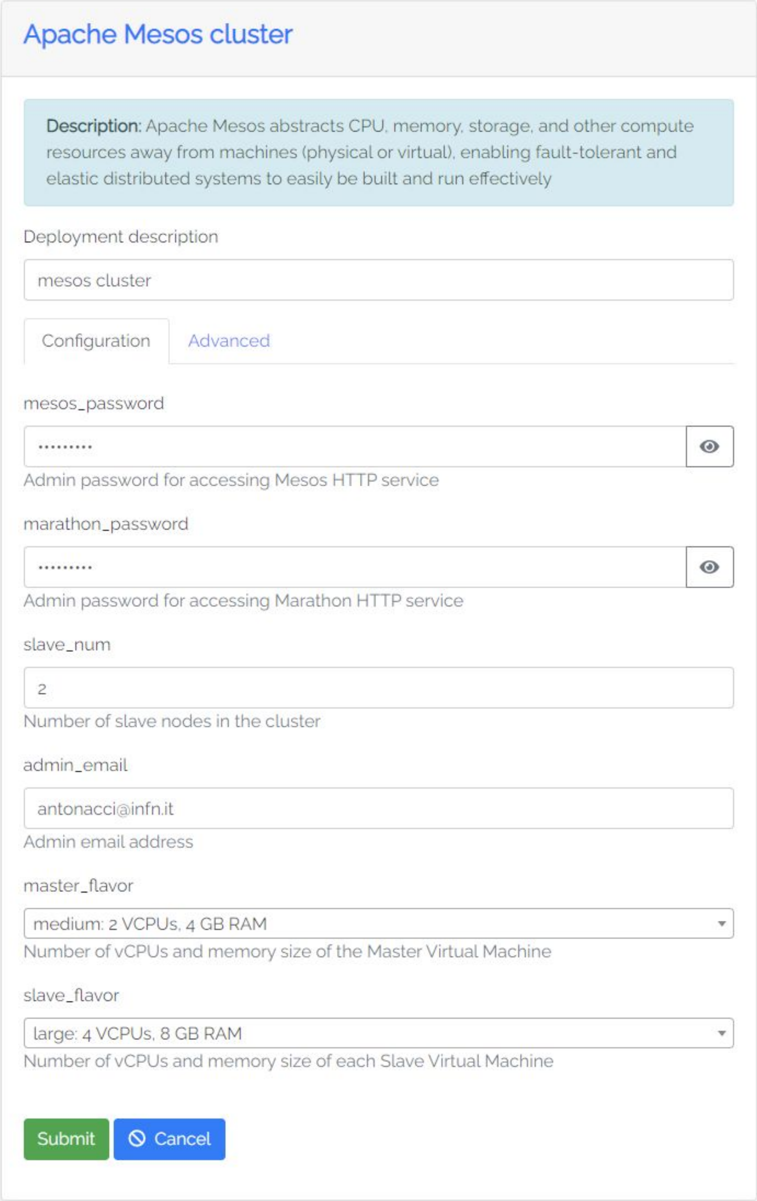

Apache Mesos cluster

Description: Apache Mesos abstracts CPU, memory, storage, and other compute resources away from machines (physical or virtual), enabling fault-tolerant and elastic distributed systems to easily be built and run effectively

Deployment description: mesos cluster

mesos_password: ·········

Admin password for accessing Mesos HTTP service

marathon_password: ·········

Admin password for accessing Marathon HTTP service

slave_num: 2

Number of slave nodes in the cluster

admin_email: antonacci@infn.it

Admin email address

master_flavor: medium: 2 VCPUs, 4 GB RAM

Number of vCPUs and memory size of the Master Virtual Machine

slave_flavor: large: 4 VCPUs, 8 GB RAM

Number of vCPUs and memory size of each Slave Virtual Machine

Deployment outputs

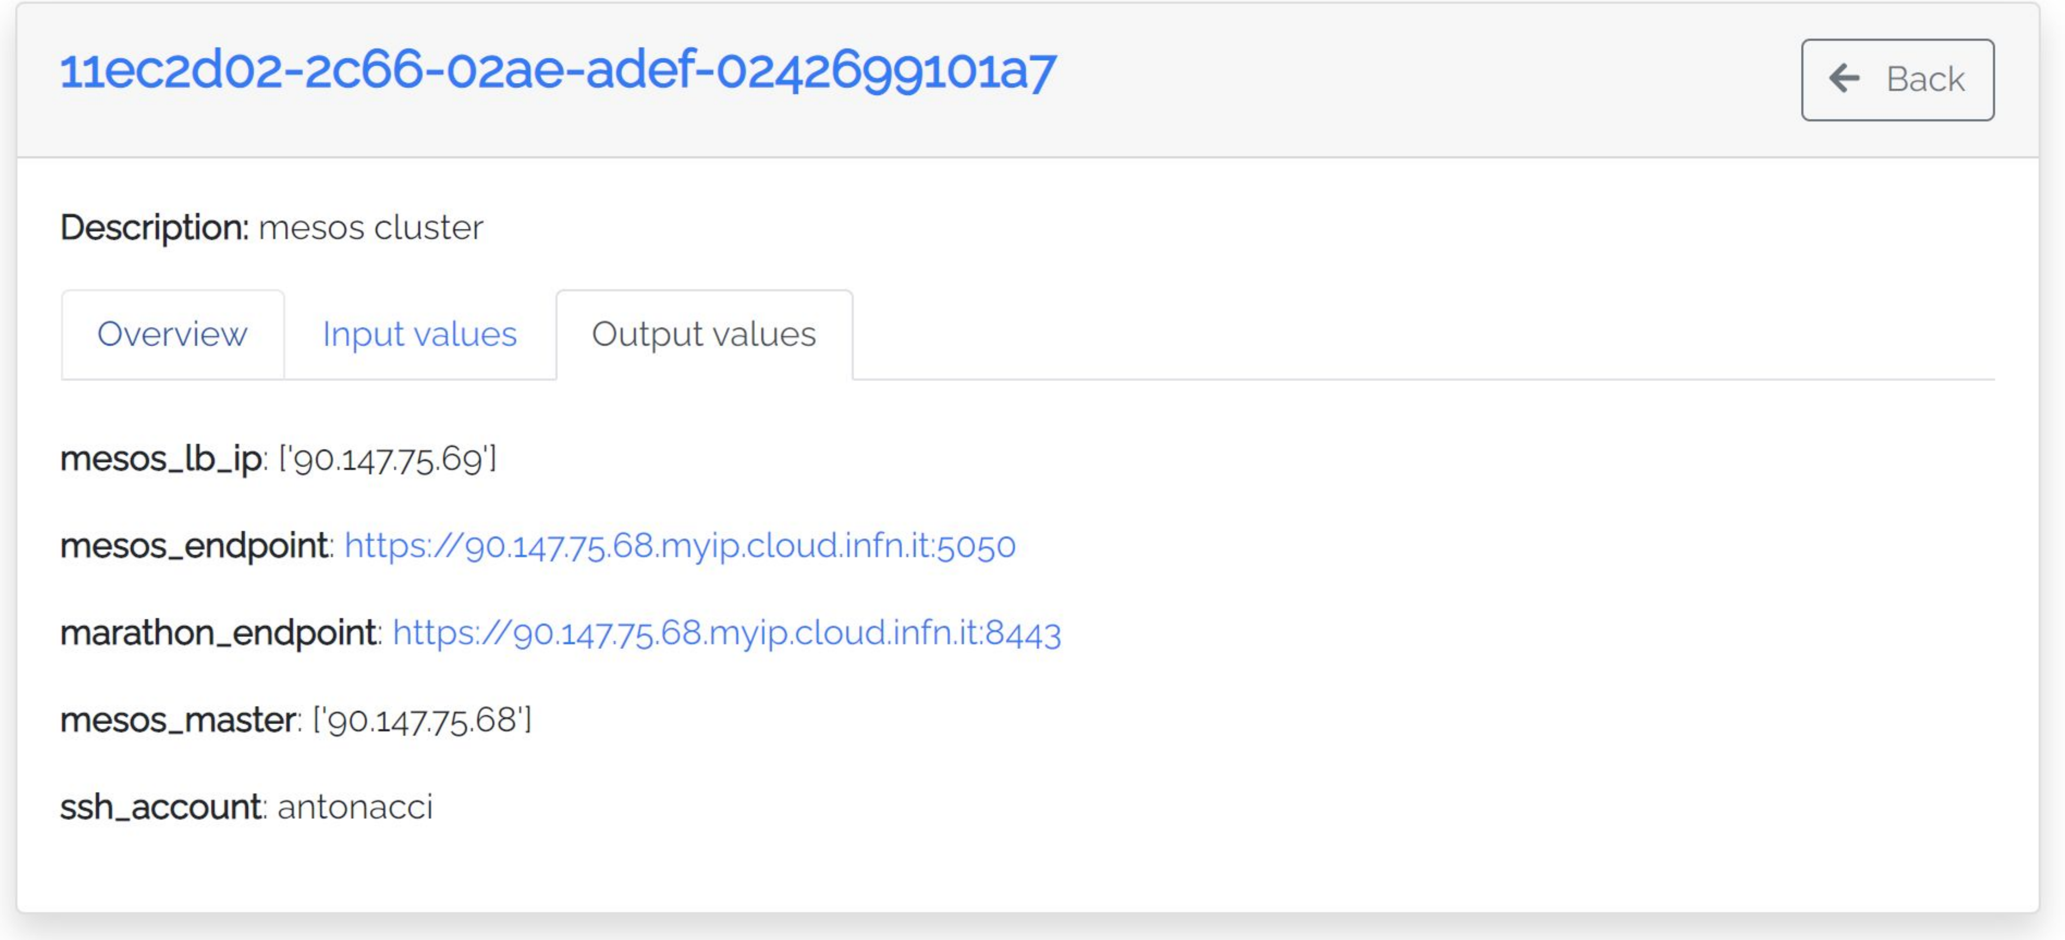

11ec2d02-2c66-02ae-adef-0242699101a7

Description: mesos cluster

mesos_lb_ip: ['90.147.75.69']

mesos_endpoint: https://90.147.75.68.myip.cloud.infn.it:5050

marathon_endpoint: https://90.147.75.68.myip.cloud.infn.it:8443

mesos_master: ['90.147.75.68']

ssh_account: antonacci

Mesos/Marathon

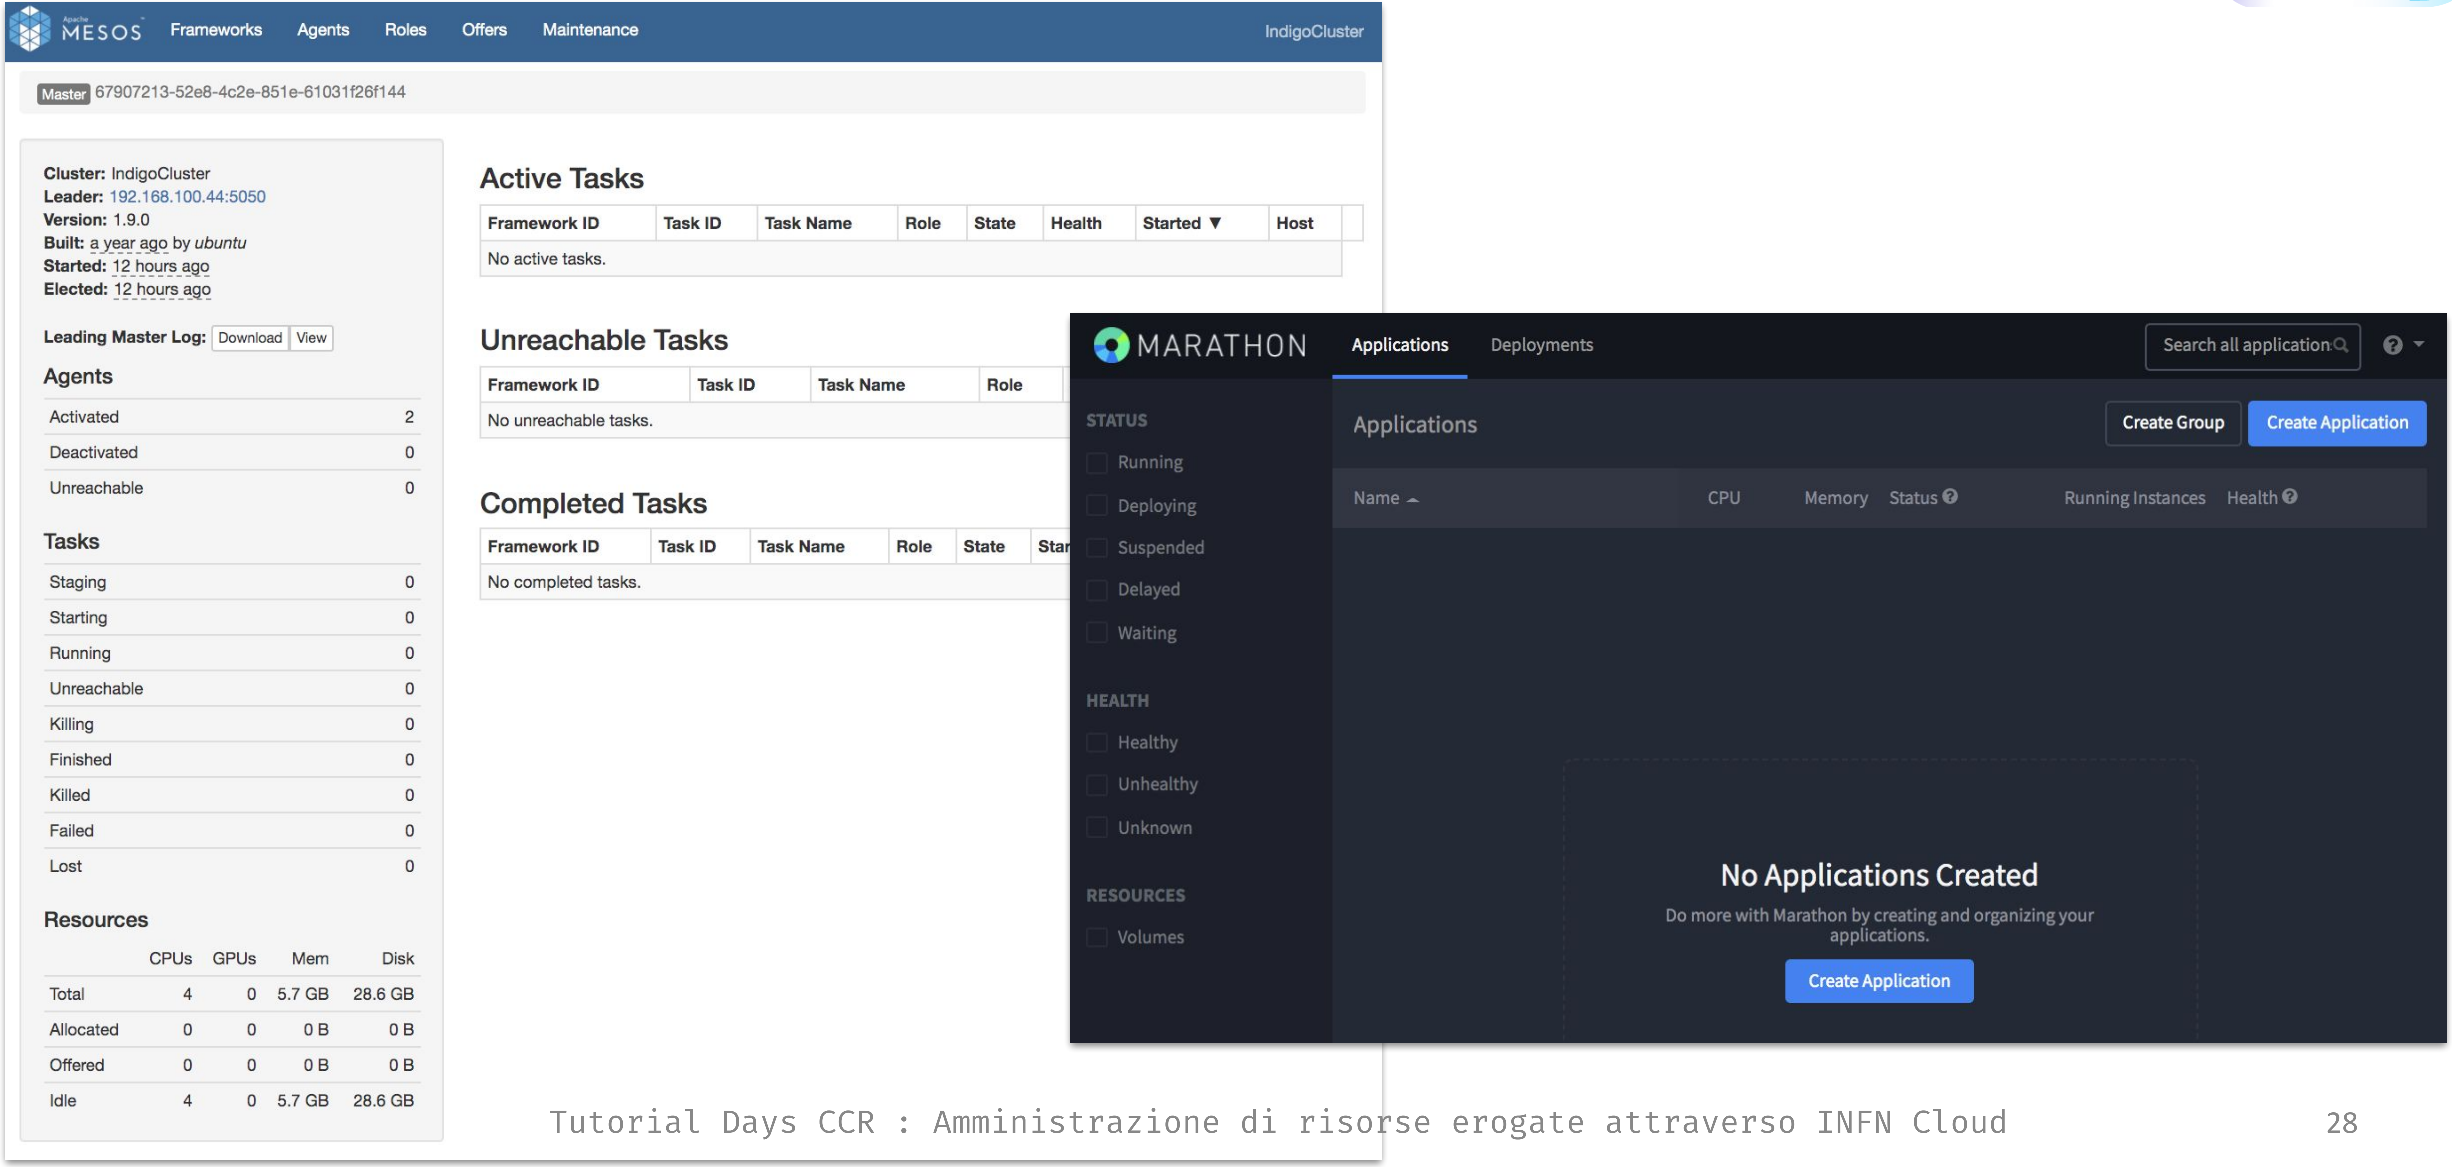

Master 67907213-52e8-4c2e-851e-61031f26f144

Cluster: IndigoCluster
Leader: 192.168.100.44:5050
Version: 1.9.0
Built: a year ago by *ubuntu*
Started: 12 hours ago
Elected: 12 hours ago

Leading Master Log: Download View

Agents

Activated	2
Deactivated	0
Unreachable	0

Tasks

Staging	0
Starting	0
Running	0
Unreachable	0
Killing	0
Finished	0
Killed	0
Failed	0
Lost	0

Resources

	CPUs	GPUs	Mem	Disk
Total	4	0	5.7 GB	28.6 GB
Allocated	0	0	0 B	0 B
Offered	0	0	0 B	0 B
Idle	4	0	5.7 GB	28.6 GB

Active Tasks

Framework ID	Task ID	Task Name	Role	State	Health	Started ▼	Host
No active tasks.							

Unreachable Tasks

Framework ID	Task ID	Task Name	Role
No unreachable tasks.			

Completed Tasks

Framework ID	Task ID	Task Name	Role	State	Star
No completed tasks.					

STATUS
□ Running
□ Deploying
□ Suspended
□ Delayed
□ Waiting

HEALTH
□ Healthy
□ Unhealthy
□ Unknown

RESOURCES
□ Volumes

Applications

Create Group · Create Application

Name ▲	CPU	Memory	Status ?	Running Instances	Health ?

No Applications Created

Do more with Marathon by creating and organizing your applications.

Create Application

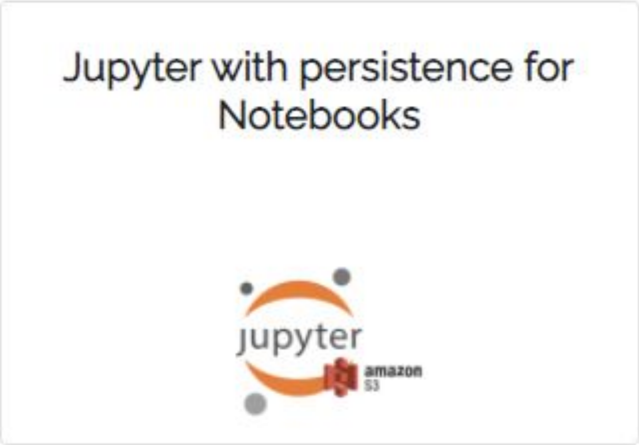

Jupyter with persistence for Notebooks

Jupyter with persistence for Notebooks

How to deploy Jupyter on a single VM enabling Notebooks persistence

Configure your service

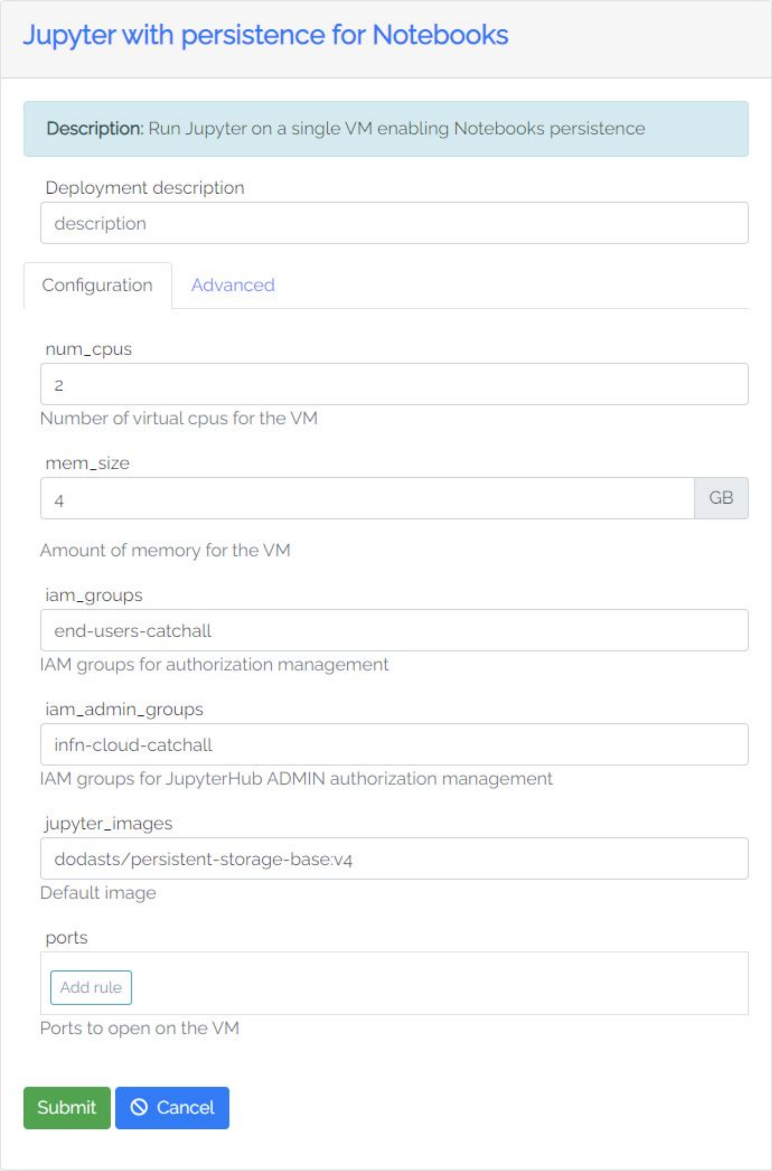

Jupyter with persistence for Notebooks

Description: Run Jupyter on a single VM enabling Notebooks persistence

Deployment description: description

num_cpus: 2
Number of virtual cpus for the VM

mem_size: 4 GB
Amount of memory for the VM

iam_groups: end-users-catchall
IAM groups for authorization management

iam_admin_groups: infn-cloud-catchall
IAM groups for JupyterHub ADMIN authorization management

jupyter_images: dodasts/persistent-storage-base:v4
Default image

ports
Ports to open on the VM

Jupyterhub is automatically installed and configured in the Virtual Machine

- **The integration with INFN Cloud IAM allows authorized users to login and spawn their containerized jupyter server starting from**

  The available default jupyter image

  A customized docker image (available locally or shared on a public registry)

- **The integration with the INFN Cloud storage provides persistence for user notebooks**

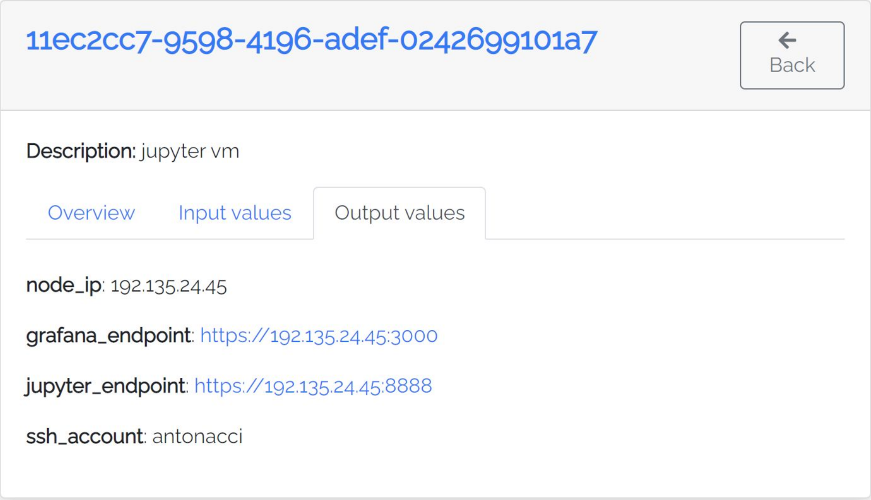

11ec2cc7-9598-4196-adef-0242699101a7

Description: jupyter vm

node_ip: 192.135.24.45

grafana_endpoint: https://192.135.24.45:3000

jupyter_endpoint: https://192.135.24.45:8888

ssh_account: antonacci

A look inside the machine

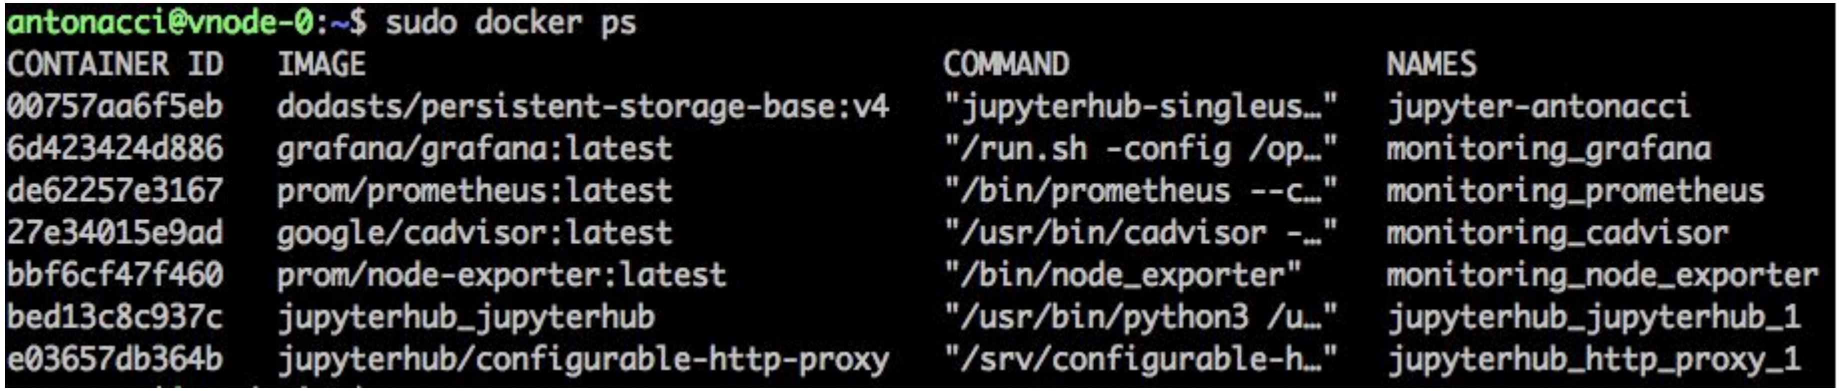

```
antonacci@vnode-0:~$ sudo docker ps
CONTAINER ID   IMAGE                                 COMMAND                   NAMES
00757aa6f5eb   dodasts/persistent-storage-base:v4    "jupyterhub-singleus…"    jupyter-antonacci
6d423424d886   grafana/grafana:latest                "/run.sh -config /op…"    monitoring_grafana
de62257e3167   prom/prometheus:latest                "/bin/prometheus --c…"    monitoring_prometheus
27e34015e9ad   google/cadvisor:latest                "/usr/bin/cadvisor -…"    monitoring_cadvisor
bbf6cf47f460   prom/node-exporter:latest             "/bin/node_exporter"      monitoring_node_exporter
bed13c8c937c   jupyterhub_jupyterhub                 "/usr/bin/python3 /u…"    jupyterhub_jupyterhub_1
e03657db364b   jupyterhub/configurable-http-proxy    "/srv/configurable-h…"    jupyterhub_http_proxy_1
```

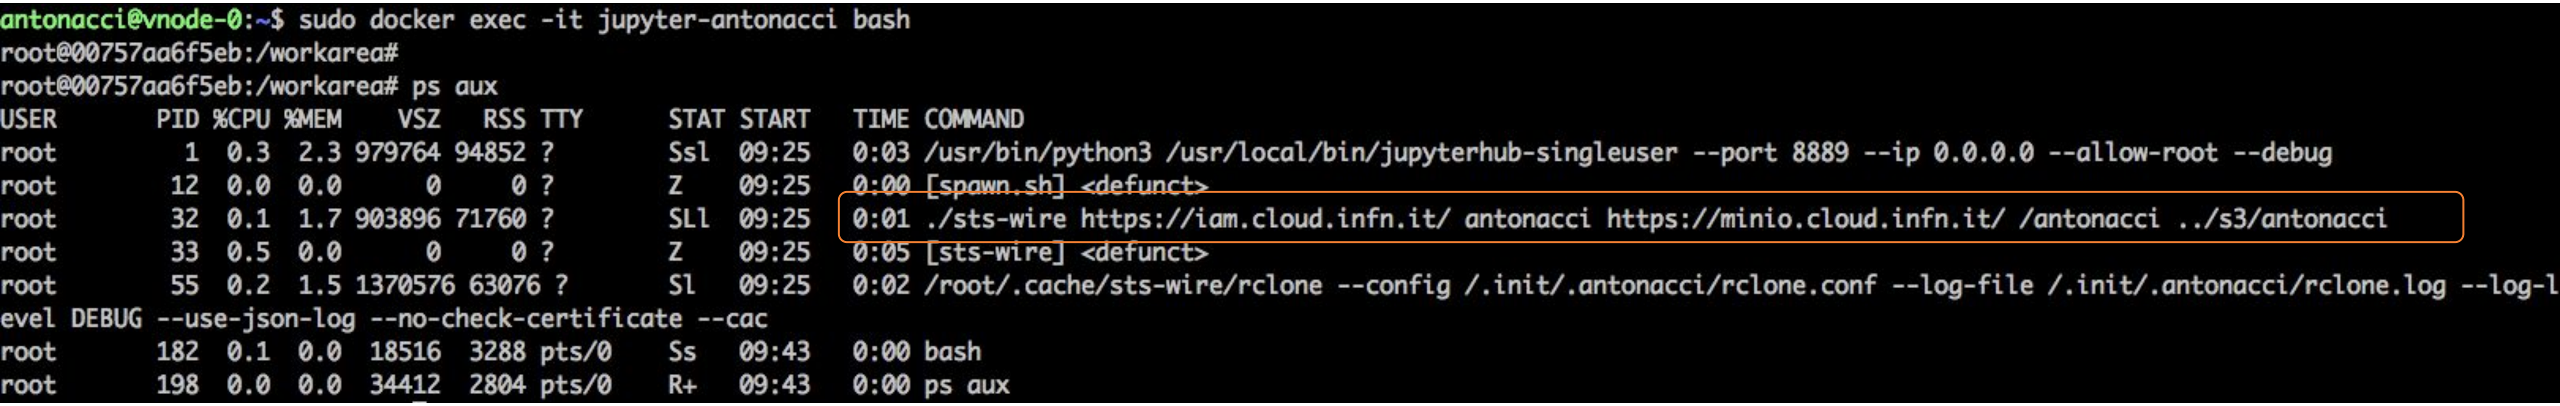

```
antonacci@vnode-0:~$ sudo docker exec -it jupyter-antonacci bash
root@00757aa6f5eb:/workarea#
root@00757aa6f5eb:/workarea# ps aux
USER       PID %CPU %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
root         1  0.3  2.3 979764 94852 ?        Ssl  09:25   0:03 /usr/bin/python3 /usr/local/bin/jupyterhub-singleuser --port 8889 --ip 0.0.0.0 --allow-root --debug
root        12  0.0  0.0      0     0 ?        Z    09:25   0:00 [spawn.sh] <defunct>
root        32  0.1  1.7 903896 71760 ?        SLl  09:25   0:01 ./sts-wire https://iam.cloud.infn.it/ antonacci https://minio.cloud.infn.it/ /antonacci ../s3/antonacci
root        33  0.5  0.0      0     0 ?        Z    09:25   0:05 [sts-wire] <defunct>
root        55  0.2  1.5 1370576 63076 ?       Sl   09:25   0:02 /root/.cache/sts-wire/rclone --config /.init/.antonacci/rclone.conf --log-file /.init/.antonacci/rclone.log --log-l
evel DEBUG --use-json-log --no-check-certificate --cac
root       182  0.1  0.0  18516  3288 pts/0    Ss   09:43   0:00 bash
root       198  0.0  0.0  34412  2804 pts/0    R+   09:43   0:00 ps aux
```

The admin panel

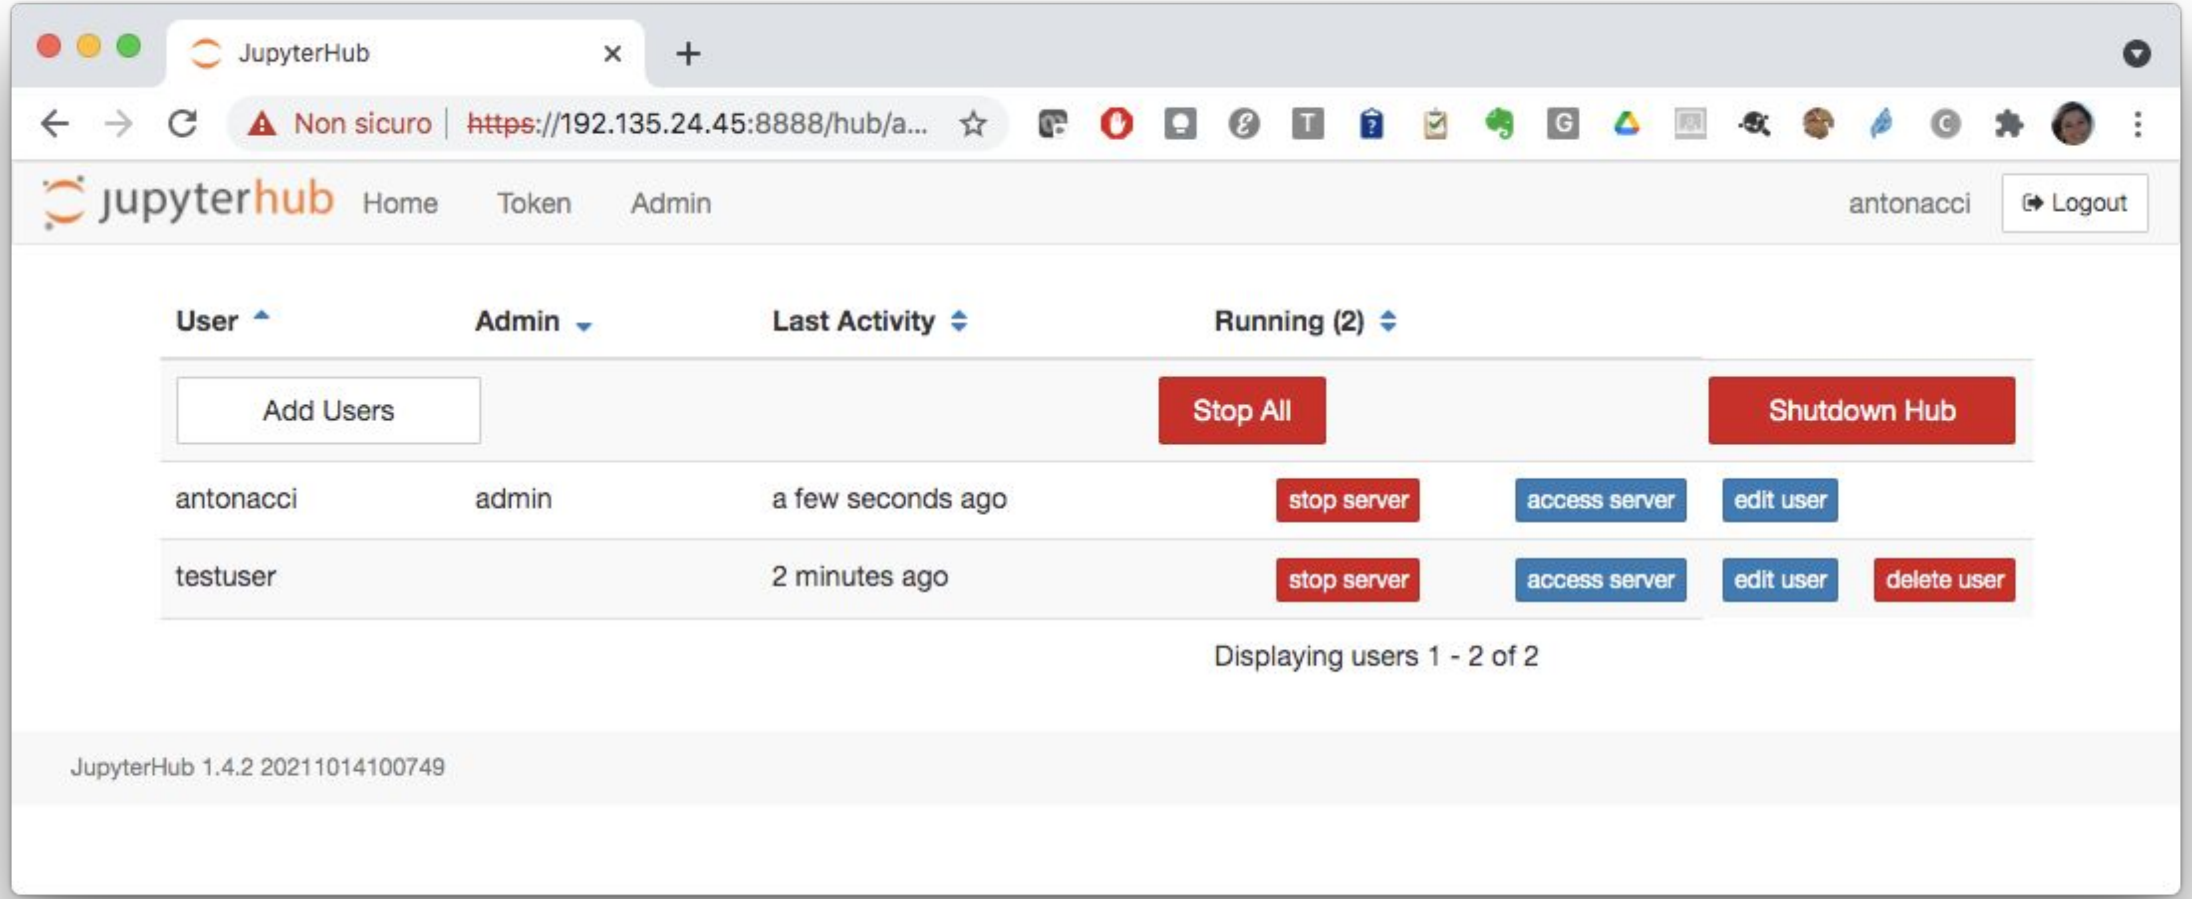

User	Admin	Last Activity	Running (2)
antonacci	admin	a few seconds ago	
testuser		2 minutes ago	

Displaying users 1 - 2 of 2

JupyterHub 1.4.2 20211014100749

The user view

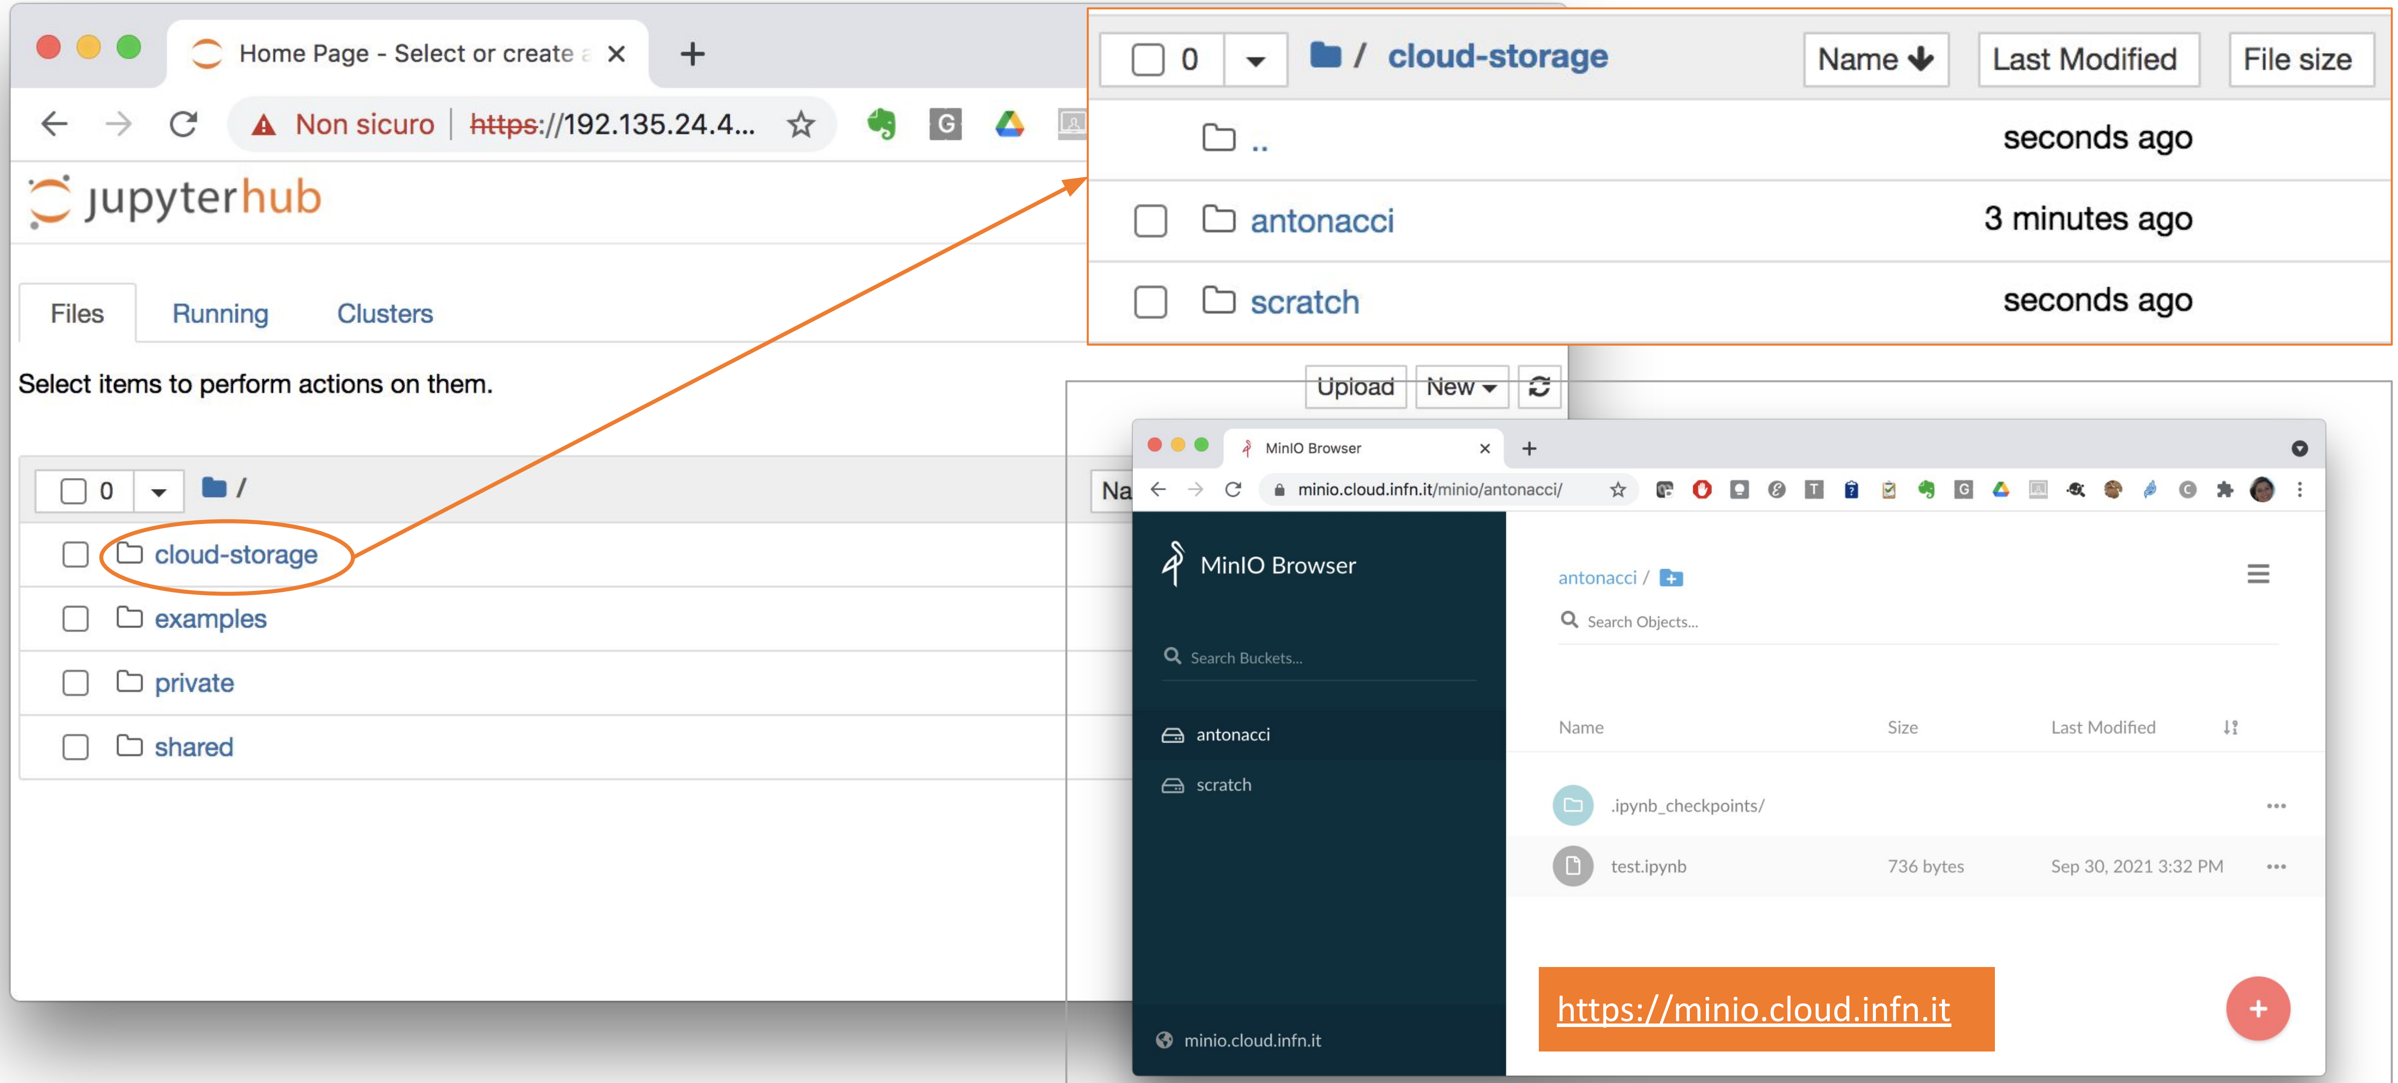

Home Page - Select or create a ✕

⚠ Non sicuro | ~~https~~://192.135.24.4...

jupyterhub

Files | Running | Clusters

Select items to perform actions on them.

Upload | New ▾ | ⟳

□ 0 ▾ /

	Name
□	cloud-storage
□	examples
□	private
□	shared

□ 0 ▾ / cloud-storage

	Name ↓	Last Modified	File size
	..	seconds ago	
□	antonacci	3 minutes ago	
□	scratch	seconds ago	

MinIO Browser ✕

minio.cloud.infn.it/minio/antonacci/

MinIO Browser

Search Buckets...

- antonacci
- scratch

minio.cloud.infn.it

antonacci /

Search Objects...

Name	Size	Last Modified
.ipynb_checkpoints/		
test.ipynb	736 bytes	Sep 30, 2021 3:32 PM

https://minio.cloud.infn.it

Monitoring

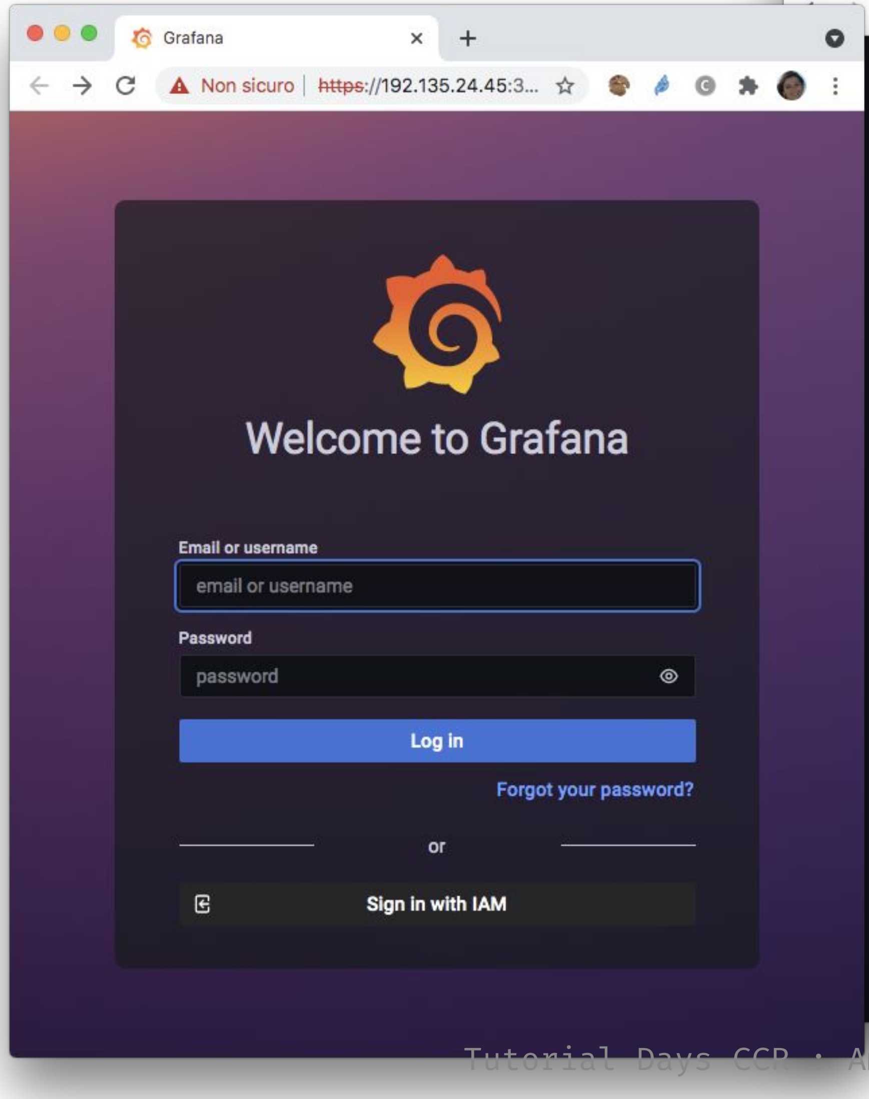

Welcome to Grafana

Email or username

email or username

Password

password

Log in

Forgot your password?

or

Sign in with IAM

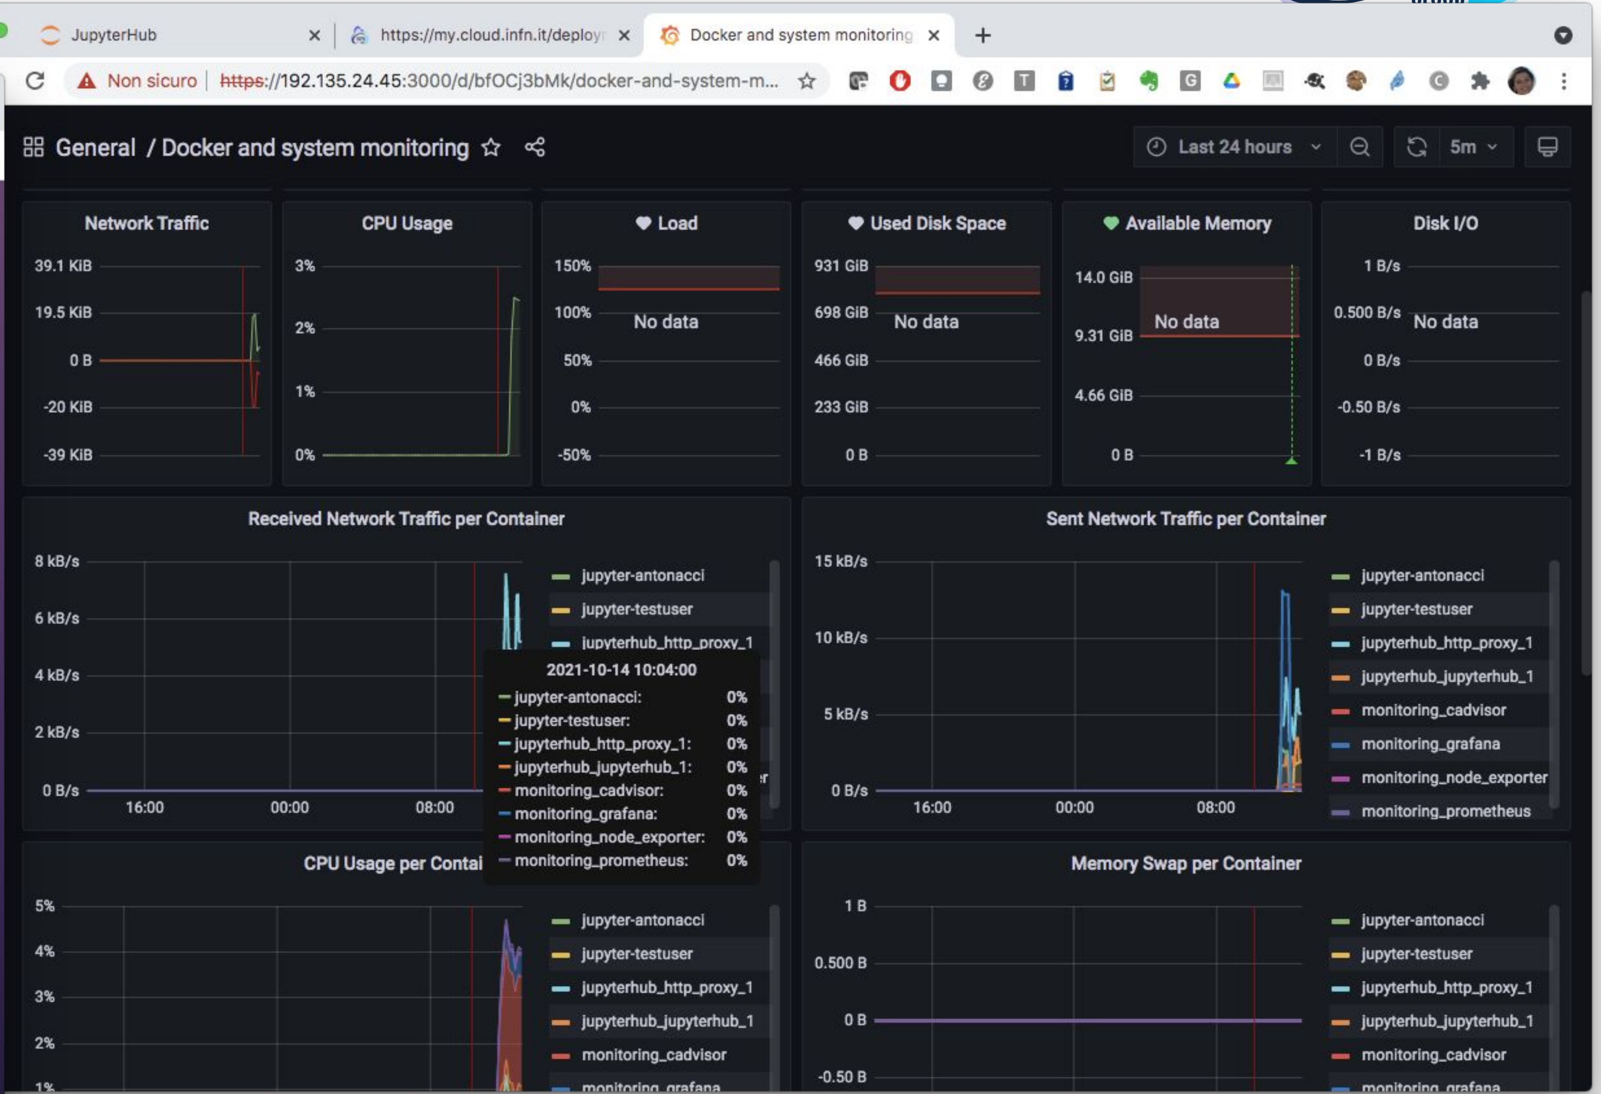

```
2021-10-14 10:04:00
jupyter-antonacci:            0%
jupyter-testuser:             0%
jupyterhub_http_proxy_1:      0%
jupyterhub_jupyterhub_1:      0%
monitoring_cadvisor:          0%
monitoring_grafana:           0%
monitoring_node_exporter:     0%
monitoring_prometheus:        0%
```

Real time collaboration - New!

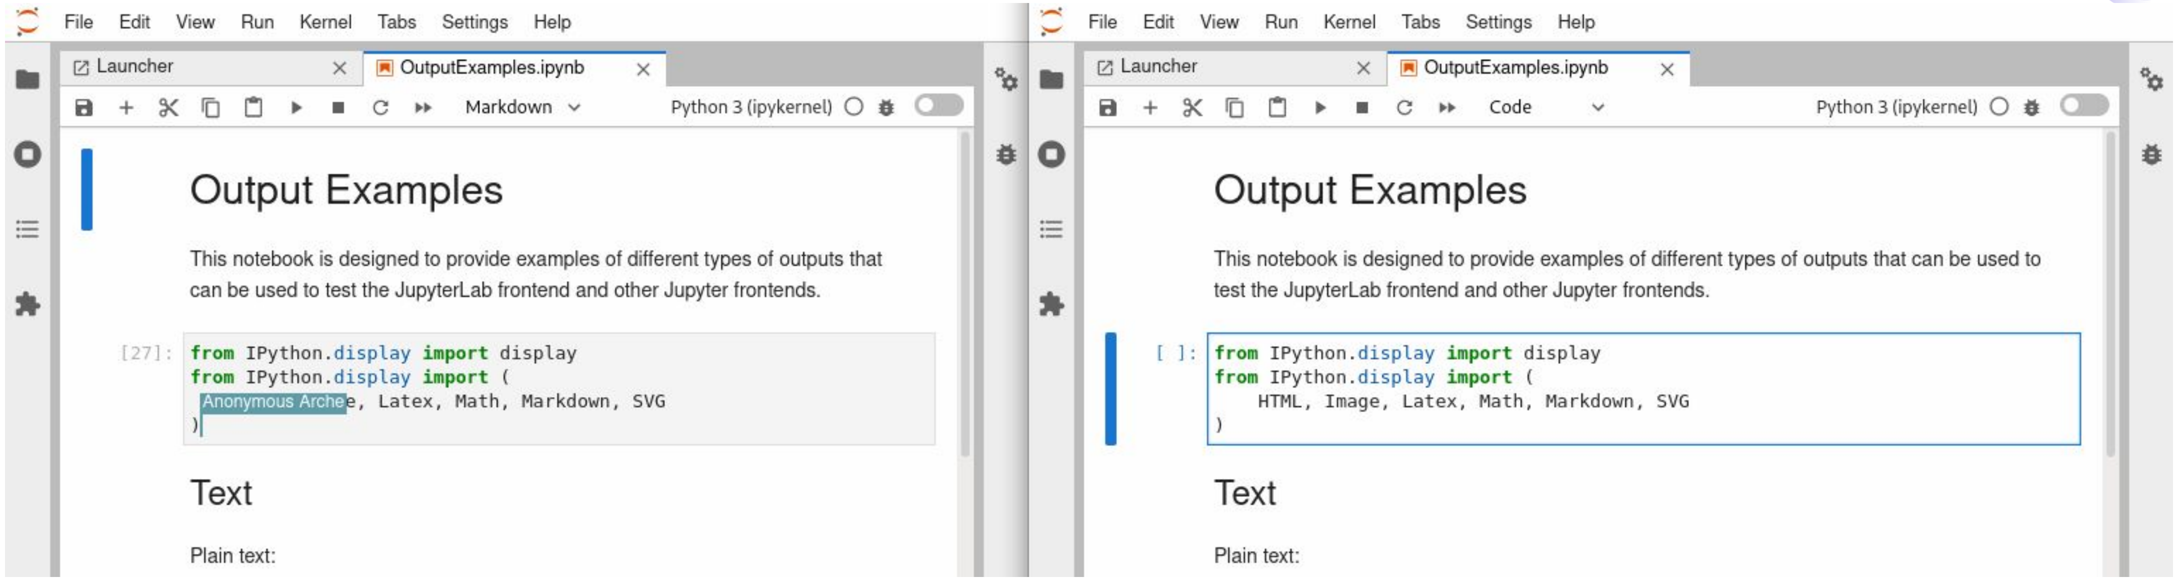

Output Examples

This notebook is designed to provide examples of different types of outputs that can be used to test the JupyterLab frontend and other Jupyter frontends.

[27]:

```
from IPython.display import display
from IPython.display import (
 Anonymous Arche, Latex, Math, Markdown, SVG
)
```

Text

Plain text:

Output Examples

This notebook is designed to provide examples of different types of outputs that can be used to test the JupyterLab frontend and other Jupyter frontends.

[ ]:

```
from IPython.display import display
from IPython.display import (
    HTML, Image, Latex, Math, Markdown, SVG
)
```

Text

Plain text:

This feature will be available soon..we are finalizing the user documentation highlighting the caveats and limitations: p.e. not all editors support sync (because, under the hood, they use a different model to represent the document’s content)

Computational enviroment for Machine Learning INFN (ML_INFN)

Working Station for CYGNO experiment

Community customizations

Built on top of the Jupyter VM use-case

ML-INFN environment

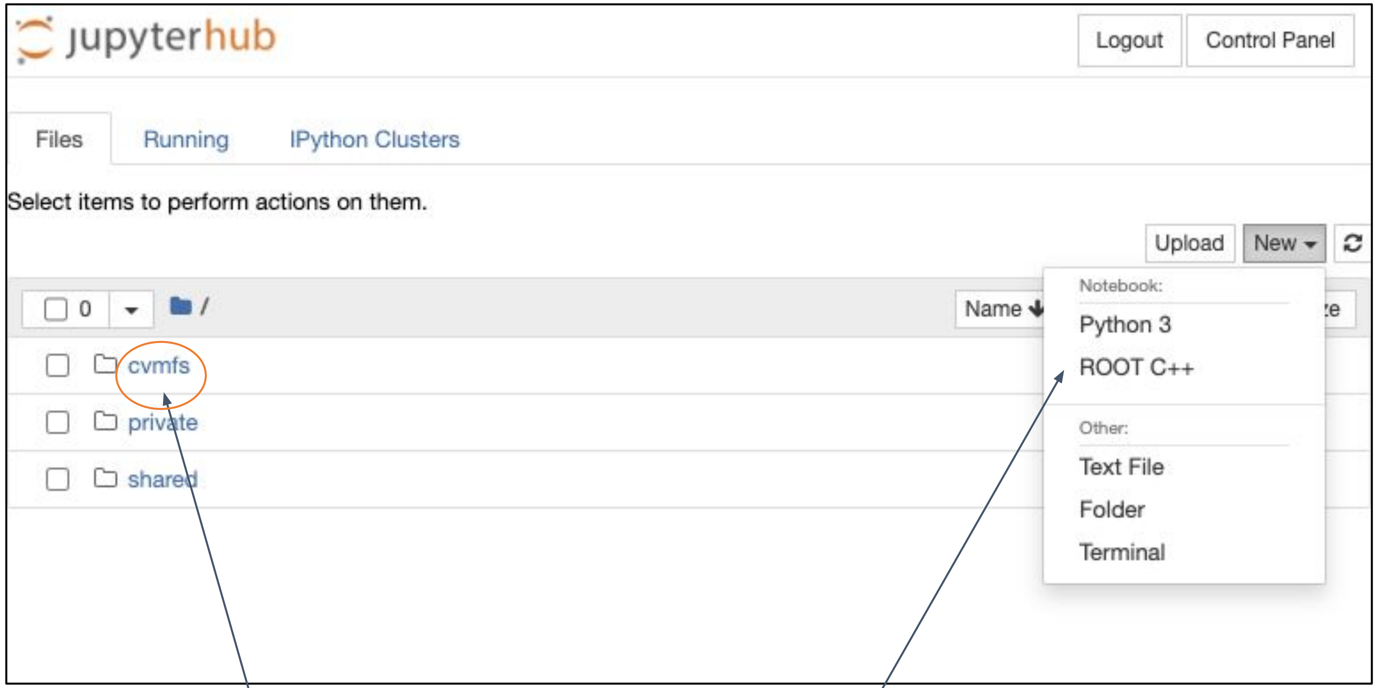

Select items to perform actions on them.

- cvmfs
- private
- shared

CVMFS mount

Customized image with specific libraries

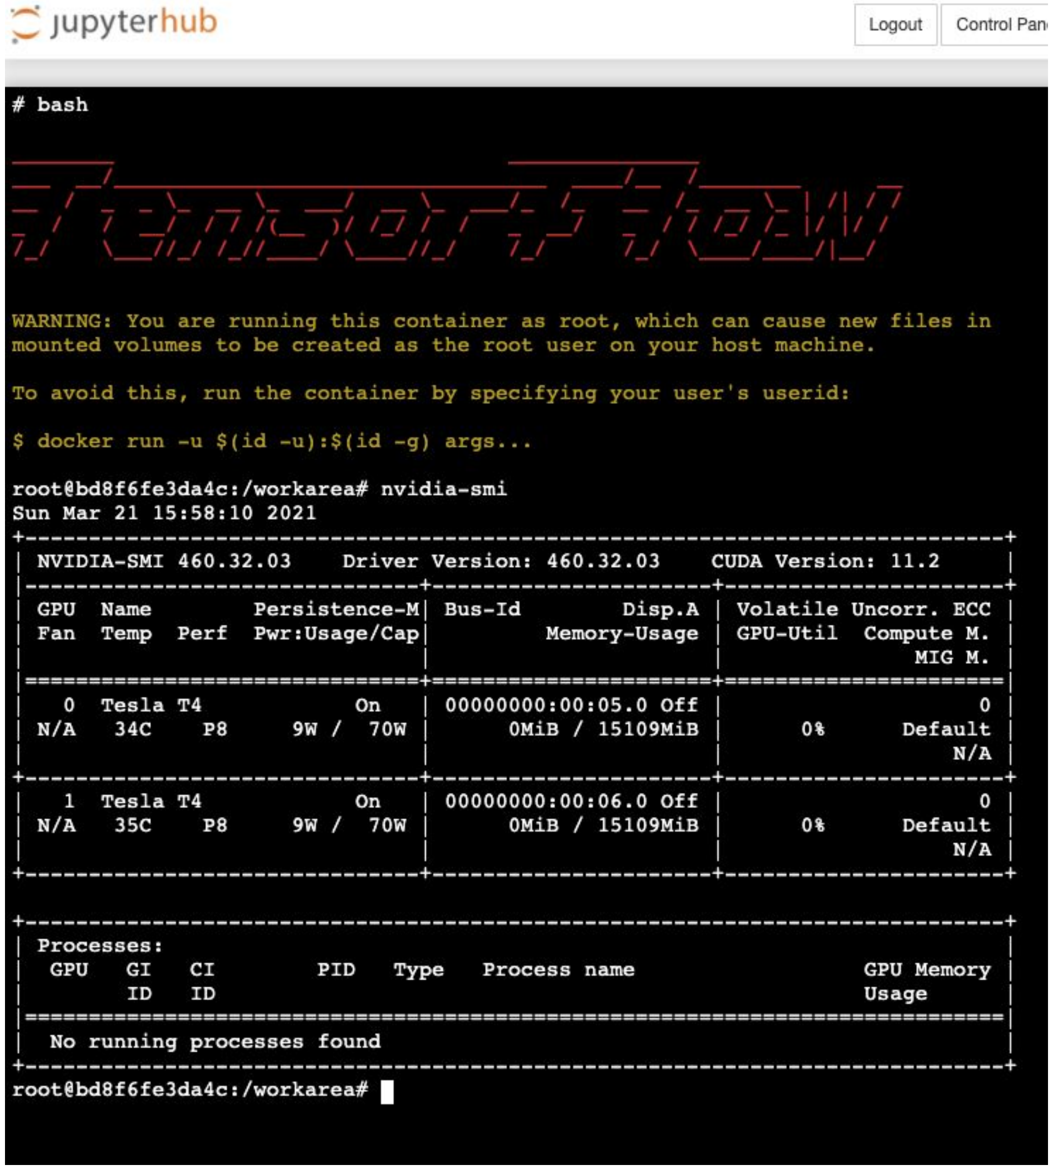

```
# bash

WARNING: You are running this container as root, which can cause new files in
mounted volumes to be created as the root user on your host machine.

To avoid this, run the container by specifying your user's userid:

$ docker run -u $(id -u):$(id -g) args...

root@bd8f6fe3da4c:/workarea# nvidia-smi
Sun Mar 21 15:58:10 2021
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 460.32.03    Driver Version: 460.32.03    CUDA Version: 11.2     |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|                               |                      |               MIG M. |
|===============================+======================+======================|
|   0  Tesla T4            On   | 00000000:00:05.0 Off |                    0 |
| N/A   34C    P8     9W /  70W |      0MiB / 15109MiB |      0%      Default |
|                               |                      |                  N/A |
+-------------------------------+----------------------+----------------------+
|   1  Tesla T4            On   | 00000000:00:06.0 Off |                    0 |
| N/A   35C    P8     9W /  70W |      0MiB / 15109MiB |      0%      Default |
|                               |                      |                  N/A |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Processes:                                                                  |
|  GPU   GI   CI        PID   Type   Process name                  GPU Memory |
|        ID   ID                                                   Usage      |
|=============================================================================|
|  No running processes found                                                 |
+-----------------------------------------------------------------------------+
root@bd8f6fe3da4c:/workarea# 
```

GPU(s) availability

Spark + Jupyter cluster

Spark + Jupyter cluster

How to deploy a complete Spark + Jupyter Notebook on top of a Kubernetes (K8s) computing cluster

Deployment architecture

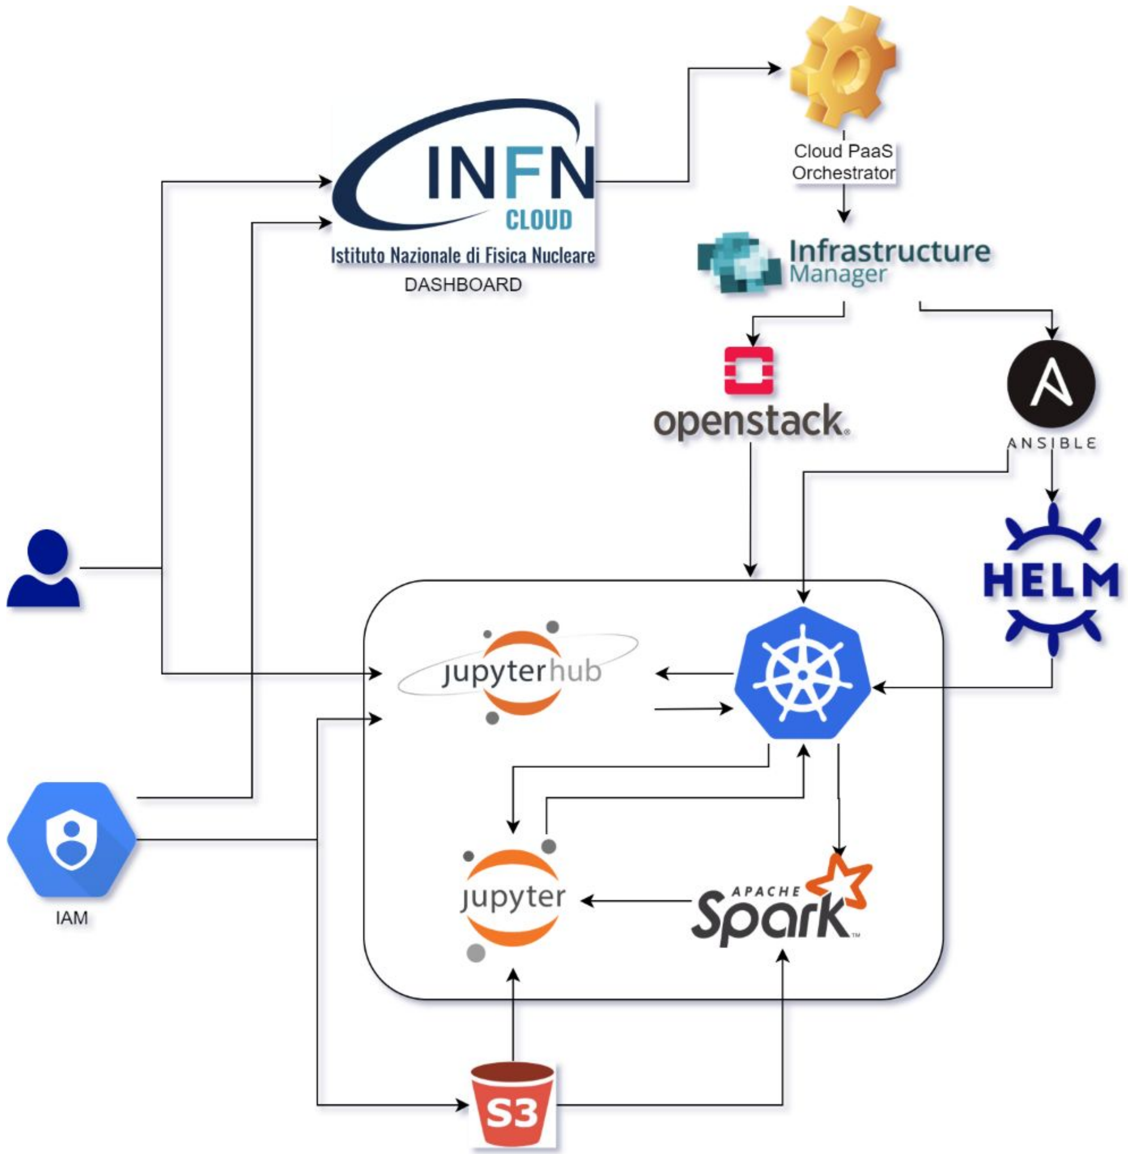

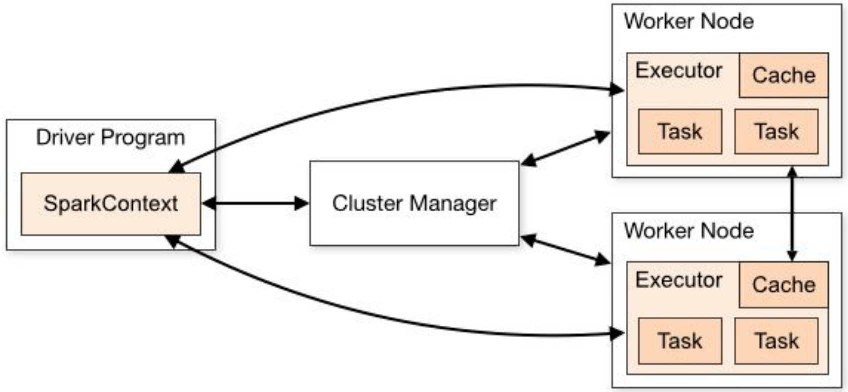

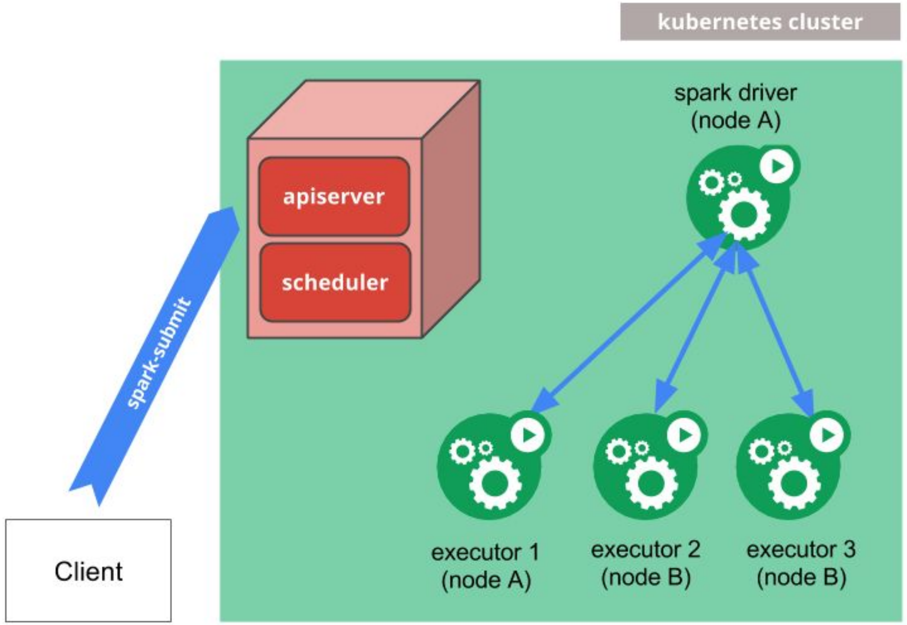

The deployment outputs

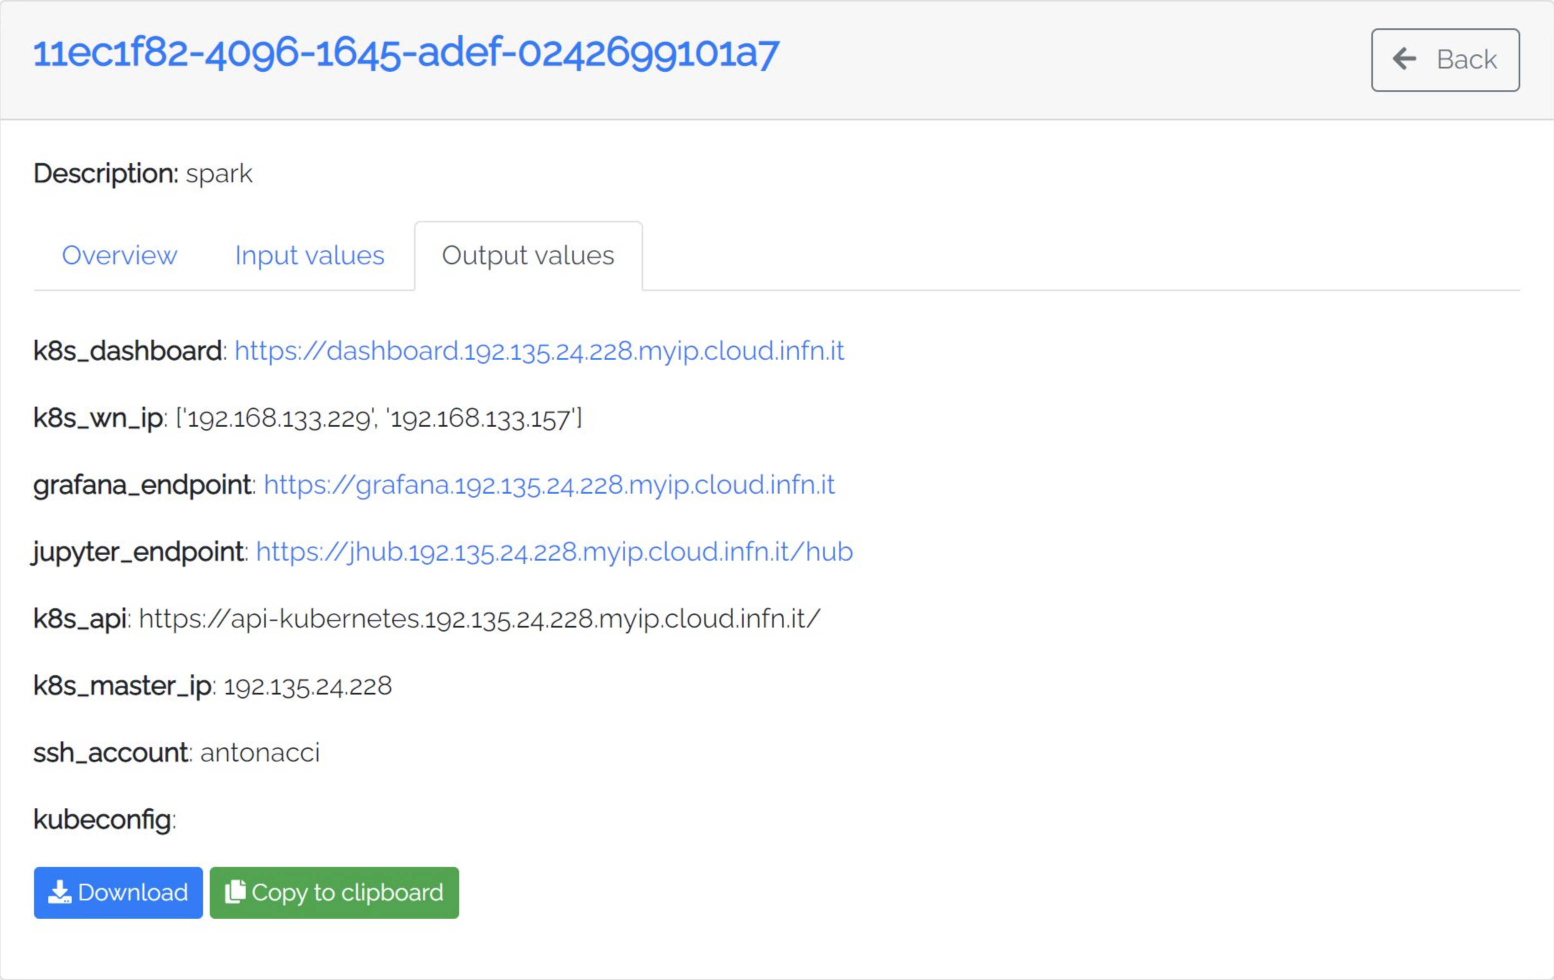

11ec1f82-4096-1645-adef-0242699101a7

Description: spark

k8s_dashboard: https://dashboard.192.135.24.228.myip.cloud.infn.it

k8s_wn_ip: ['192.168.133.229', '192.168.133.157']

grafana_endpoint: https://grafana.192.135.24.228.myip.cloud.infn.it

jupyter_endpoint: https://jhub.192.135.24.228.myip.cloud.infn.it/hub

k8s_api: https://api-kubernetes.192.135.24.228.myip.cloud.infn.it/

k8s_master_ip: 192.135.24.228

ssh_account: antonacci

kubeconfig:

And more…

Further use-cases/services

- Elasticsearch + Kibana - already available
- HTCondor cluster - will be available soon
- *Notebook as a Service (NaaS) - beta testing will start soon*
- *Deployment on private network - will be available in the next future*

INFN-Cloud allows you to implement new services and/or customize existing ones

❑ please contact us at cloud-support@infn.it and you will be redirected to the proper INFN Cloud support team.

References

Web site:
https://www.cloud.infn.it

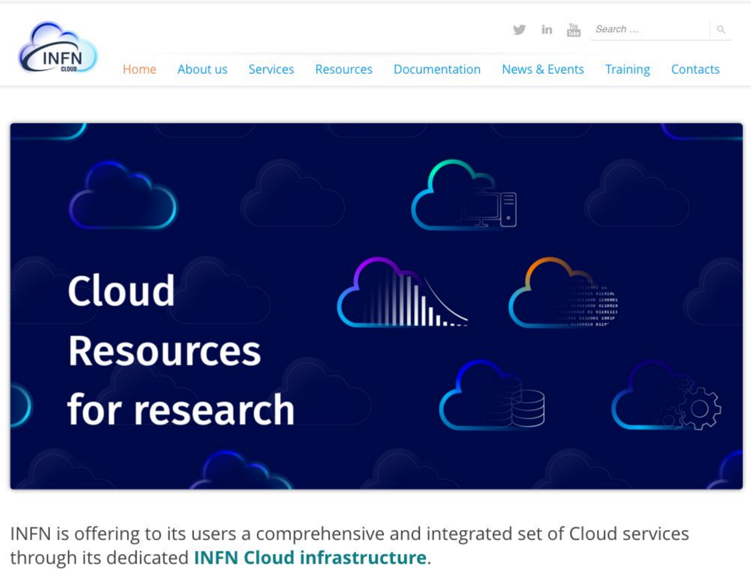

Cloud Resources for research

INFN is offering to its users a comprehensive and integrated set of Cloud services through its dedicated **INFN Cloud infrastructure**.

Documentation :
https://guides.cloud.infn.it/docs/users-guides/en/latest/

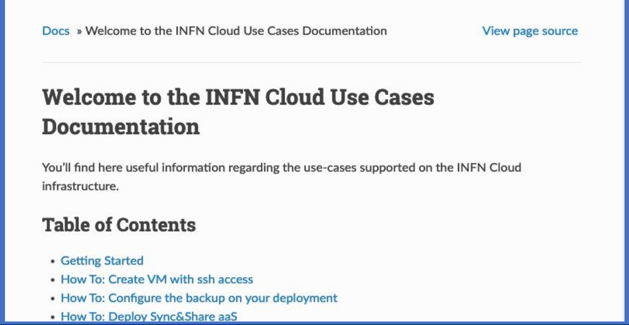

Welcome to the INFN Cloud Use Cases Documentation

You'll find here useful information regarding the use-cases supported on the INFN Cloud infrastructure.

Table of Contents

- Getting Started
- How To: Create VM with ssh access
- How To: Configure the backup on your deployment
- How To: Deploy Sync&Share aaS

Support :
https://servicedesk.cloud.infn.it or ✉ cloud-support@infn.it

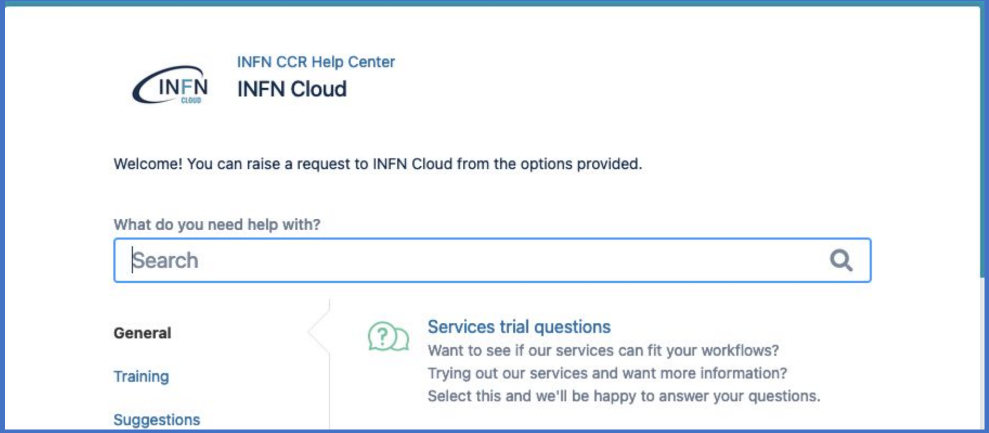

INFN CCR Help Center
INFN Cloud

Welcome! You can raise a request to INFN Cloud from the options provided.

What do you need help with?

Services trial questions
Want to see if our services can fit your workflows?
Trying out our services and want more information?
Select this and we'll be happy to answer your questions.

Thank you

for your attention!