INFN Cloud centralized services

Corso di formazione per neoassunti nelle attività di computing – INFN DATACloud
06 Marzo 2025

Federica Fanzago – INFN Padova

Outline

- Short description
  - INFN Cloud object storage (Ceph-RGW)
  - INFN Cloud Registry (Harbor)
  - Notebook as a Service (NaaS)
  - Healthchecks as a Service

Available services

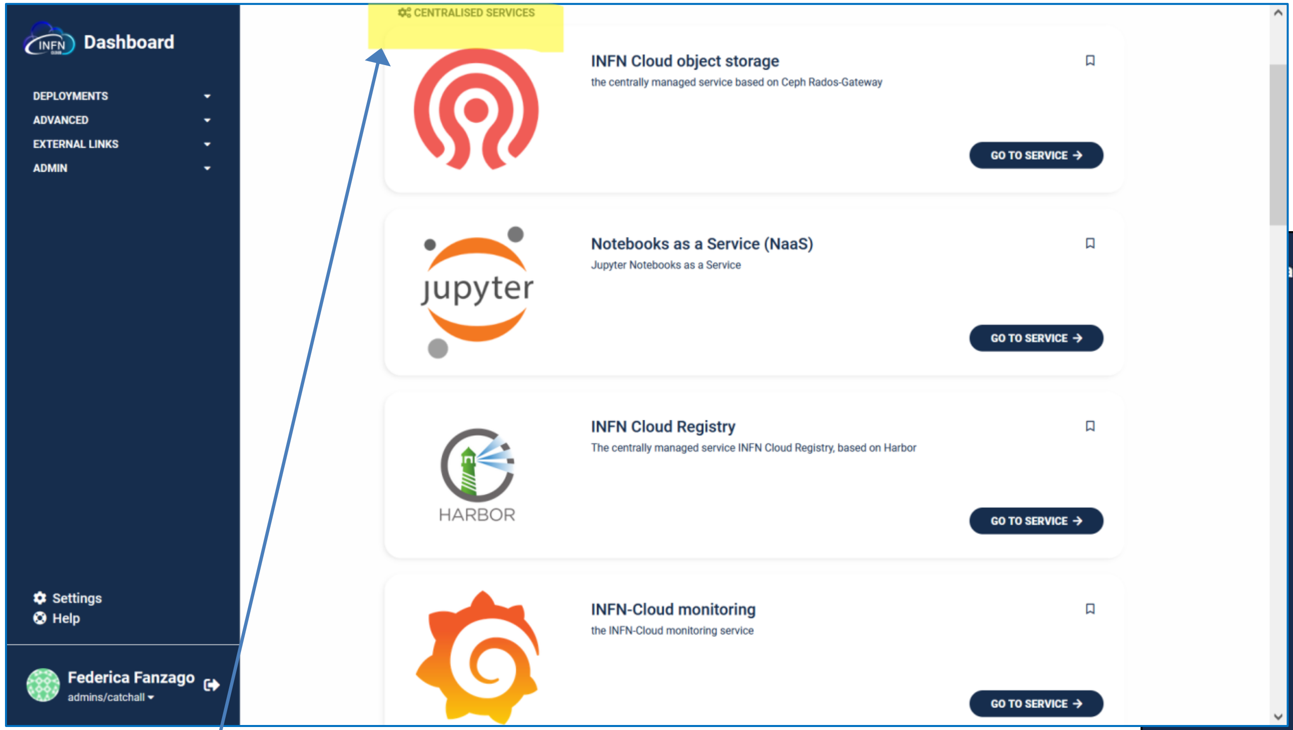

Dashboard

- DEPLOYMENTS
- ADVANCED
- EXTERNAL LINKS
- ADMIN

CENTRALISED SERVICES

INFN Cloud object storage
the centrally managed service based on Ceph Rados-Gateway
GO TO SERVICE →

Notebooks as a Service (NaaS)
Jupyter Notebooks as a Service
GO TO SERVICE →

INFN Cloud Registry
The centrally managed service INFN Cloud Registry, based on Harbor
GO TO SERVICE →

INFN-Cloud monitoring
the INFN-Cloud monitoring service
GO TO SERVICE →

Settings
Help

Federica Fanzago
admins/catchall

Same view for all the registered INFN Cloud users

View depends on user's iam group

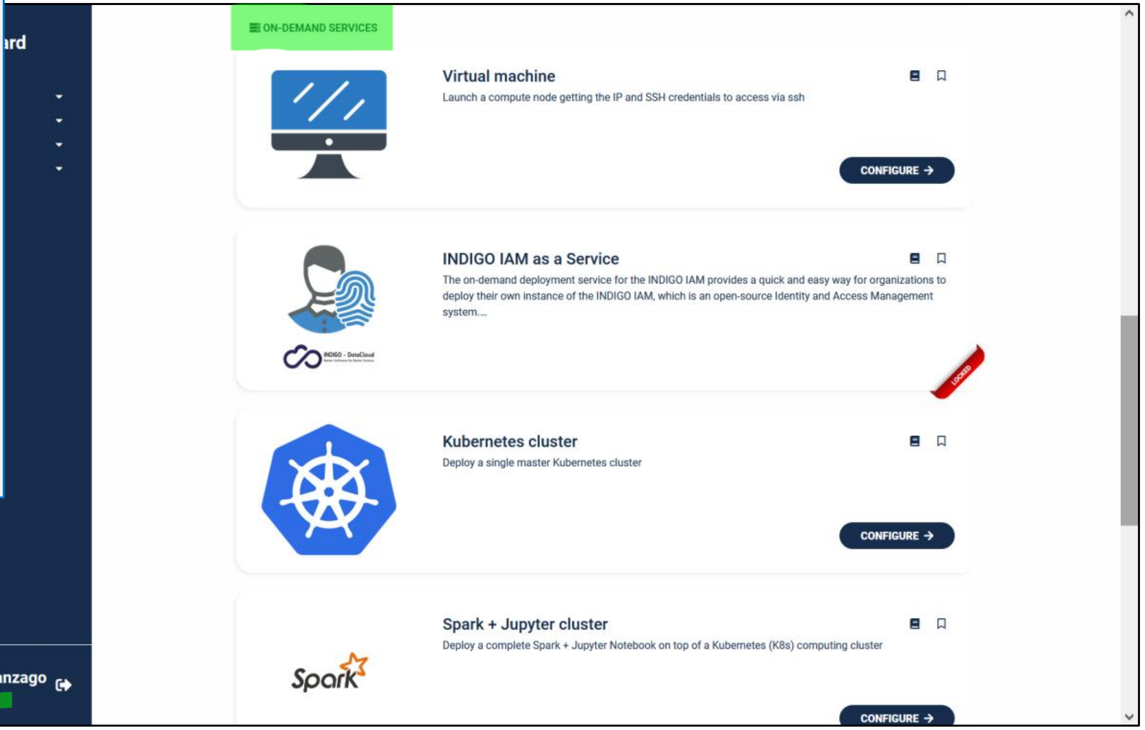

ON-DEMAND SERVICES

Virtual machine
Launch a compute node getting the IP and SSH credentials to access via ssh
CONFIGURE →

INDIGO IAM as a Service
The on-demand deployment service for the INDIGO IAM provides a quick and easy way for organizations to deploy their own instance of the INDIGO IAM, which is an open-source Identity and Access Management system....
LOCKED

Kubernetes cluster
Deploy a single master Kubernetes cluster
CONFIGURE →

Spark + Jupyter cluster
Deploy a complete Spark + Jupyter Notebook on top of a Kubernetes (K8s) computing cluster
CONFIGURE →

Settings
Help

Federica Fanzago
admins/ai-infn

Centrally managed services

- Fully-managed
  - their installation, configuration, update and security are under «INFN Cloud» responsibility (dedicated team).
- Installed on the INFN Cloud backbone, CNAF and Bari.
- Available and ready to use for all the registered users in INFN Cloud
  - Iam groups added by default: users/catchall, users/s3 and users/naas
  - the system admin nomination «nomina ad amministratore INFN Cloud (quale utente amministratore)» isn't required

INFN Cloud object storage

INFN Cloud object storage

the centrally managed service based on Ceph Rados-Gateway

- https://s3webui.cloud.infn.it
- https://guides.cloud.infn.it/docs/users-guides/en/latest/users_guides/centralised/objectstorage.html
- INFN Cloud provides object storage as centrally managed service. Ceph Storage Cluster is the backend
  - Migrated from Openstack Swift with Minio-gw in June 2024
- This storage is the persistent one for INFN Cloud.
- The object storage is replicated in the two sites of backbone, CNAF and Bari (no distributed cluster)
  - It guarantees the redundancy of data and disaster recovery (replica via sync)
  - No data backup

RADOS Gateway (RGW)

- INFN Cloud is using the Ceph Rados Gateway (RESTful gateway for object storage)
  - It provides access via S3 compatible api
- There are three instances of RGW in each site (High availability)
- Only authorized users can access the storage service (authentication and authorization via IAM) .
- A web interface to access the storage, create buckets and manage files has been developed using S3 protocol and OAuth2/OpenID Connect

Login to webui via OpenID

https://s3webui.cloud.infn.it

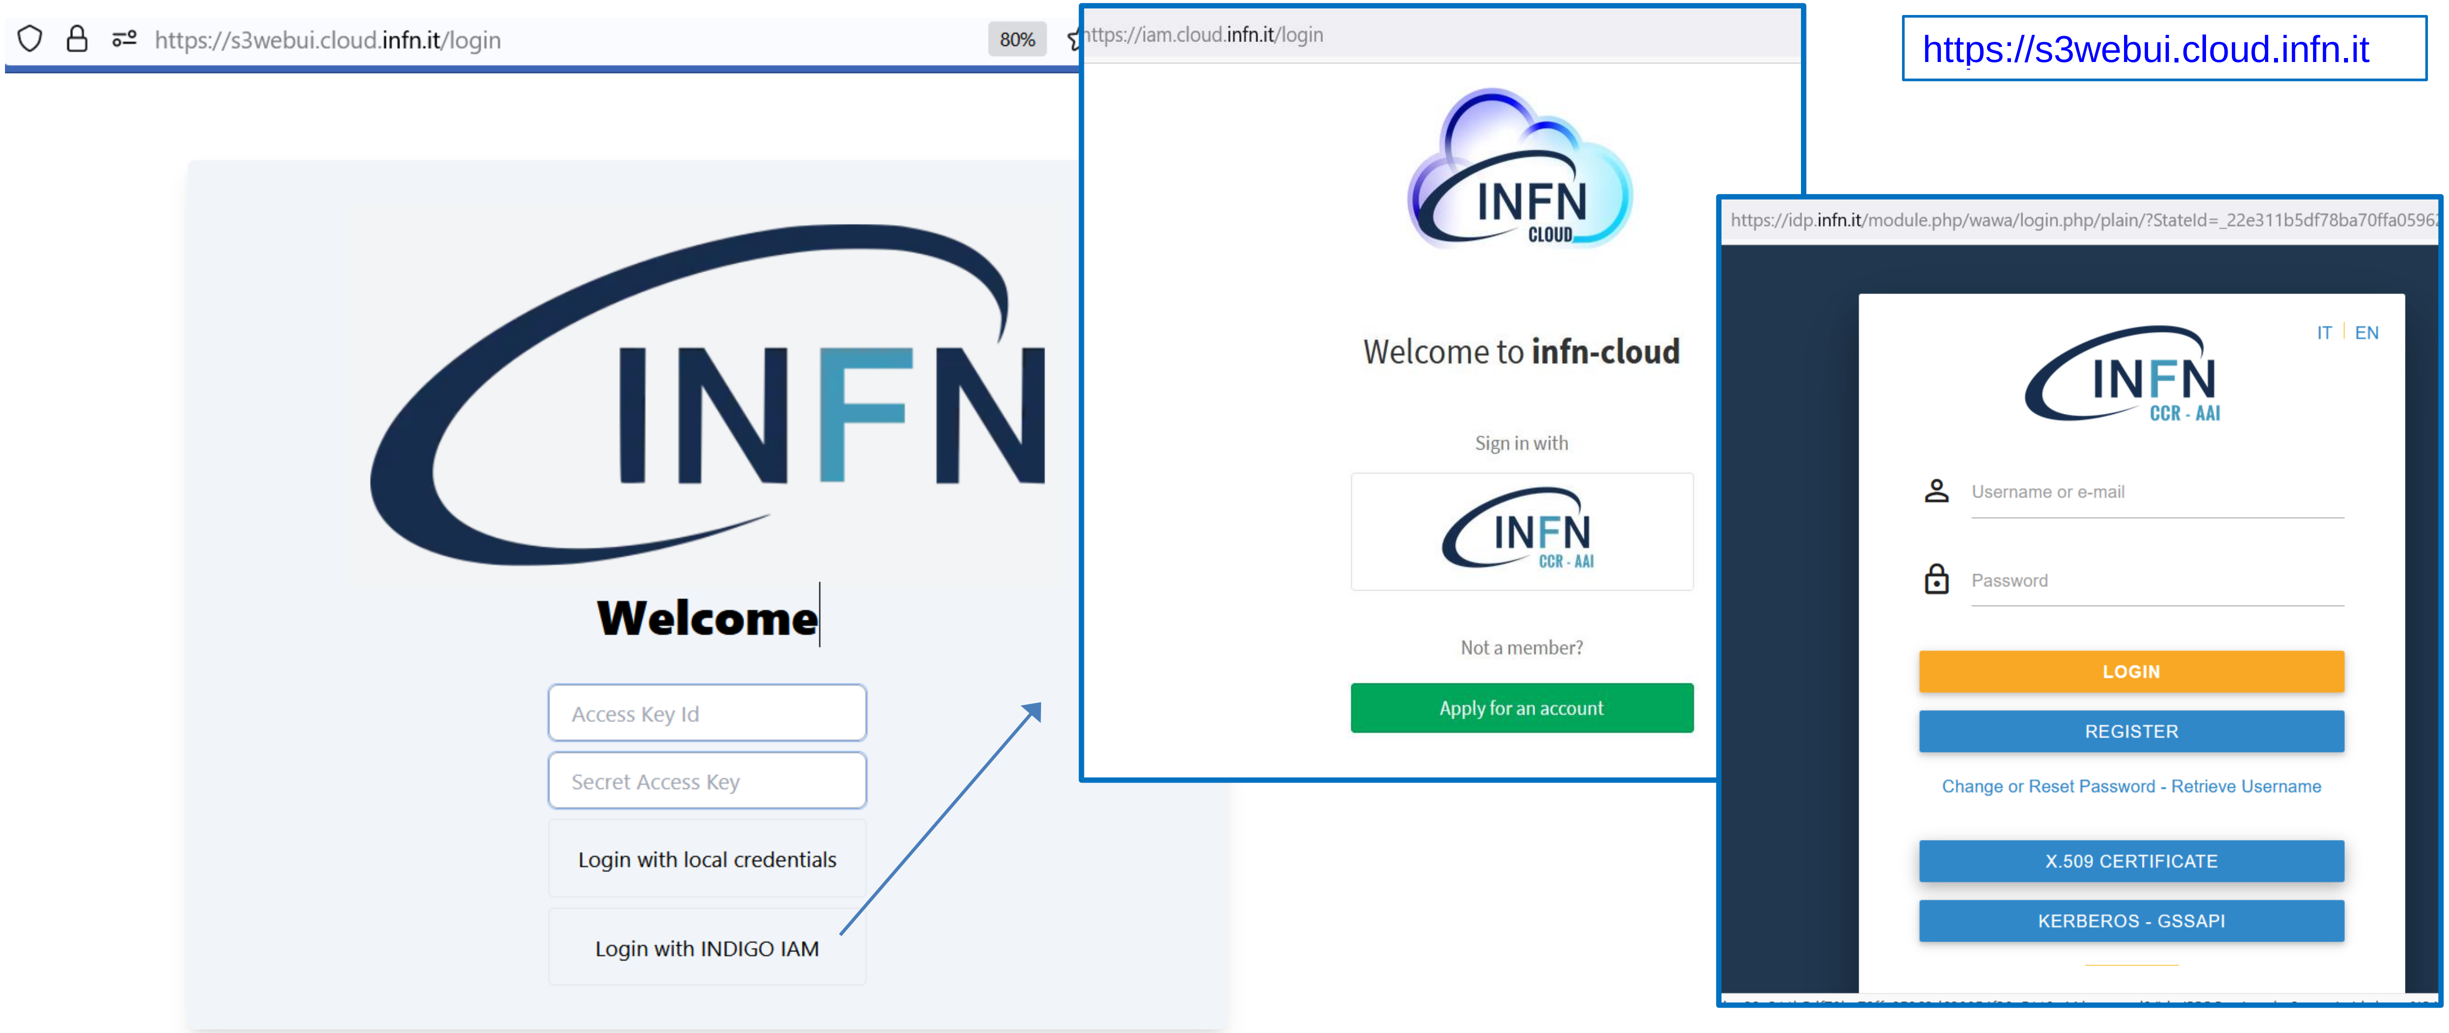

https://s3webui.cloud.infn.it/login

Welcome

Access Key Id

Secret Access Key

Login with local credentials

Login with INDIGO IAM

https://iam.cloud.infn.it/login

Welcome to **infn-cloud**

Sign in with

INFN CCR - AAI

Not a member?

Apply for an account

https://idp.infn.it/module.php/wawa/login.php/plain/?StateId=_22e311b5df78ba70ffa0596

IT | EN

Username or e-mail

Password

LOGIN

REGISTER

Change or Reset Password - Retrieve Username

X.509 CERTIFICATE

KERBEROS - GSSAPI

Persistent data with object storage

- The storage configuration allows users to have a personal bucket (directory) and an area under /scratch labeled as their INFN AAI username.
- The quota for each user is 200GB.
- Data stored in the personal bucket are private, data under /scratch are visible and downloadable by all.
- The storage is integrated with the Open Policy Agent (OPA) that enforces bucket policies rules.

Personal bucket

- The personal bucket can by create by user accessing the Ceph RGW webui otherwise it is automatically created the first time user instantiates a Jupyter notebook via INFN Cloud.

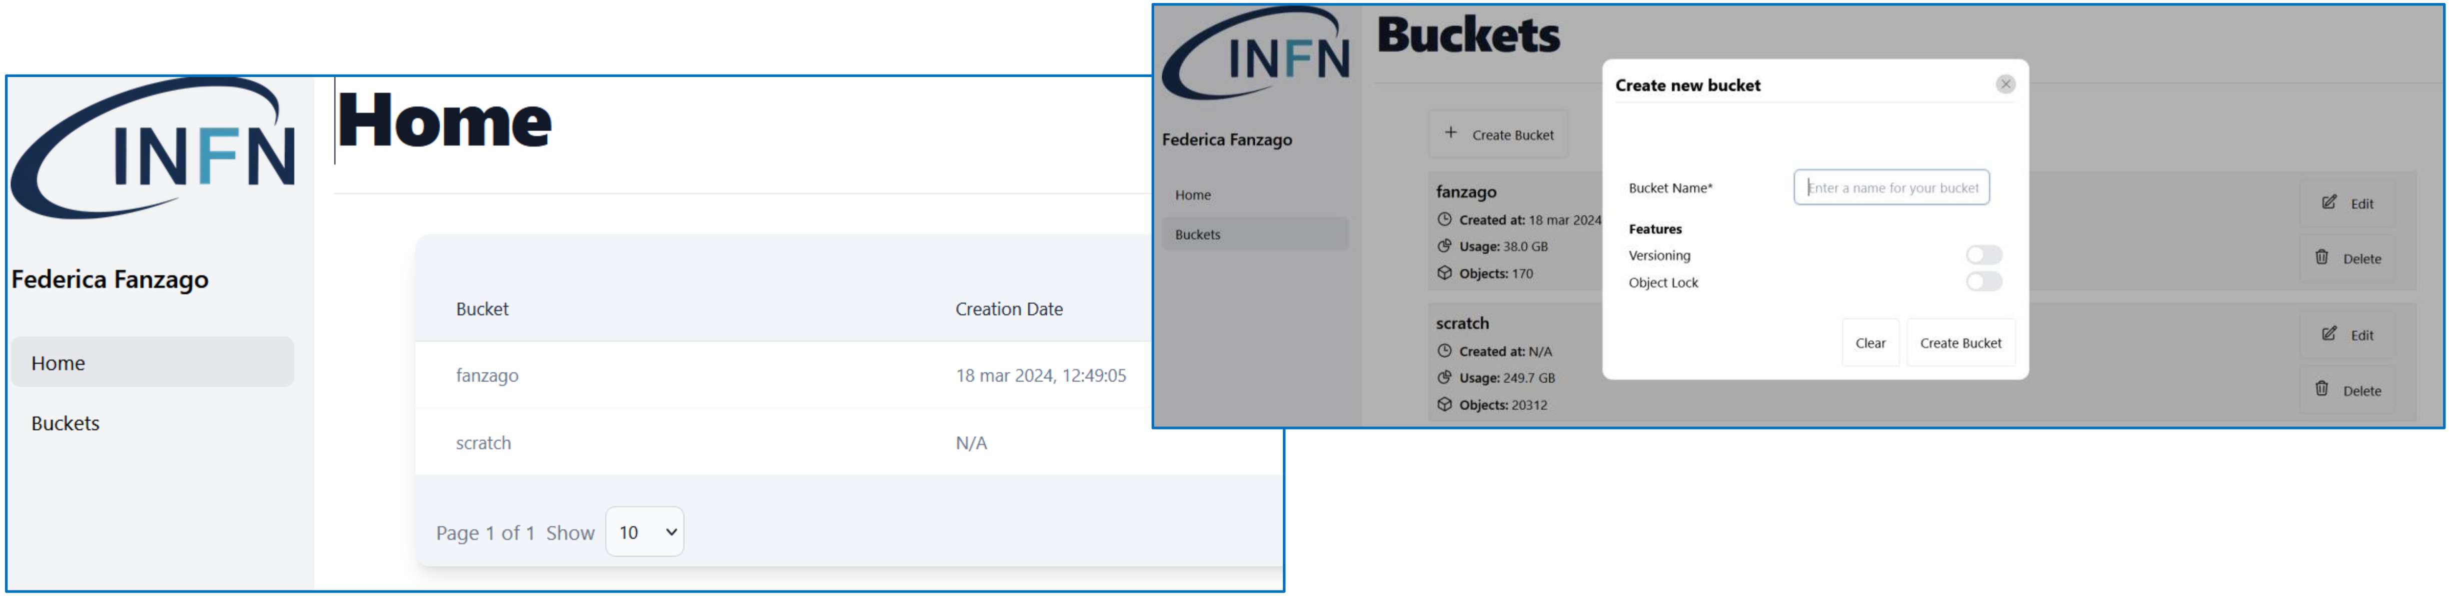

Home

Federica Fanzago

Bucket	Creation Date
fanzago	18 mar 2024, 12:49:05
scratch	N/A

Page 1 of 1 Show 10

Buckets

Federica Fanzago

fanzago
Created at: 18 mar 2024
Usage: 38.0 GB
Objects: 170

scratch
Created at: N/A
Usage: 249.7 GB
Objects: 20312

Create new bucket

Bucket Name* Enter a name for your bucket

Features

Versioning

Object Lock

Personal area under scratch

- The area under /scratch has to be created by the user

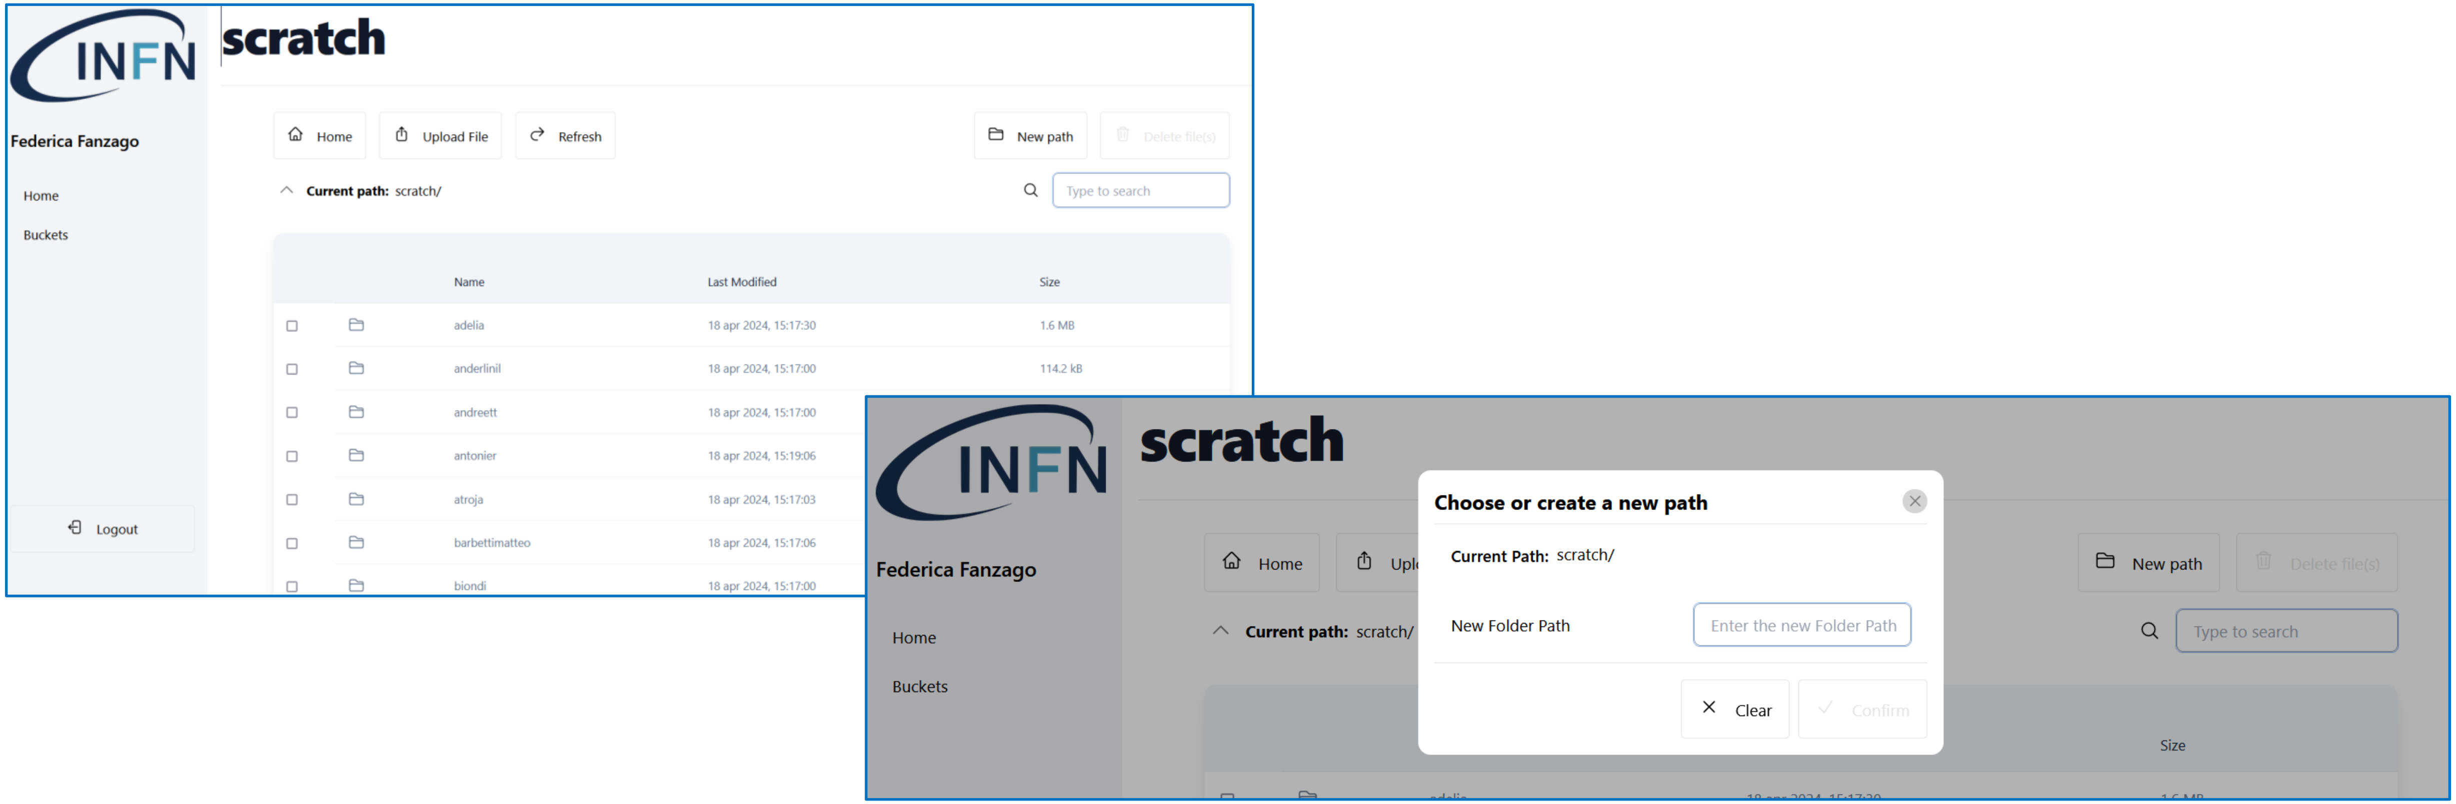

Current path: scratch/

		Name	Last Modified	Size
□		adelia	18 apr 2024, 15:17:30	1.6 MB
□		anderlinil	18 apr 2024, 15:17:00	114.2 kB
□		andreett	18 apr 2024, 15:17:00	
□		antonier	18 apr 2024, 15:19:06	
□		atroja	18 apr 2024, 15:17:03	
□		barbettimatteo	18 apr 2024, 15:17:06	
□		biondi	18 apr 2024, 15:17:00	

Choose or create a new path

Current Path: scratch/

New Folder Path: Enter the new Folder Path

S3 api and user configuration backup

- Prerequisites to get the IAM token with audience «object»:
  - Oidc-agent installed
  - Oidc profile configured
- S3 api to access data via code (boto3 api + Secure Token Service)
- A system for the backup offline for users to save files/configurations via Rclone is available. Image in Habor (the container registry)
- Via Rclone the object storage is accessible as local file system
  - rclone ls profile:/<bucket>

Object storage also for

- Cvmfs users repo i.e /cvmfs/username.infn.it
- Cvmfs software repo «datacloud.infn.it»
- Harbor repos

CernVM-File system (cvmfs) is a a read-only file system

INFN Cloud Registry

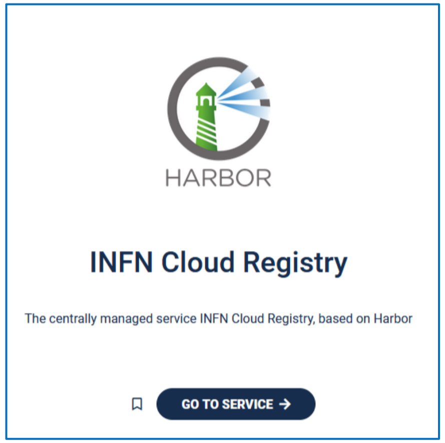

INFN Cloud Registry

The centrally managed service INFN Cloud Registry, based on Harbor

- https://harbor.cloud.infn.it
- https://guides.cloud.infn.it/docs/users-guides/en/latest/users_guides/centralised/harbor.html
- The INFN Cloud Registry for container images and helm charts, based on Harbor software.
    - «Project Harbor is an an open-source trusted cloud native registry project that stores, signs, and scans content»
- Harbor extends some functionalities of Docker Hub (i.e checks vulnerabilities) and overcomes its pull rate limit (i.e replica, proxy-cache).

Harbor: access

- Installed on the backbone (CNAF and Bari), high availability
- Users can access the service via oidc provider, using INFN-AAI credentials.
  - Users can see projects associated with their iam groups (if not visible, projects have to be created by Harbor admin).
- Two types of project supported:
  - Public: any user can pull images from this project
  - Private: only users who are members of the project can pull images

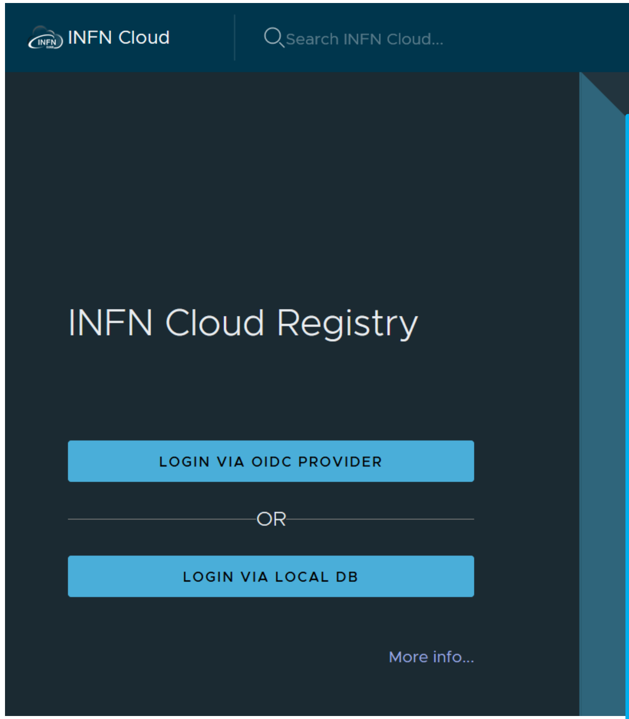

INFN Cloud Registry

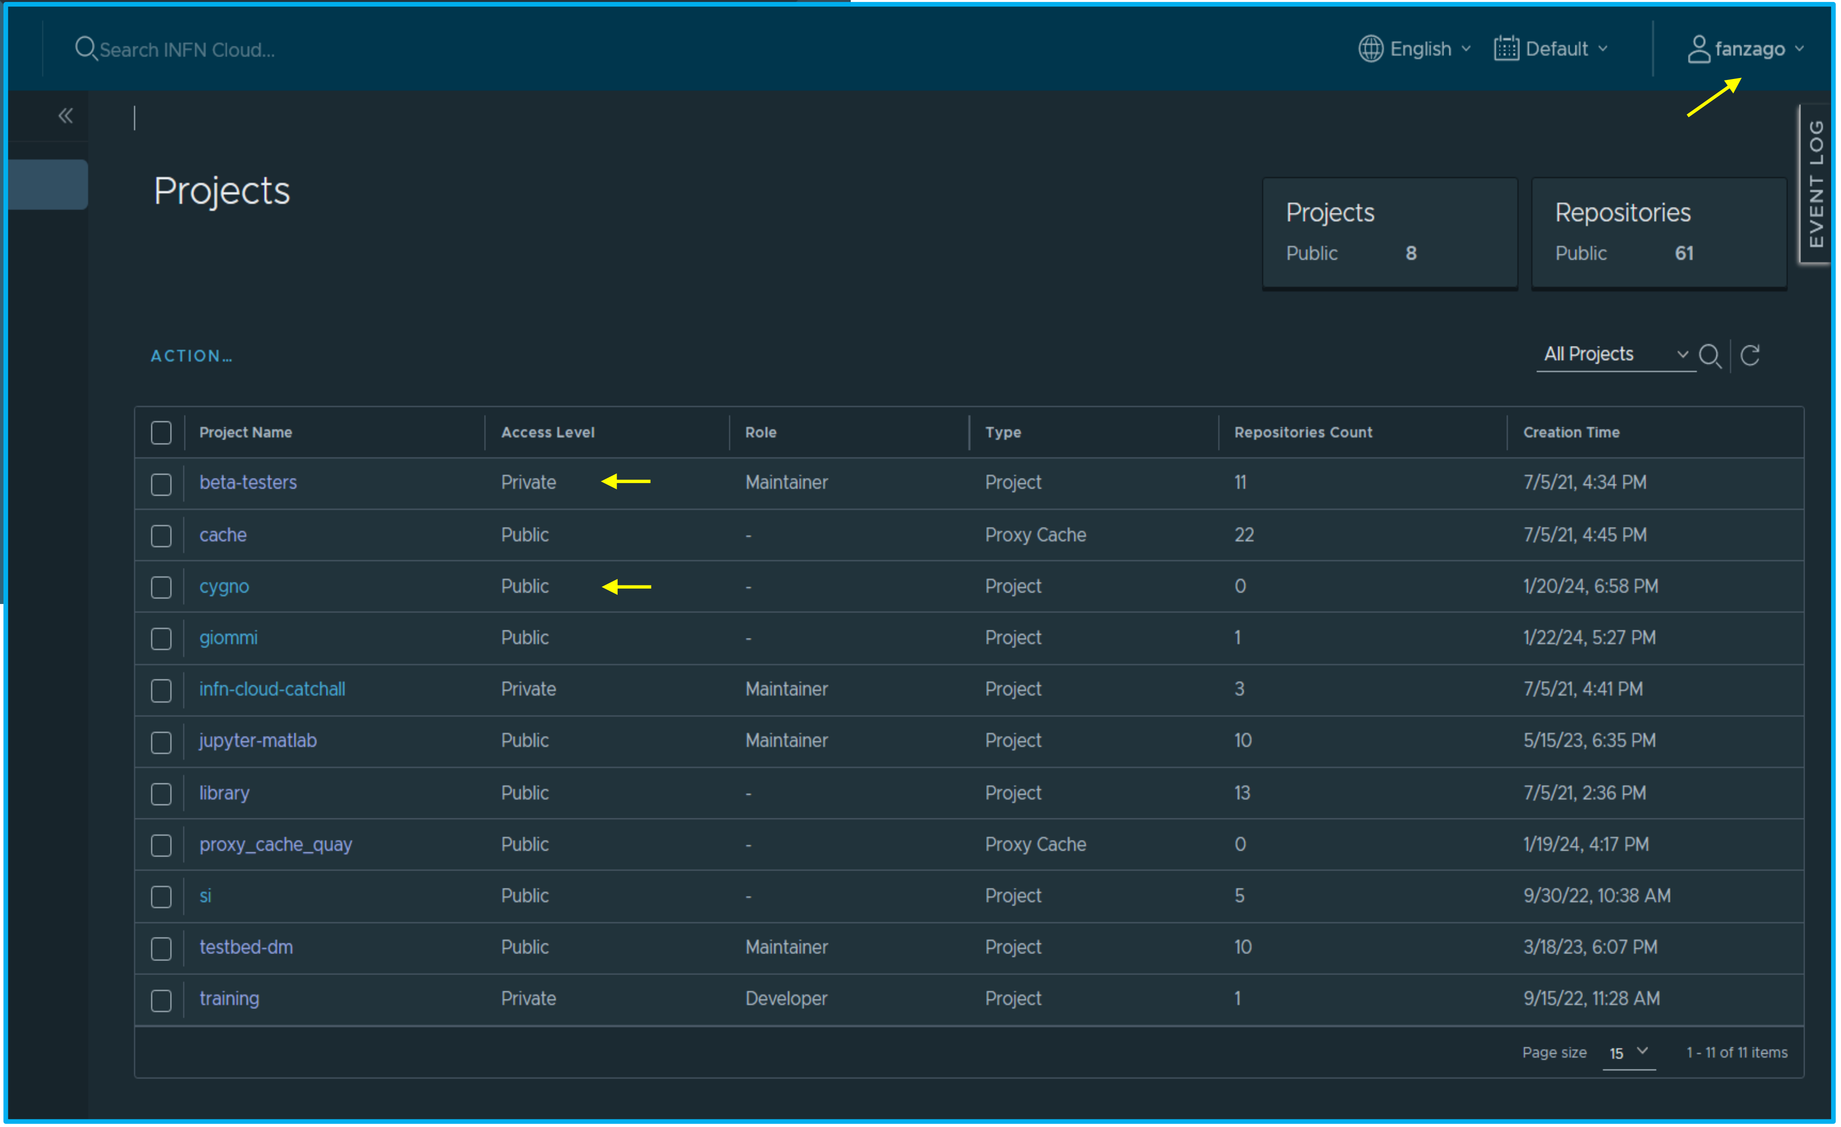

Projects

Projects	
Public	8

Repositories	
Public	61

	Project Name	Access Level	Role	Type	Repositories Count	Creation Time
□	beta-testers	Private	Maintainer	Project	11	7/5/21, 4:34 PM
□	cache	Public	-	Proxy Cache	22	7/5/21, 4:45 PM
□	cygno	Public	-	Project	0	1/20/24, 6:58 PM
□	giommi	Public	-	Project	1	1/22/24, 5:27 PM
□	infn-cloud-catchall	Private	Maintainer	Project	3	7/5/21, 4:41 PM
□	jupyter-matlab	Public	Maintainer	Project	10	5/15/23, 6:35 PM
□	library	Public	-	Project	13	7/5/21, 2:36 PM
□	proxy_cache_quay	Public	-	Proxy Cache	0	1/19/24, 4:17 PM
□	si	Public	-	Project	5	9/30/22, 10:38 AM
□	testbed-dm	Public	Maintainer	Project	10	3/18/23, 6:07 PM
□	training	Private	Developer	Project	1	9/15/22, 11:28 AM

Users have different abilities depending on the role they have in a project.

Limited guest, Guest, Developer, Master, Maintainer, Project Admin

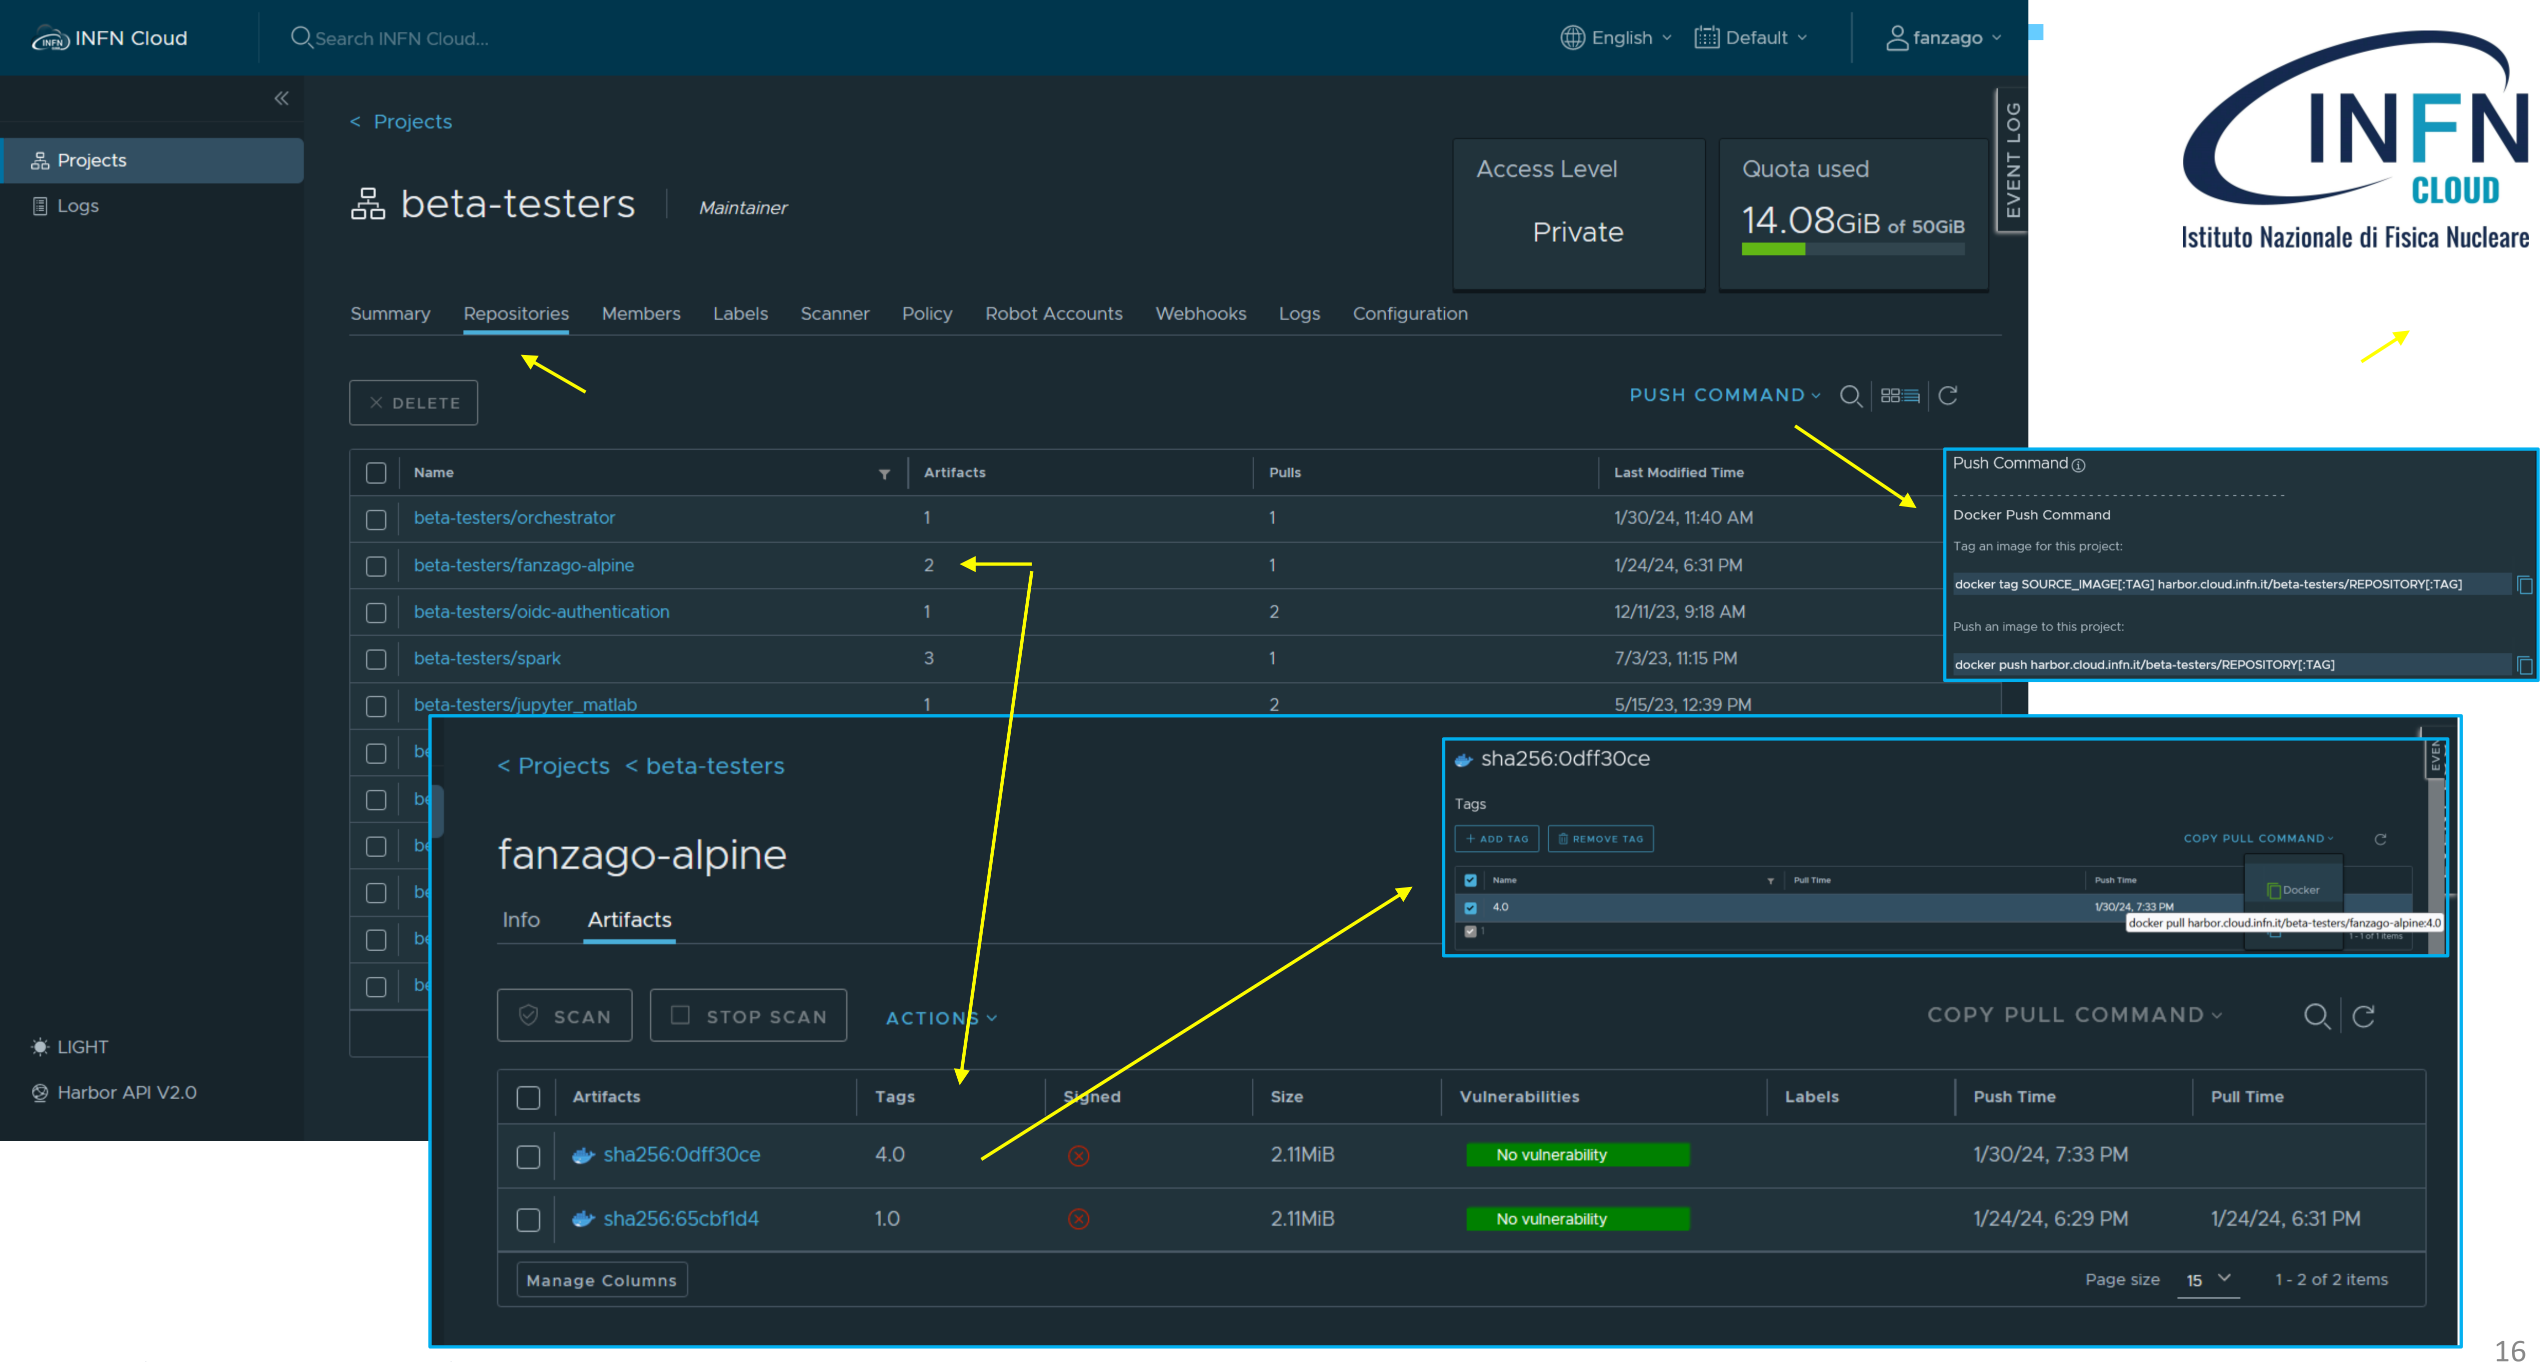

< Projects

beta-testers | *Maintainer*

Access Level	Quota used
Private	14.08GiB of 50GiB

□	Name	Artifacts	Pulls	Last Modified Time
□	beta-testers/orchestrator	1	1	1/30/24, 11:40 AM
□	beta-testers/fanzago-alpine	2	1	1/24/24, 6:31 PM
□	beta-testers/oidc-authentication	1	2	12/11/23, 9:18 AM
□	beta-testers/spark	3	1	7/3/23, 11:15 PM
□	beta-testers/jupyter_matlab	1	2	5/15/23, 12:39 PM

Push Command ⓘ

Docker Push Command

Tag an image for this project:

```
docker tag SOURCE_IMAGE[:TAG] harbor.cloud.infn.it/beta-testers/REPOSITORY[:TAG]
```

Push an image to this project:

```
docker push harbor.cloud.infn.it/beta-testers/REPOSITORY[:TAG]
```

< Projects < beta-testers

fanzago-alpine

□	Artifacts	Tags	Signed	Size	Vulnerabilities	Labels	Push Time	Pull Time
□	sha256:0dff30ce	4.0	⊗	2.11MiB	No vulnerability		1/30/24, 7:33 PM	
□	sha256:65cbf1d4	1.0	⊗	2.11MiB	No vulnerability		1/24/24, 6:29 PM	1/24/24, 6:31 PM

sha256:0dff30ce

Tags

☑	Name	Pull Time	Push Time
☑	4.0		1/30/24, 7:33 PM

```
docker pull harbor.cloud.infn.it/beta-testers/fanzago-alpine:4.0
```

Federica Fanzago INFN Padova

Harbor: push and pull

- How to push and pull of files to/from repos
  - Needed a terminal with docker up and running
  - Login into the service via docker command using username and password (CLI secret) shown clicking on username and then “user profile”

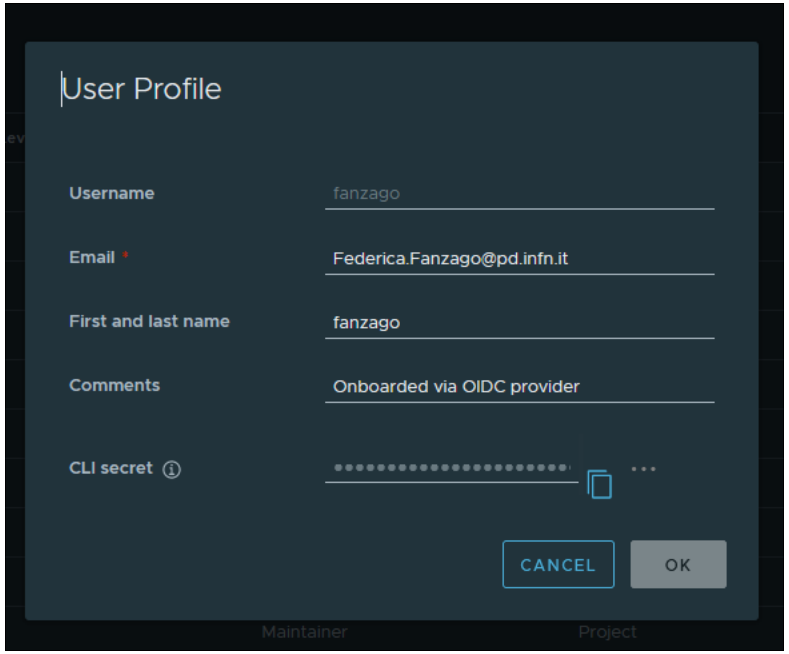

User Profile

Username	fanzago
Email *	Federica.Fanzago@pd.infn.it
First and last name	fanzago
Comments	Onboarded via OIDC provider
CLI secret	••••••••••••••••••••••••

```
## login
$ docker login harbor.cloud.infn.it
## tag the image
$ docker tag SOURCE_IMAGE[:TAG] harbor.cloud.infn.it/project/REPOSITORY[:TAG]
## push the image
$ docker push harbor.cloud.infn.it/project/REPOSITORY[:TAG]
## pull the image
$ docker pull harbor.cloud.infn.it/project/REPOSITORY[:TAG]
```

Harbor: other details

- Images are stored in the object storage of INFN Cloud, a quota can be guarantee for projects. The bucket is unique.
- Proxy cache configured:
  - when a pull request comes to a proxy cache project, if the image is not cached, Harbor pulls the image from the target registry and serves the pull command as if it is a local image from the proxy cache project

```
$ harbor.cloud.infn.it/cache/<docker_repo_name>/<image>:<tag>
```

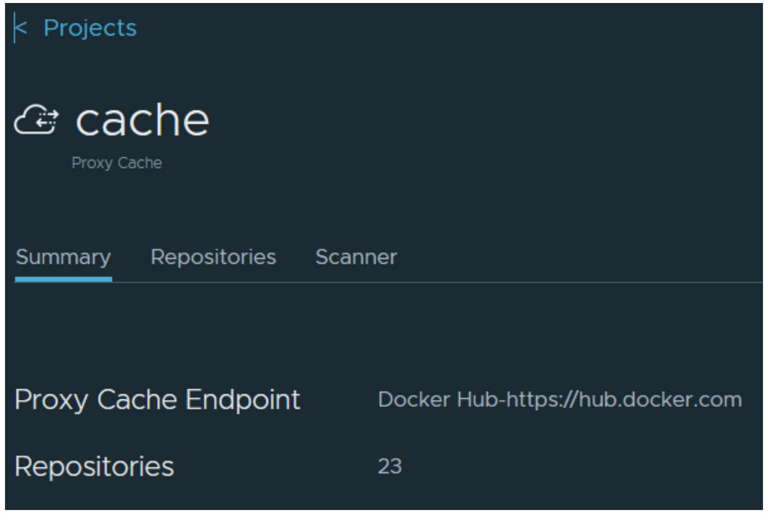

Proxy Cache Endpoint Docker Hub-https://hub.docker.com

Repositories 23

Notebook as a service (NaaS)

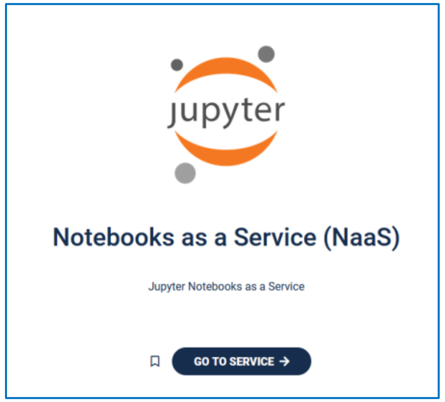

Notebooks as a Service (NaaS)

Jupyter Notebooks as a Service

Service
under maintenance

- https://hub.cloud.infn.it
- https://guides.cloud.infn.it/docs/users-guides/en/latest/users_guides/centralised/naas.html
- It is installed on the backbone (CNAF and Bari), high availability
  - JupyterHub server
- After login through IAM, user starts the “personal” notebook using the docker image made available for the Jupyter Notebook
  - INFN-Cloud base image: harbor.cloud.infn.it/datacloud-templates/jaas_user_containers:1.3.0-2
- Automatic mount of S3 RGW object storage user area (cloud-storage)

Jupyter notebook

- Project started in 2014 by the Jupyter project organization https://jupyter.org
- Open-source web application that provides an interactive environment to create documents called notebooks (.ipynb)
  - A notebook can contain formatted text, equations, images and code
  - The Jupyter notebook provides interpreters for various programming languages as python, R, Julia (kernels, default ipython)
  - The interactive code is executed via browser
  - Single user, notebooks can be exported and shared via git, email…
- Very useful for data science and scientific computing

- Analysis facility
- AI_INFN
- CYGNO
- QUAX …

JupyterLab and JupyterHub

- The JupyterLab is a more extensible and composable interactive computing interface for more complex workflows.
  - It is an integrated environment that aggregates in a single interface more tools, included Notebook.
- The JupyterHub allows group of users to use “private” Jupyter notebook server on shared resources.
- Some INFN Cloud services are implementing Jupyter (hub + notebook (lab))

INFN Cloud services based on Jupyter

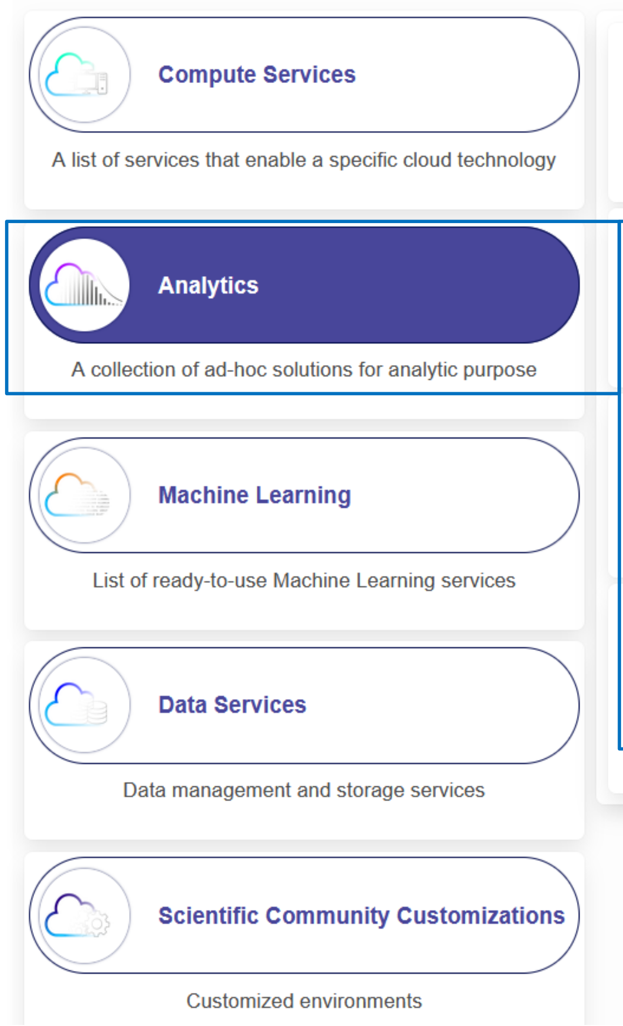

Compute Services

A list of services that enable a specific cloud technology

Analytics

A collection of ad-hoc solutions for analytic purpose

Machine Learning

List of ready-to-use Machine Learning services

Data Services

Data management and storage services

Scientific Community Customizations

Customized environments

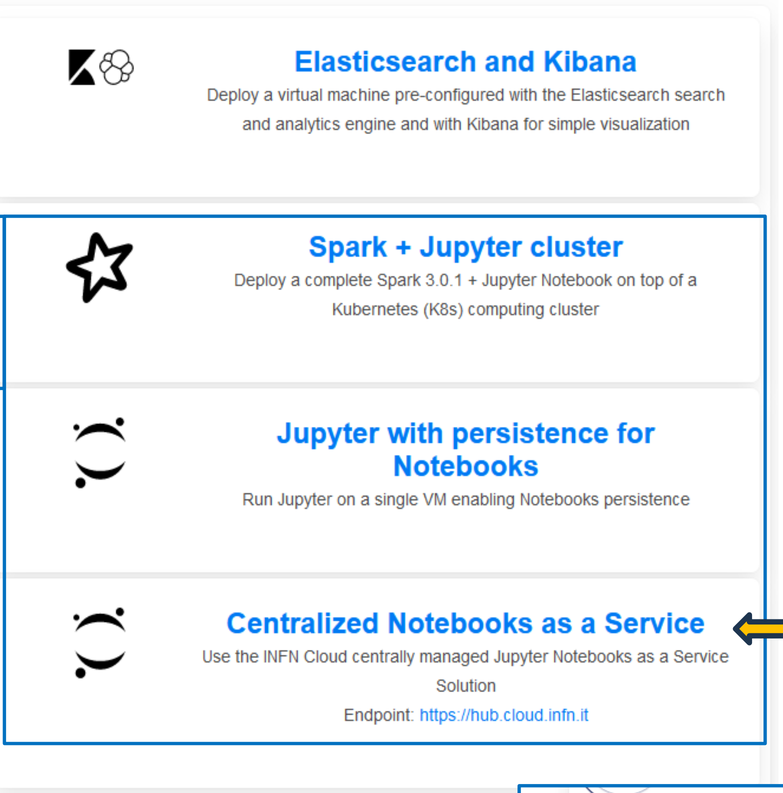

Elasticsearch and Kibana

Deploy a virtual machine pre-configured with the Elasticsearch search and analytics engine and with Kibana for simple visualization

Spark + Jupyter cluster

Deploy a complete Spark 3.0.1 + Jupyter Notebook on top of a Kubernetes (K8s) computing cluster

Jupyter with persistence for Notebooks

Run Jupyter on a single VM enabling Notebooks persistence

Centralized Notebooks as a Service

Use the INFN Cloud centrally managed Jupyter Notebooks as a Service Solution

Endpoint: https://hub.cloud.infn.it

https://www.cloud.infn.it/service-catalogue

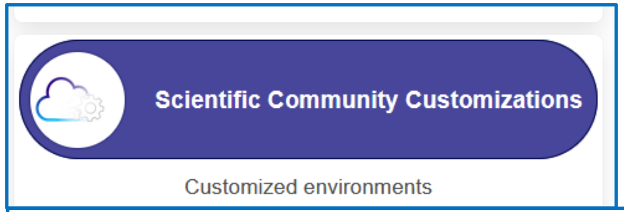

Scientific Community Customizations

Customized environments

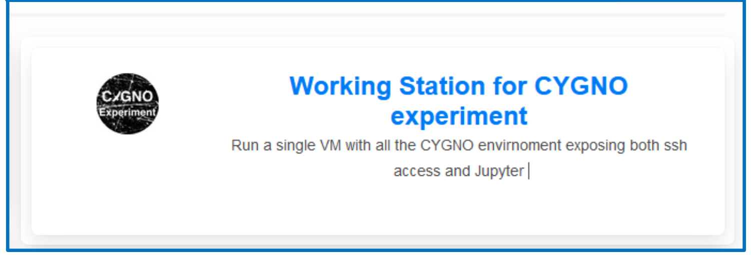

Working Station for CYGNO experiment

Run a single VM with all the CYGNO envirnoment exposing both ssh access and Jupyter

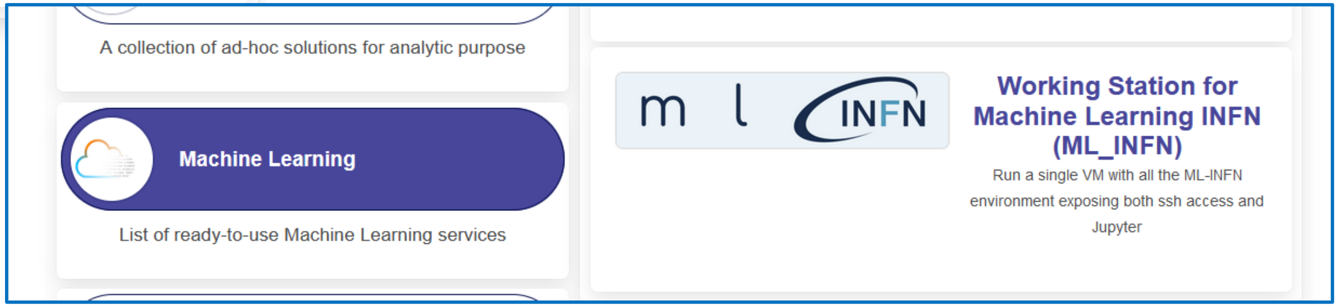

A collection of ad-hoc solutions for analytic purpose

Machine Learning

List of ready-to-use Machine Learning services

Working Station for Machine Learning INFN (ML_INFN)

Run a single VM with all the ML-INFN environment exposing both ssh access and Jupyter

NaaS Hub Interface

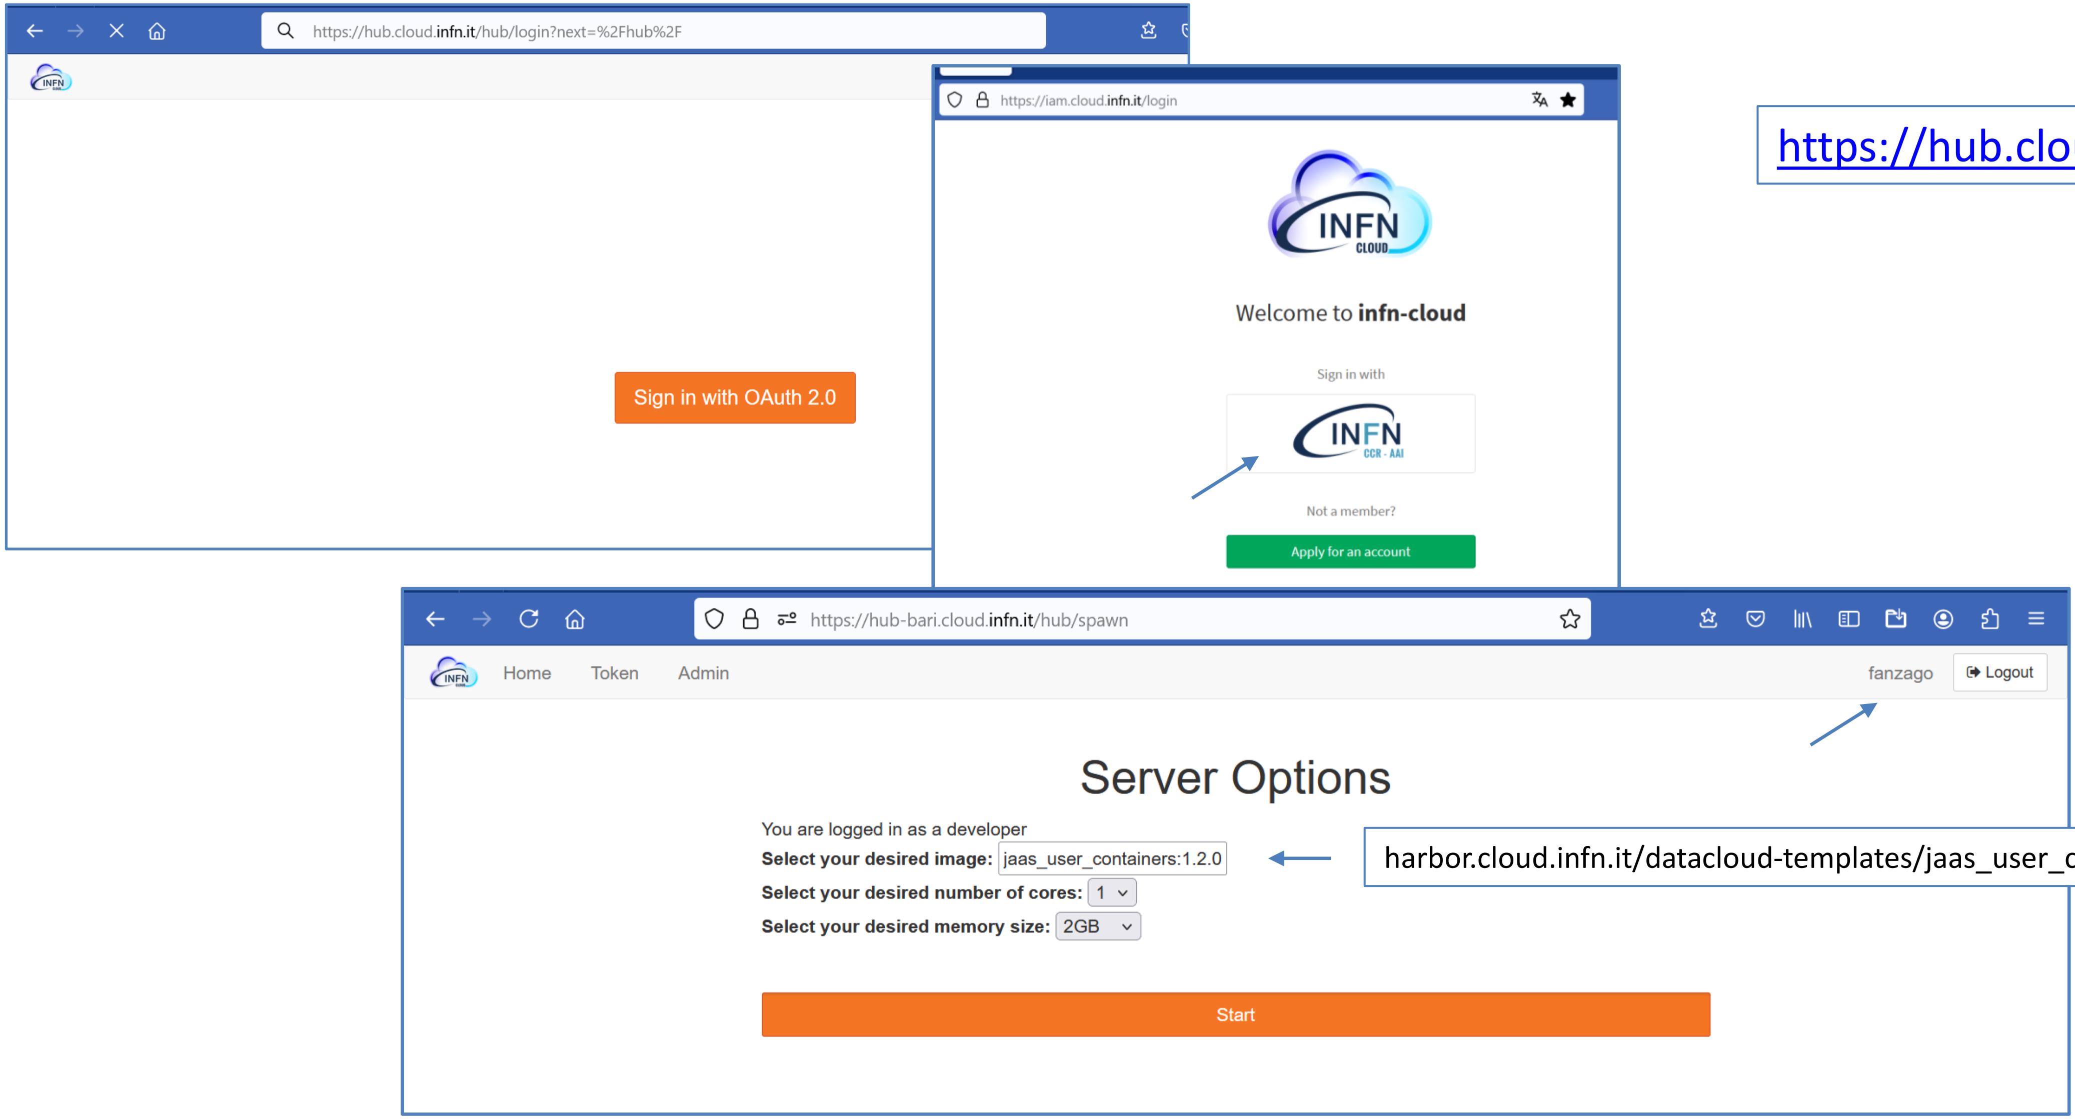

https://hub.cloud.infn.it/hub/login?next=%2Fhub%2F

https://iam.cloud.infn.it/login

Welcome to **infn-cloud**

https://hub-bari.cloud.infn.it/hub/spawn

Server Options

You are logged in as a developer

Select your desired image: jaas_user_containers:1.2.0

Select your desired number of cores: 1

Select your desired memory size: 2GB

https://hub.cloud.infn.it

harbor.cloud.infn.it/datacloud-templates/jaas_user_containers:1.3.0-2

Jupyter notebook interface

- The Jupyter Notebook interface allows the management of kernels and notebooks

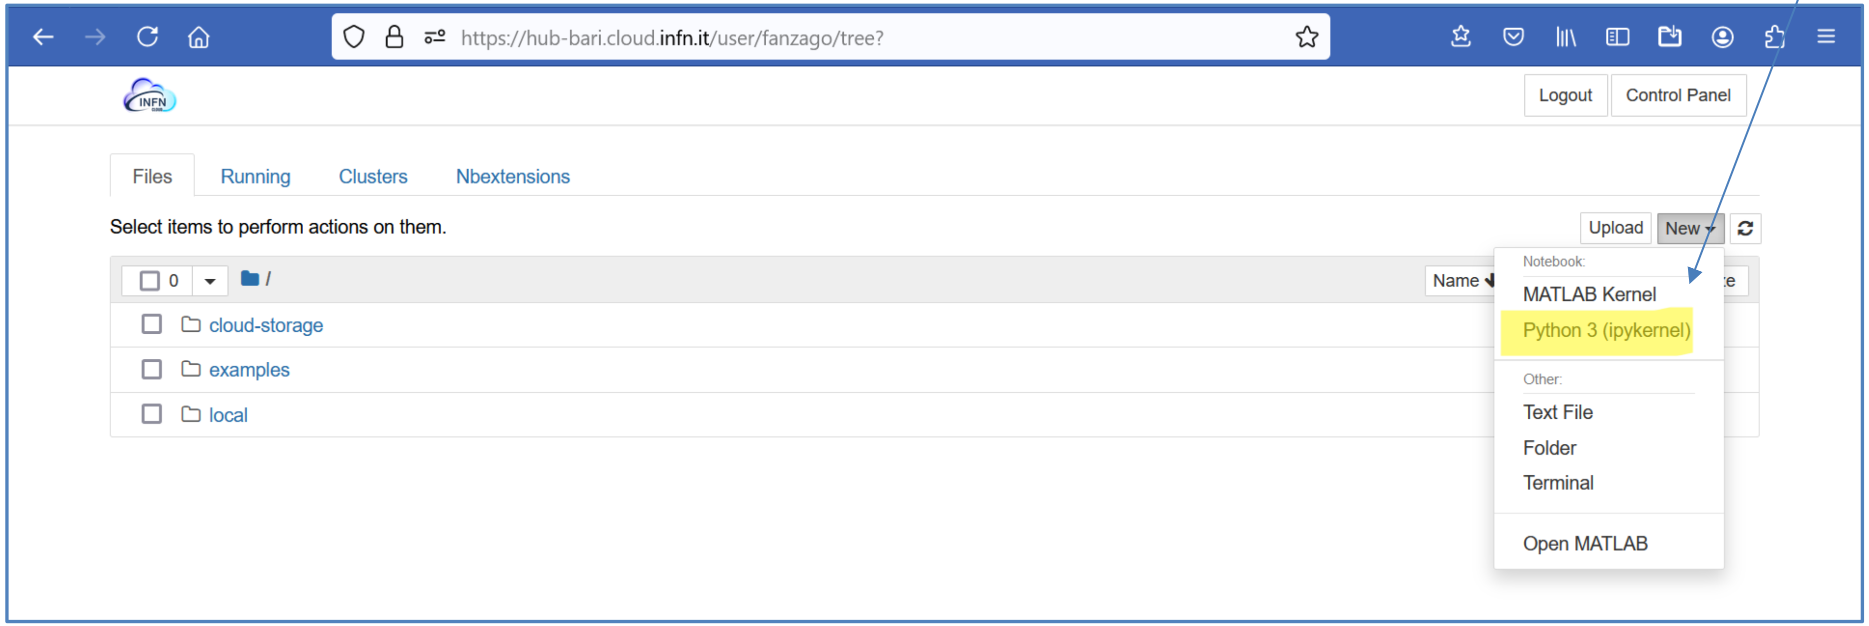

https://hub-bari.cloud.infn.it/user/fanzago/tree?

Select items to perform actions on them.

□ cloud-storage

□ examples

□ local

Jupyter notebook interface

- The interactive work is done inside “cells” that can contain and execute code, commands and formatted text
- Each cell is executed in sequential way. During the execution, an [*] is shown. After the execution, its sequential number [x].

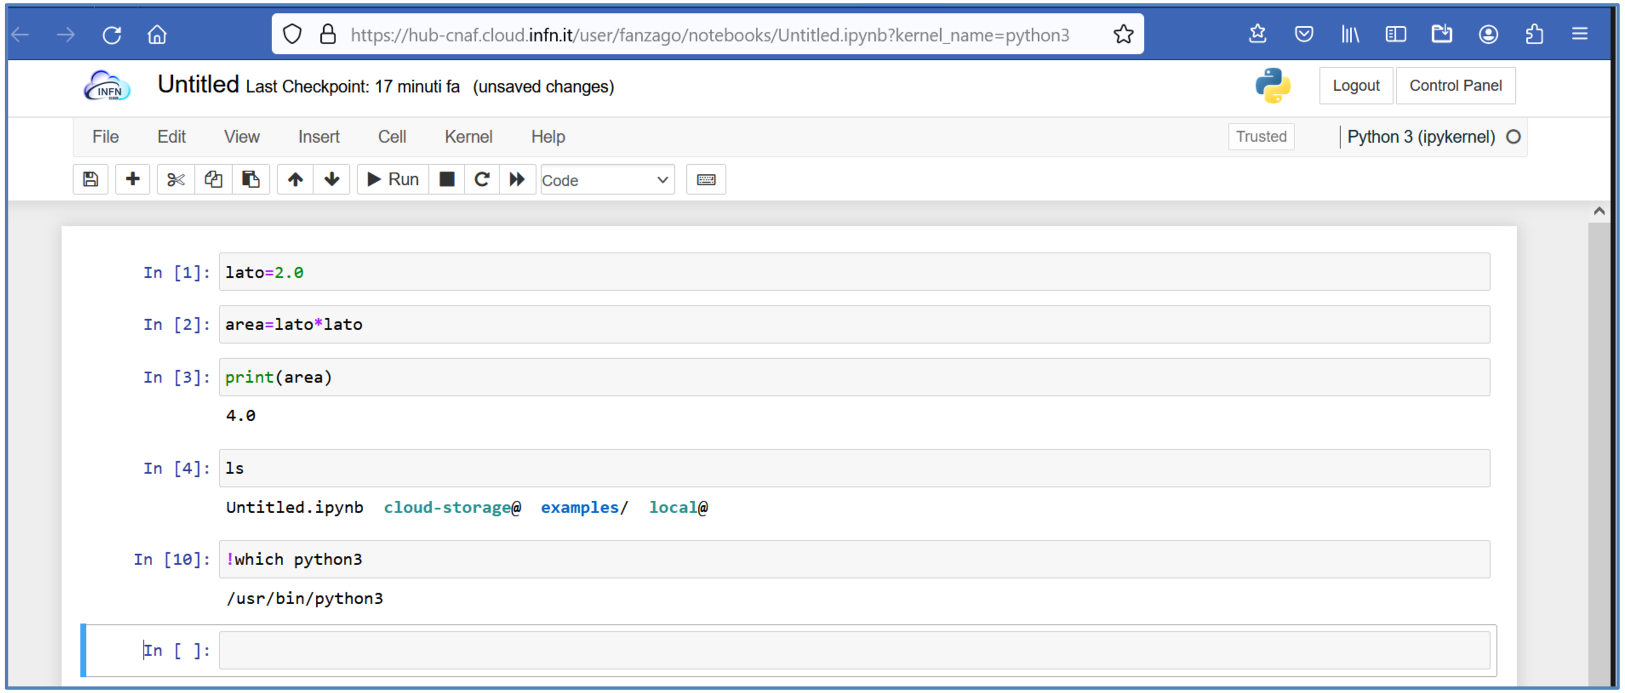
https://hub-cnaf.cloud.infn.it/user/fanzago/notebooks/Untitled.ipynb?kernel_name=python3
Untitled Last Checkpoint: 17 minuti fa (unsaved changes)
Logout
Control Panel
File
Edit
View
Insert
Cell
Kernel
Help
Trusted
Python 3 (ipykernel)
Run
Code

```
In [1]: lato=2.0

In [2]: area=lato*lato

In [3]: print(area)
        4.0

In [4]: ls
        Untitled.ipynb  cloud-storage@  examples/  local@

In [10]: !which python3
         /usr/bin/python3

In [ ]:
```

Inside cells, select code:
Shift + enter - cell execution
ESC a,b - add cells
ESC dd - delete cell

%lsmagic
%run <script.py>
%load <script.py>
%who - list of python variables
%pinfo <var> - description of variable
%env - environment variables

Pip freeze - python modules installed
Pip install <module>

Where to store notebook documents?

- Saving document: the default path is the local “jupyter-workspace” directory of the system where the notebook server is running
  - It runs in a container: data are lost if the container is removed
- Need to have a “permanent” space where to store them and guarantee persistence of data.
  - In INFN Cloud the object storage meets this request

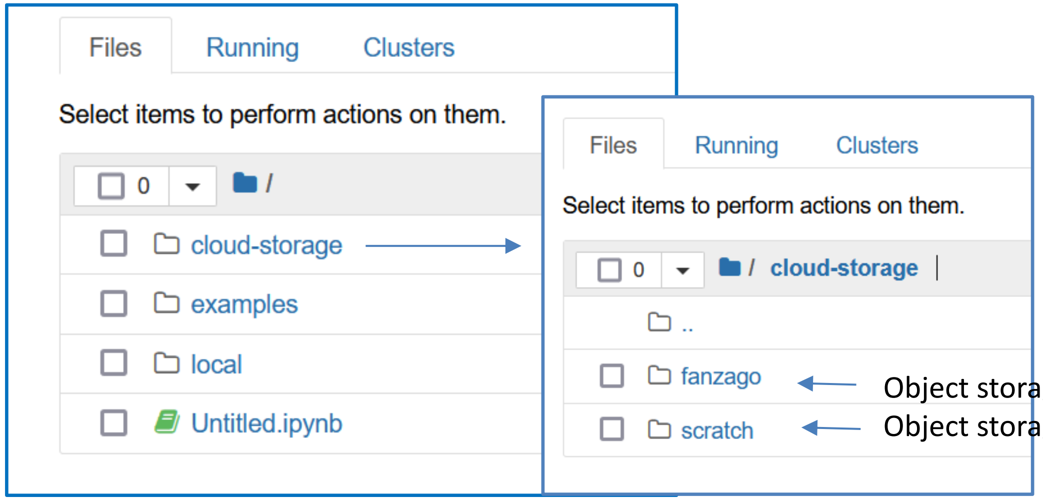

Select items to perform actions on them.

□ 0 ▾ 📁 /

□ 🗀 cloud-storage

□ 🗀 examples

□ 🗀 local

□ Untitled.ipynb

Select items to perform actions on them.

□ 0 ▾ 📁 / cloud-storage

🗀 ..

□ 🗀 fanzago

□ 🗀 scratch

Object storage user bucket

Object storage scratch area

Local = directory «local» in the server (bari or cnaf)

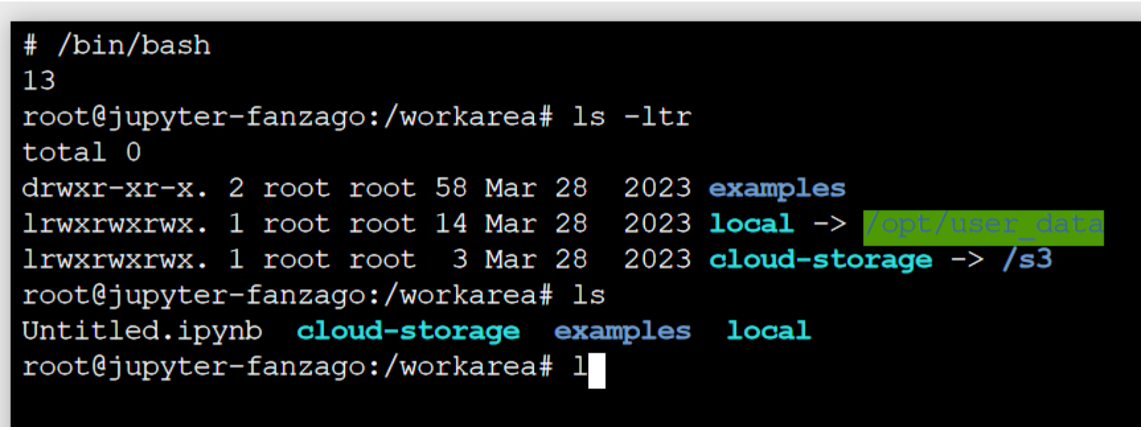

```
# /bin/bash
13
root@jupyter-fanzago:/workarea# ls -ltr
total 0
drwxr-xr-x. 2 root root 58 Mar 28  2023 examples
lrwxrwxrwx. 1 root root 14 Mar 28  2023 local -> /opt/user_data
lrwxrwxrwx. 1 root root  3 Mar 28  2023 cloud-storage -> /s3
root@jupyter-fanzago:/workarea# ls
Untitled.ipynb  cloud-storage  examples  local
root@jupyter-fanzago:/workarea# l
```

Persistent directory for data

- All the services running notebooks mount these areas, so they are visible in the file system as posix directories.
  - cloud-storage dir

Using sts-wire, fuse and Rclone

- The personal bucket can by create by user accessing the Ceph-RGW webui otherwise it is automatically created the first time user instantiates a notebook.

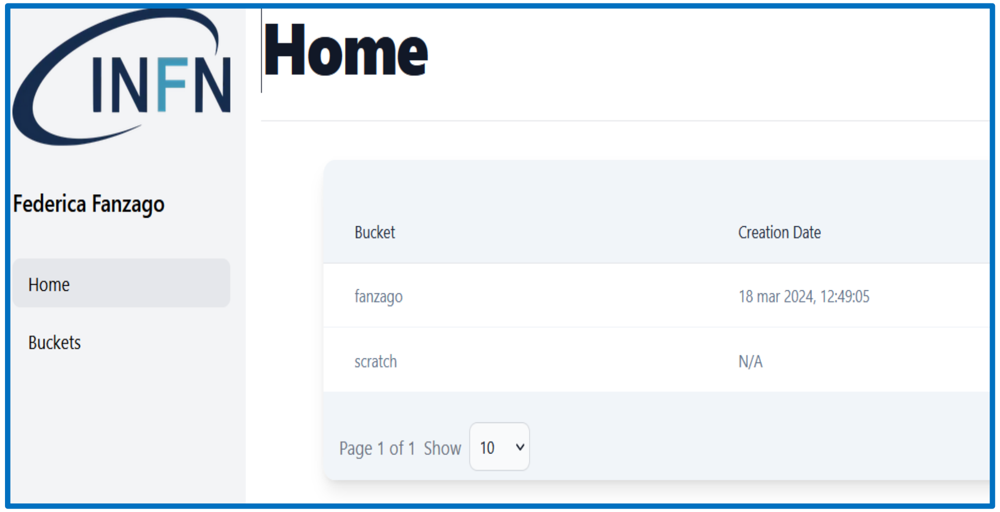

Bucket	Creation Date
fanzago	18 mar 2024, 12:49:05
scratch	N/A

Comparing with “Jupyter with persistence” deployment (self-managed service)

- Login in the INFN dashboard https://my.cloud.infn.it and select the “Jupyter with persistence” button. Then configure the service filling the form.

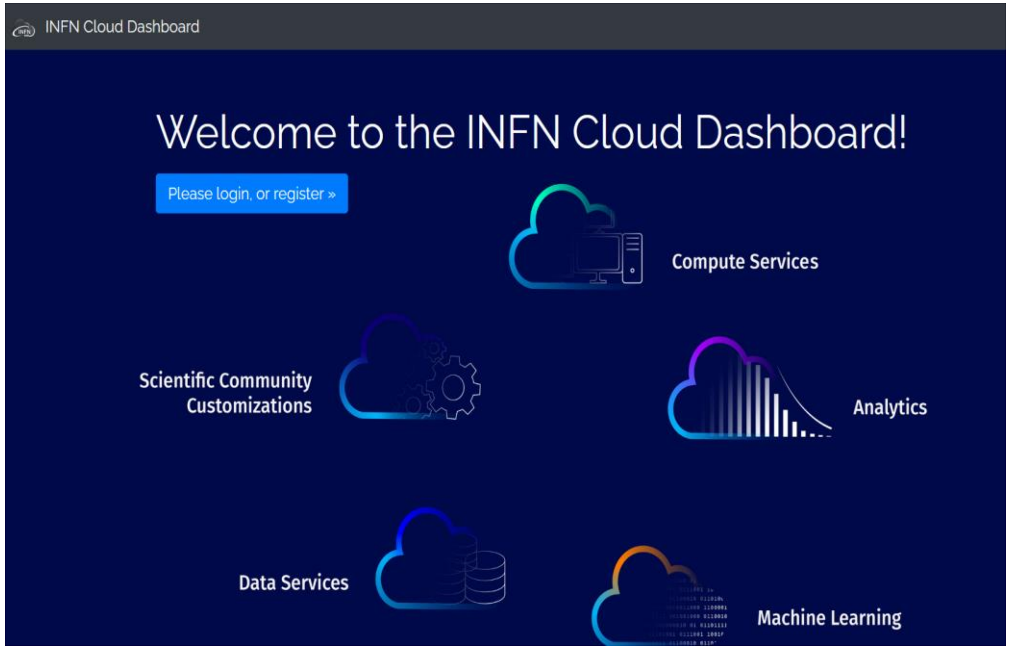

Welcome to the INFN Cloud Dashboard!

Authentication and authorization via IAM

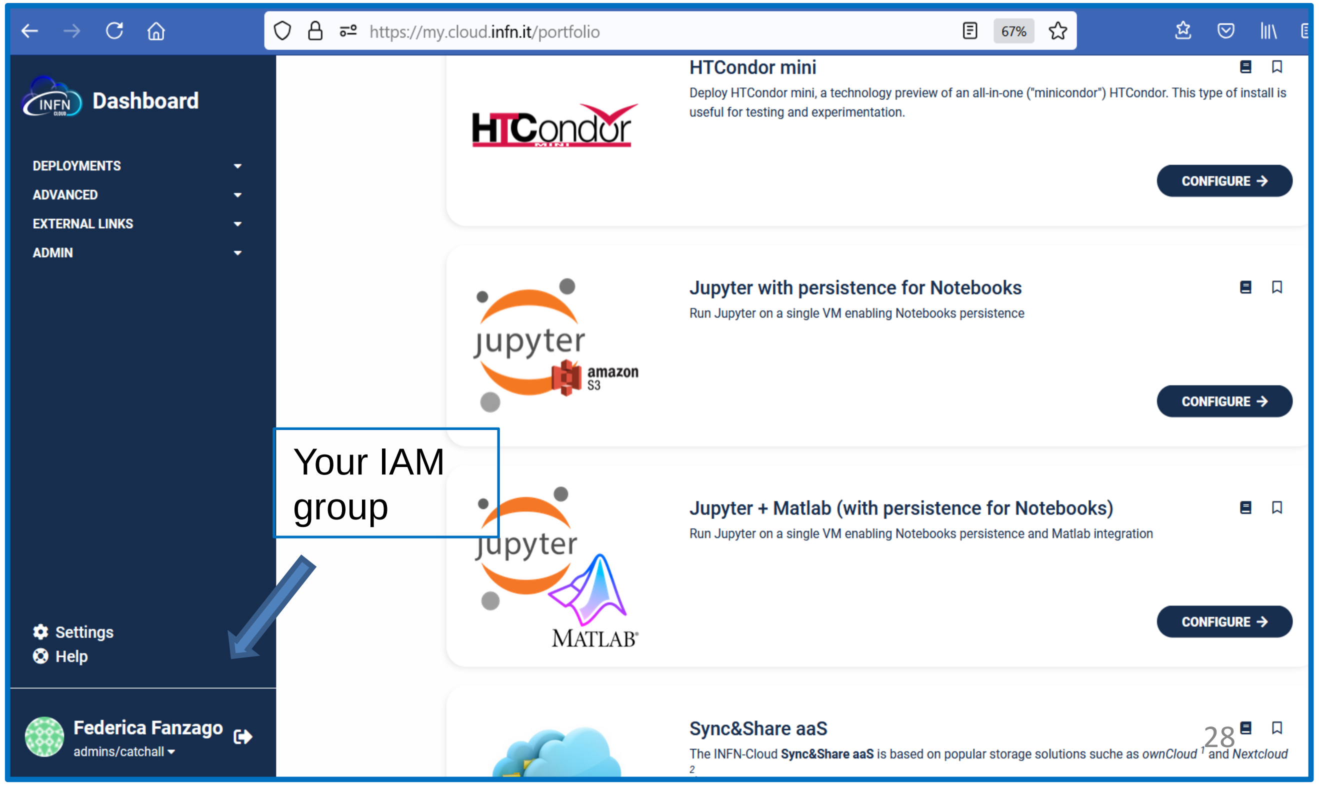

HTCondor mini

Deploy HTCondor mini, a technology preview of an all-in-one ("minicondor") HTCondor. This type of install is useful for testing and experimentation.

Jupyter with persistence for Notebooks

Run Jupyter on a single VM enabling Notebooks persistence

Jupyter + Matlab (with persistence for Notebooks)

Run Jupyter on a single VM enabling Notebooks persistence and Matlab integration

Sync&Share aaS

The INFN-Cloud **Sync&Share aaS** is based on popular storage solutions suche as *ownCloud* [1] and *Nextcloud* [2]

Your IAM group

Filling the form…

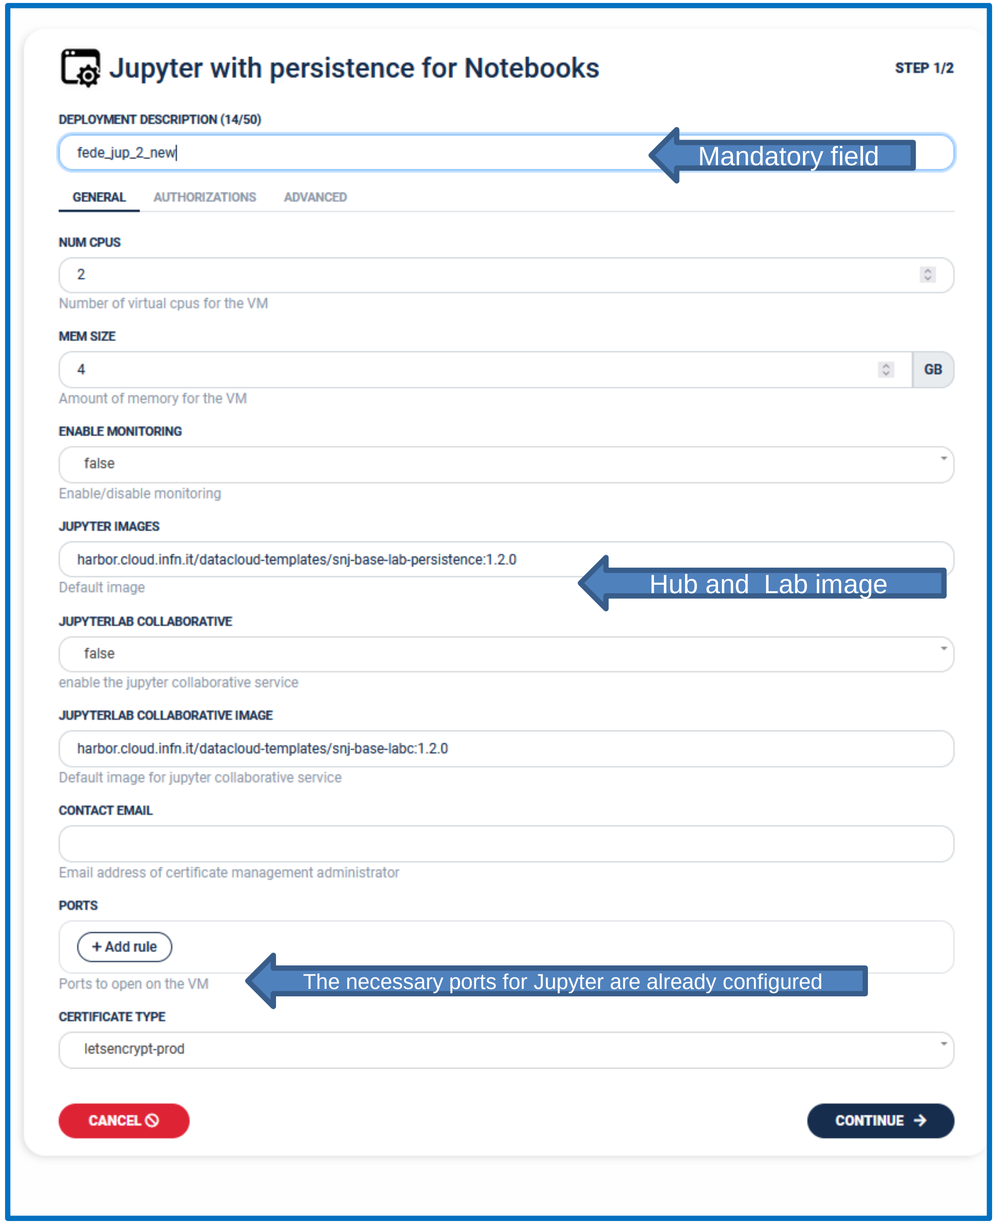

Jupyter with persistence for Notebooks

STEP 1/2

DEPLOYMENT DESCRIPTION (14/50)

fede_jup_2_new

Mandatory field

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

harbor.cloud.infn.it/datacloud-templates/snj-base-lab-persistence:1.2.0

Default image

Hub and Lab image

JUPYTERLAB COLLABORATIVE

false

enable the jupyter collaborative service

JUPYTERLAB COLLABORATIVE IMAGE

harbor.cloud.infn.it/datacloud-templates/snj-base-labc:1.2.0

Default image for jupyter collaborative service

CONTACT EMAIL

Email address of certificate management administrator

PORTS

+ Add rule

Ports to open on the VM

The necessary ports for Jupyter are already configured

CERTIFICATE TYPE

letsencrypt-prod

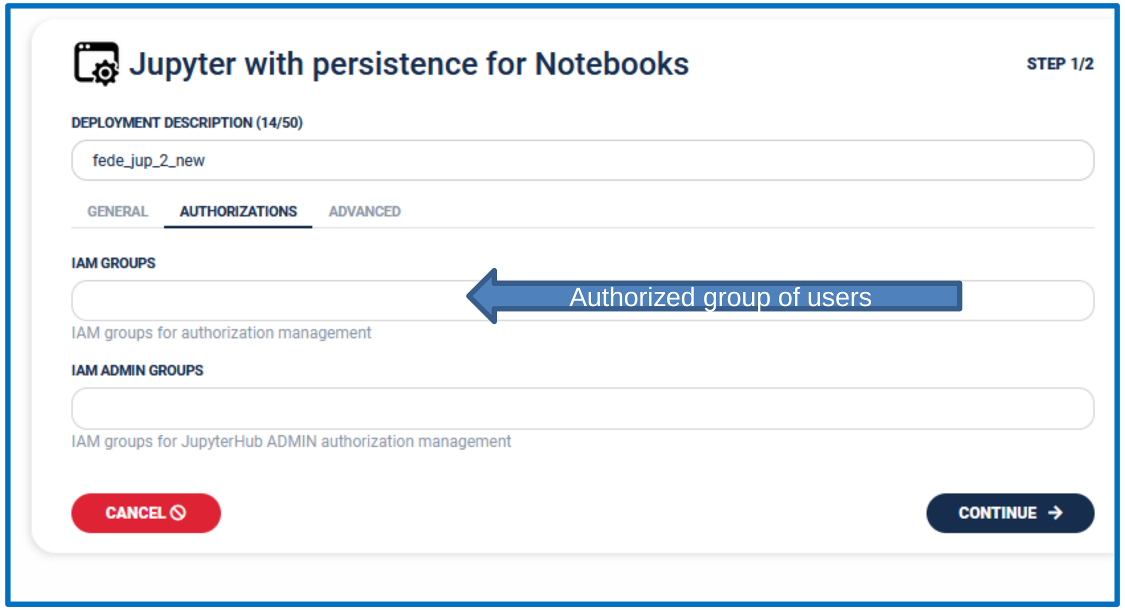

Jupyter with persistence for Notebooks

STEP 1/2

DEPLOYMENT DESCRIPTION (14/50)

fede_jup_2_new

GENERAL | AUTHORIZATIONS | ADVANCED

IAM GROUPS

Authorized group of users

IAM groups for authorization management

IAM ADMIN GROUPS

IAM groups for JupyterHub ADMIN authorization management

User guide:
https://guides.cloud.infn.it/docs/users-guides/en/latest/users_guides/sysadmin/compute/jh_with_persistence.html

…then submit the deployment

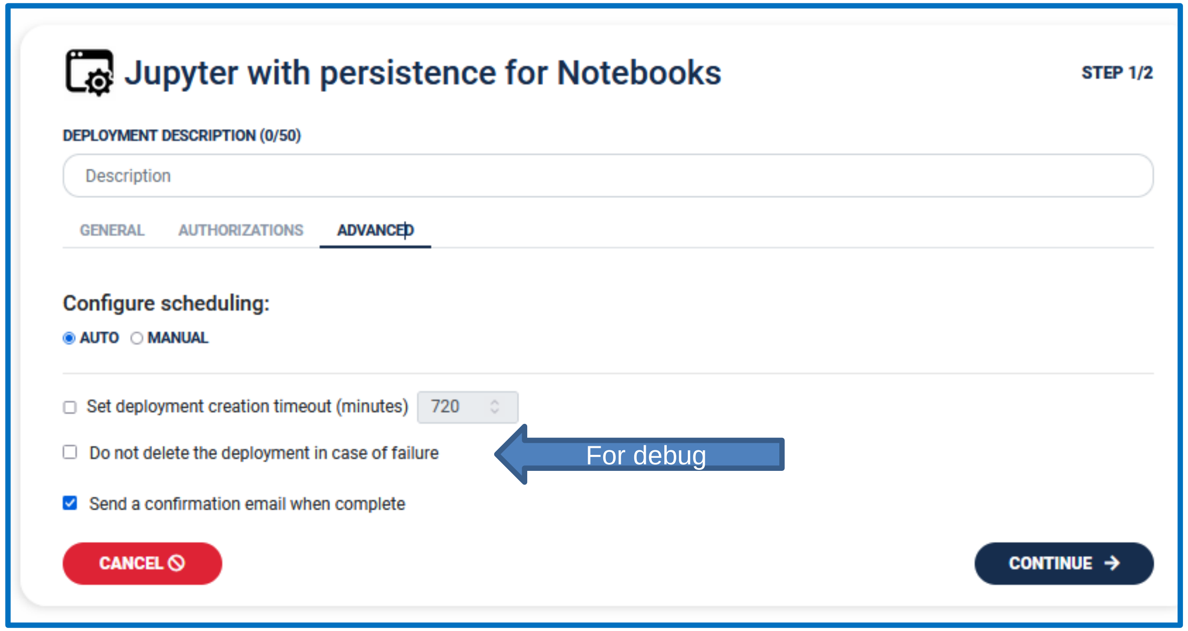

Jupyter with persistence for Notebooks

STEP 1/2

DEPLOYMENT DESCRIPTION (0/50)

Description

Configure scheduling:

● AUTO ○ MANUAL

□ Set deployment creation timeout (minutes) 720

□ Do not delete the deployment in case of failure

For debug

☑ Send a confirmation email when complete

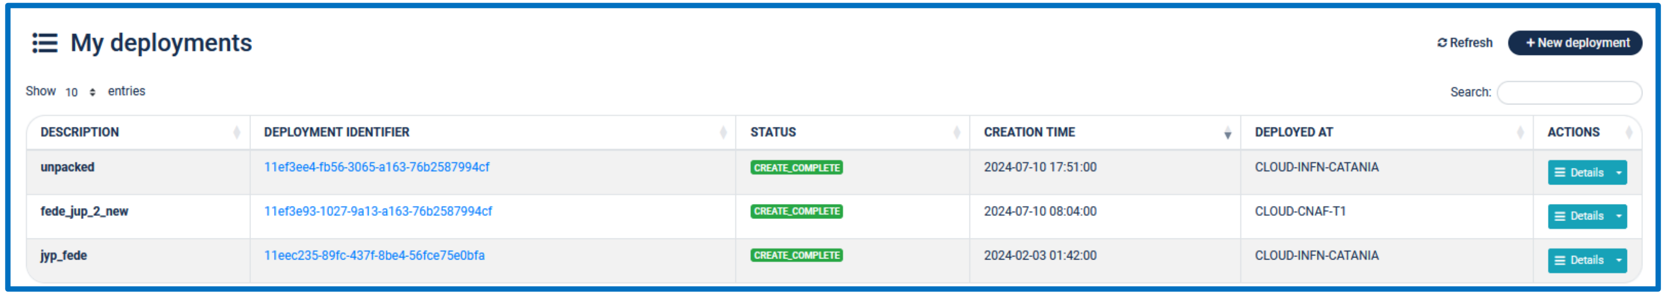

My deployments

Show 10 entries

Search:

DESCRIPTION	DEPLOYMENT IDENTIFIER	STATUS	CREATION TIME	DEPLOYED AT	ACTIONS
unpacked	11ef3ee4-fb56-3065-a163-76b2587994cf	CREATE_COMPLETE	2024-07-10 17:51:00	CLOUD-INFN-CATANIA	Details
fede_jup_2_new	11ef3e93-1027-9a13-a163-76b2587994cf	CREATE_COMPLETE	2024-07-10 08:04:00	CLOUD-CNAF-T1	Details
jyp_fede	11eec235-89fc-437f-8be4-56fce75e0bfa	CREATE_COMPLETE	2024-02-03 01:42:00	CLOUD-INFN-CATANIA	Details

Deployment details

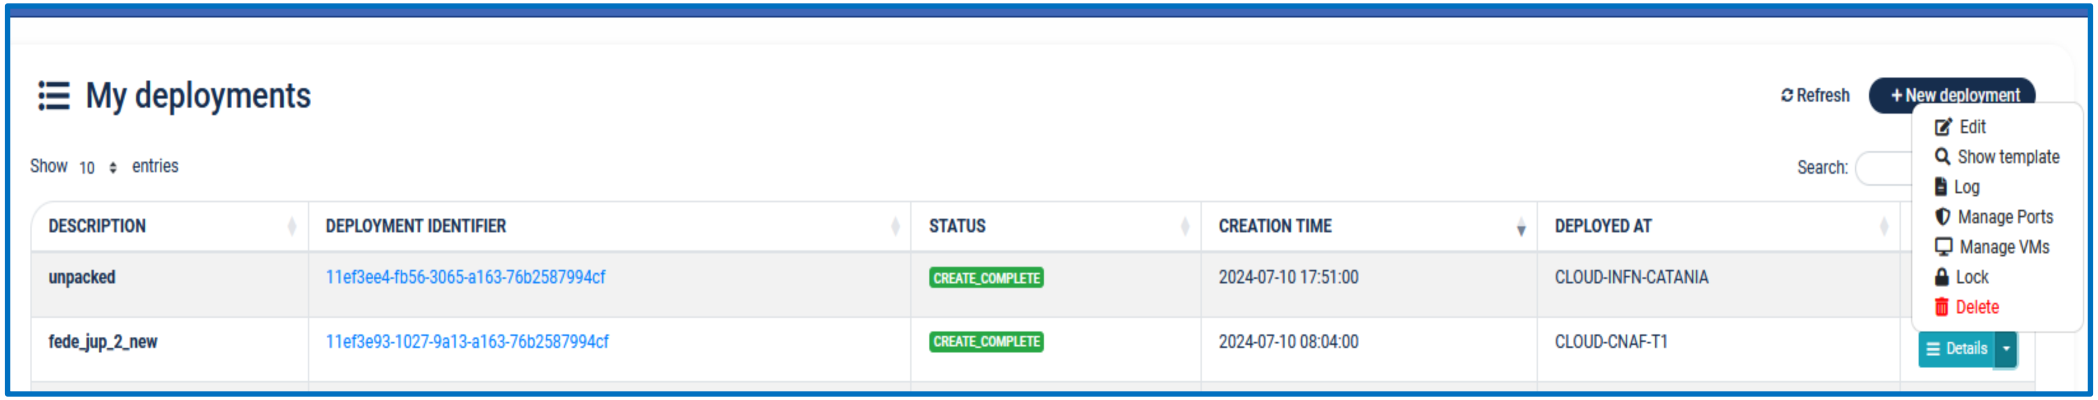

My deployments

Show 10 entries

Search:

DESCRIPTION	DEPLOYMENT IDENTIFIER	STATUS	CREATION TIME	DEPLOYED AT
unpacked	11ef3ee4-fb56-3065-a163-76b2587994cf	CREATE_COMPLETE	2024-07-10 17:51:00	CLOUD-INFN-CATANIA
fede_jup_2_new	11ef3e93-1027-9a13-a163-76b2587994cf	CREATE_COMPLETE	2024-07-10 08:04:00	CLOUD-CNAF-T1

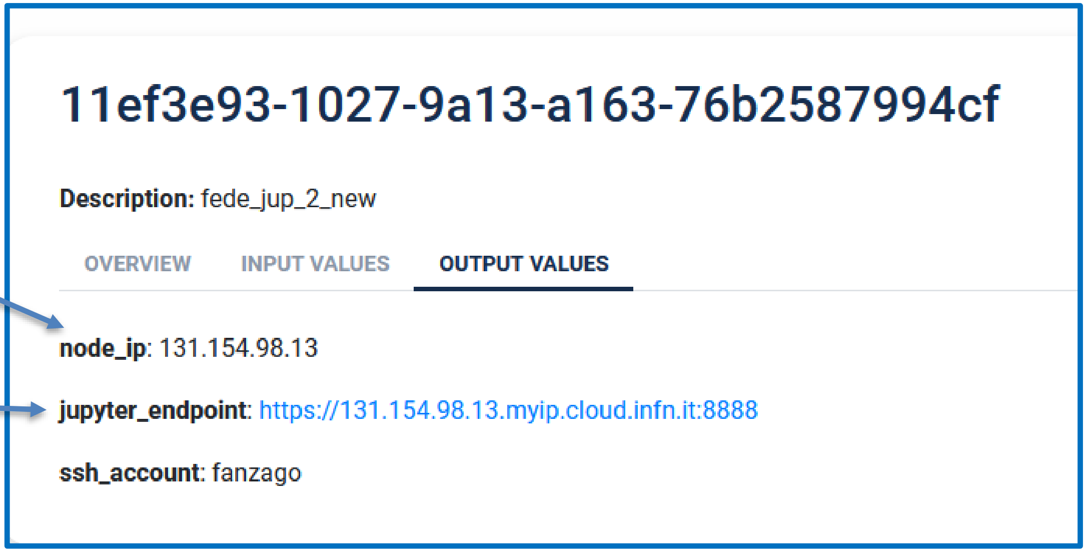

11ef3e93-1027-9a13-a163-76b2587994cf

Description: fede_jup_2_new

node_ip: 131.154.98.13

jupyter_endpoint: https://131.154.98.13.myip.cloud.infn.it:8888

ssh_account: fanzago

Remember your ssh key to access the Vm

JupyterHub url

Your JupyterHub

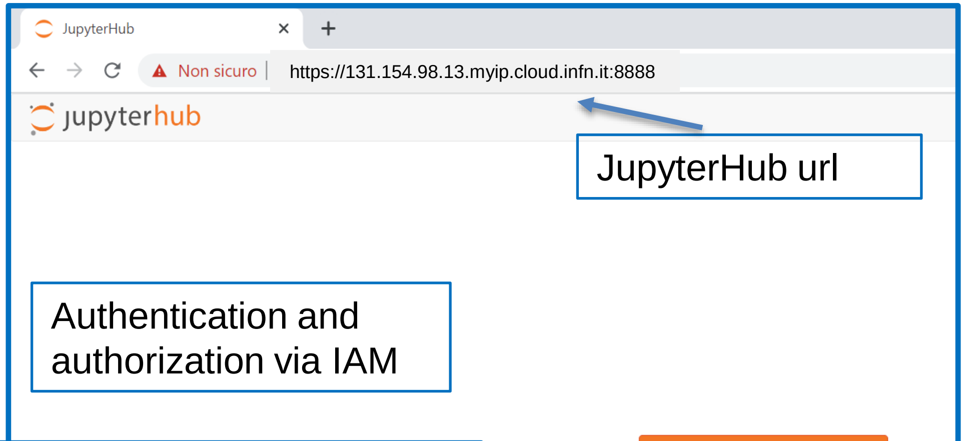

https://131.154.98.13.myip.cloud.infn.it:8888

JupyterHub url

Authentication and authorization via IAM

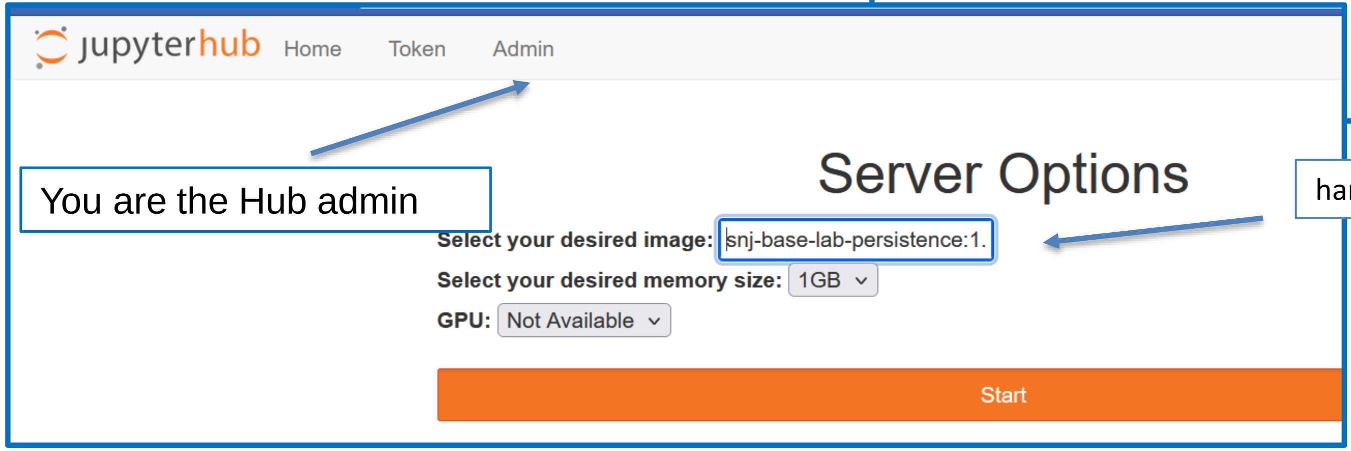

Server Options

Select your desired image: snj-base-lab-persistence:1.

Select your desired memory size: 1GB

GPU: Not Available

You are the Hub admin

harbor.cloud.infn.it/datacloud-templates/snj-base-lab-persistence:1.2.0

Image provided to start the JupyterLab server. You can use your private one built from the default one

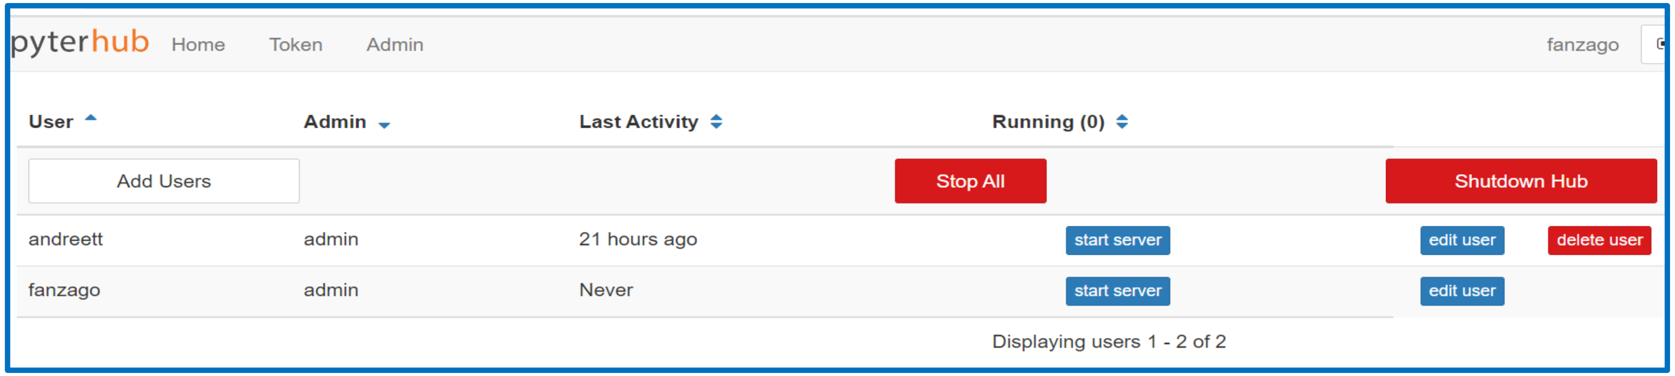

User	Admin	Last Activity	Running (0)		
Add Users			Stop All	Shutdown Hub	
andreett	admin	21 hours ago	start server	edit user	delete user
fanzago	admin	Never	start server	edit user	

Displaying users 1 - 2 of 2

Hub’s admin can control users and notebooks

And your JupyterLab

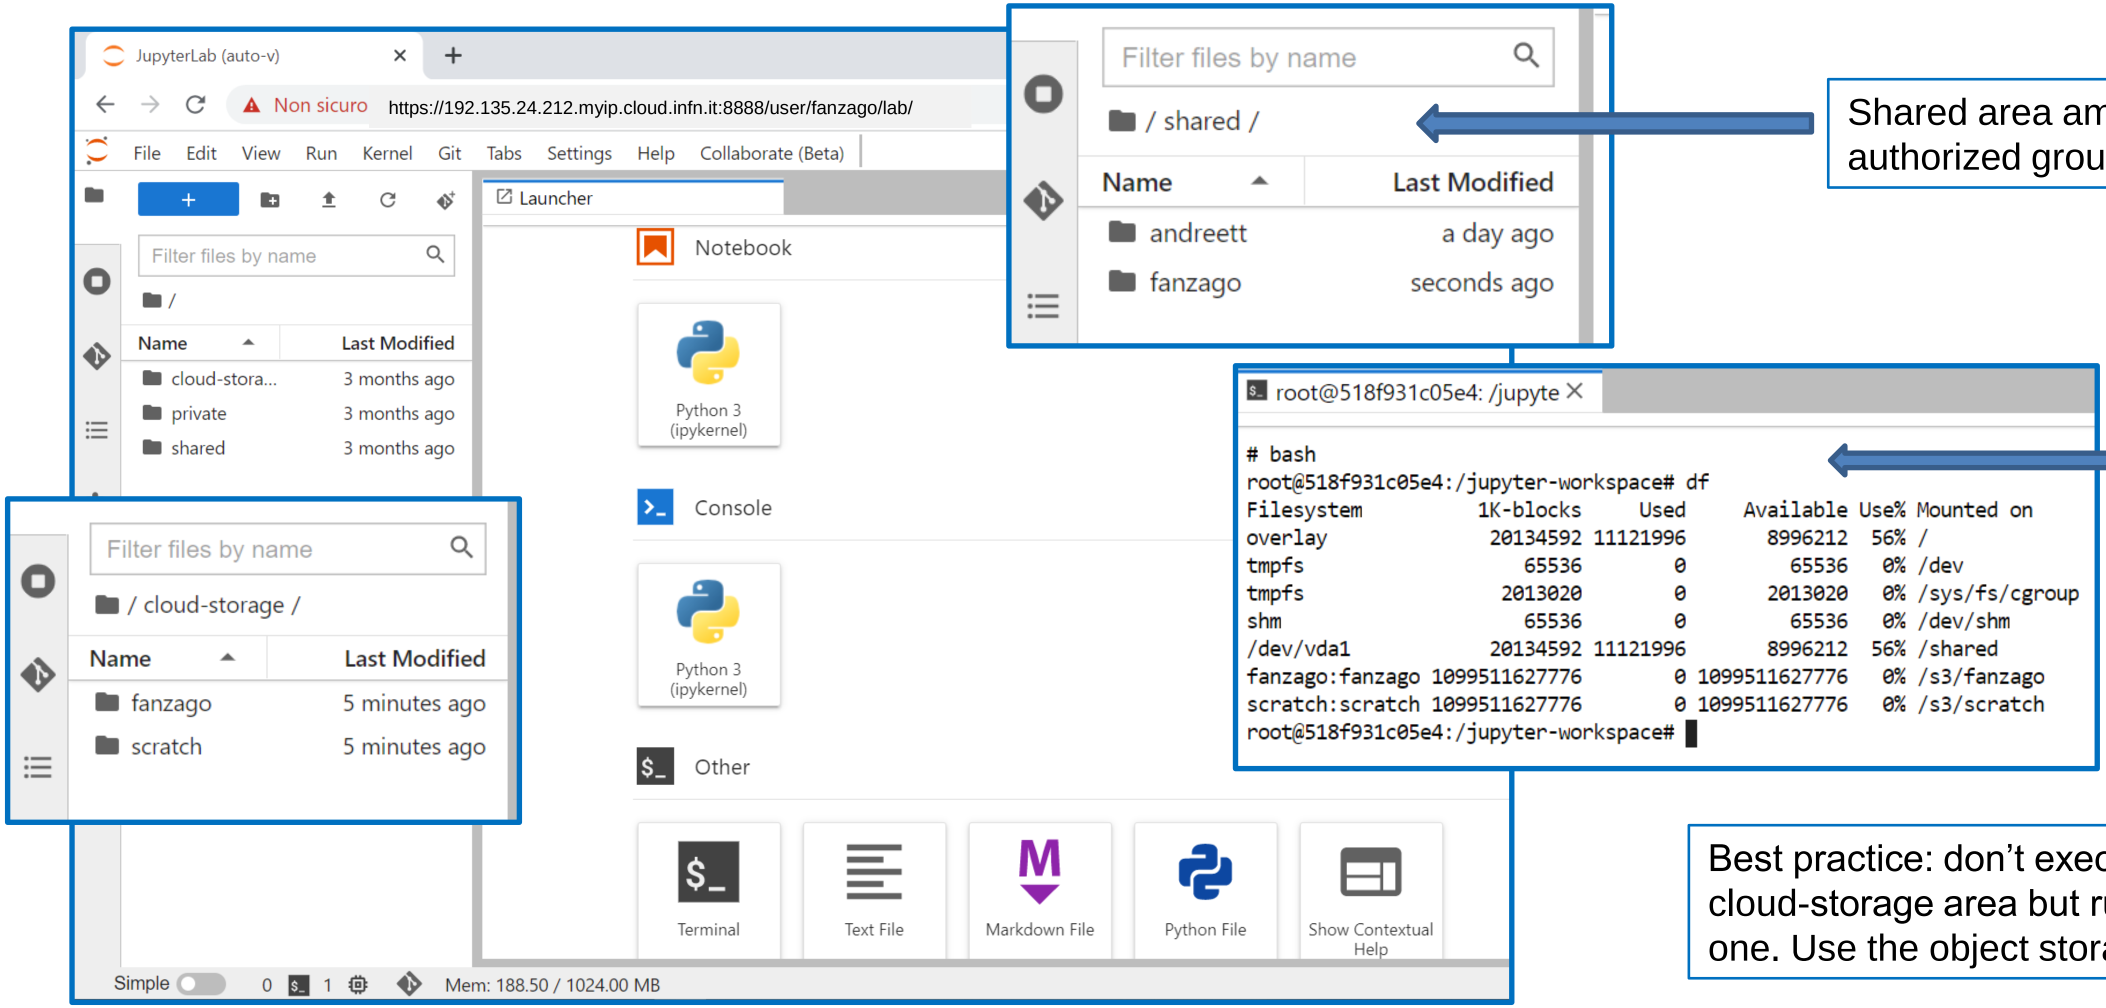

/

Name	Last Modified
cloud-stora...	3 months ago
private	3 months ago
shared	3 months ago

/ shared /

Name	Last Modified
andreett	a day ago
fanzago	seconds ago

/ cloud-storage /

Name	Last Modified
fanzago	5 minutes ago
scratch	5 minutes ago

root@518f931c05e4: /jupyte

```
# bash
root@518f931c05e4:/jupyter-workspace# df
Filesystem          1K-blocks     Used     Available Use% Mounted on
overlay              20134592 11121996       8996212  56% /
tmpfs                   65536        0         65536   0% /dev
tmpfs                 2013020        0       2013020   0% /sys/fs/cgroup
shm                     65536        0         65536   0% /dev/shm
/dev/vda1            20134592 11121996       8996212  56% /shared
fanzago:fanzago 1099511627776        0 1099511627776   0% /s3/fanzago
scratch:scratch 1099511627776        0 1099511627776   0% /s3/scratch
root@518f931c05e4:/jupyter-workspace#
```

Shared area among users of authorized group

Inside Lab container

Best practice: don’t execute notebooks in the cloud-storage area but run them in the “local” one. Use the object storage only as archive.

Containers in the VM

```
>> ssh  -i <key> <ssh_account>@node_ip
```

Container Lab

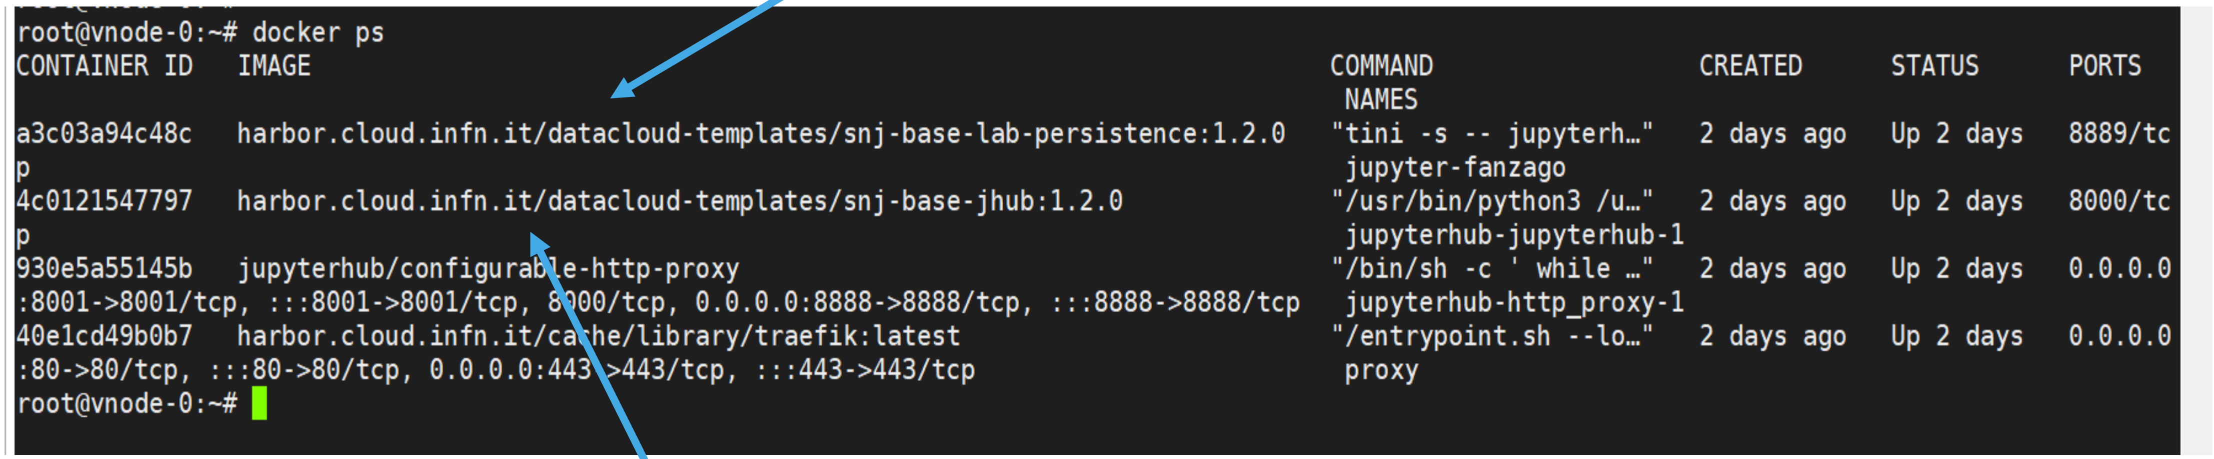

```
root@vnode-0:~# docker ps
CONTAINER ID   IMAGE                                                                     COMMAND                  CREATED        STATUS       PORTS
                                                                                          NAMES
a3c03a94c48c   harbor.cloud.infn.it/datacloud-templates/snj-base-lab-persistence:1.2.0   "tini -s -- jupyterh…"   2 days ago   Up 2 days   8889/tc
p                                                                                          jupyter-fanzago
4c0121547797   harbor.cloud.infn.it/datacloud-templates/snj-base-jhub:1.2.0              "/usr/bin/python3 /u…"   2 days ago   Up 2 days   8000/tc
p                                                                                          jupyterhub-jupyterhub-1
930e5a55145b   jupyterhub/configurable-http-proxy                                        "/bin/sh -c ' while …"   2 days ago   Up 2 days   0.0.0.0
:8001->8001/tcp, :::8001->8001/tcp, 8000/tcp, 0.0.0.0:8888->8888/tcp, :::8888->8888/tcp   jupyterhub-http_proxy-1
40e1cd49b0b7   harbor.cloud.infn.it/cache/library/traefik:latest                         "/entrypoint.sh --lo…"   2 days ago   Up 2 days   0.0.0.0
:80->80/tcp, :::80->80/tcp, 0.0.0.0:443->443/tcp, :::443->443/tcp                          proxy
root@vnode-0:~#
```

Container Hub

Healthchecks as a service

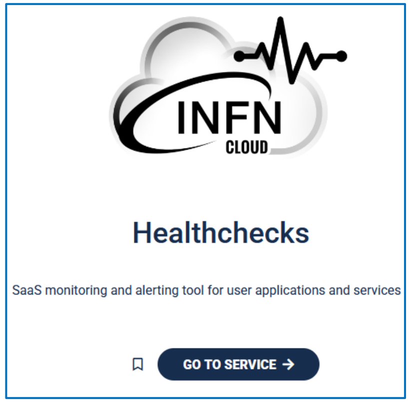

Healthchecks

SaaS monitoring and alerting tool for user applications and services

- https://healthchecks.cloud.infn.it
- https://guides.cloud.infn.it/docs/users-guides/en/latest/users_guides/centralised/healthchecks.html
- Healthchecks is an open-source powerful monitoring and alerting tool for applications and services.
  - It allows to easily monitor the availability of a system or the execution of a program by periodically sending HTTP “keep-alive” requests to custom endpoints.
- It is installed on the backbone (CNAF and Bari), high availability
- After login through IAM, user is redirected to the service webpage where projects (check) can be created.

Healthchecks interface

- The page displays your projects and allows to create new ones.

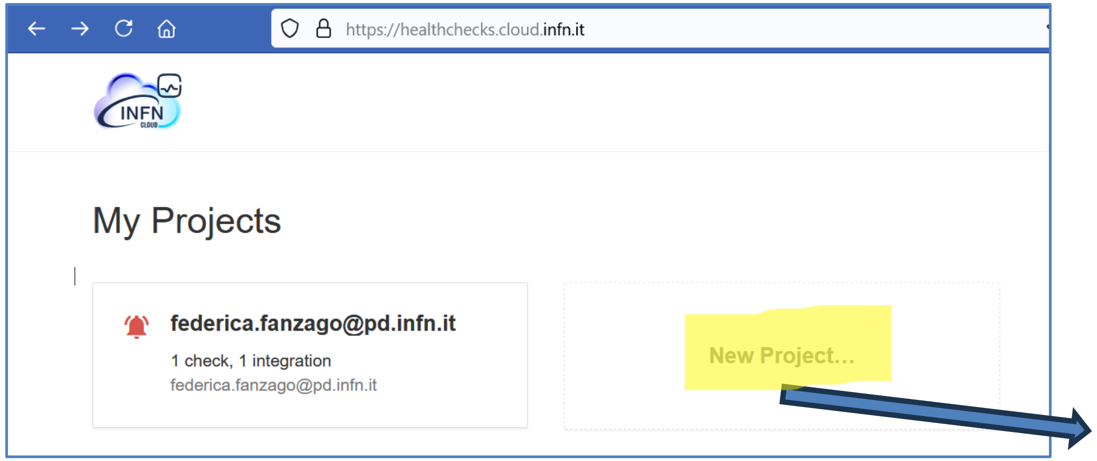

https://healthchecks.cloud.infn.it

My Projects

federica.fanzago@pd.infn.it
1 check, 1 integration
federica.fanzago@pd.infn.it

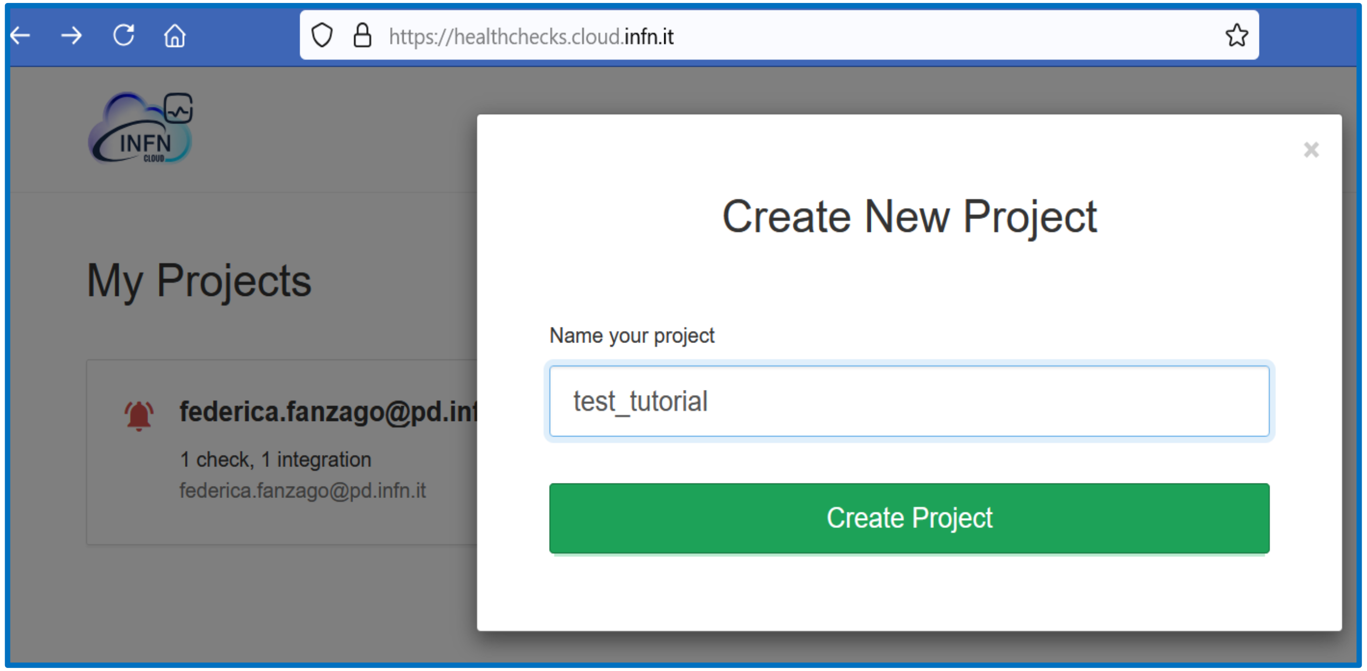

https://healthchecks.cloud.infn.it

My Projects

federica.fanzago@pd.inf
1 check, 1 integration
federica.fanzago@pd.infn.it

Create New Project

Name your project: test_tutorial

Authentication and authorization via IAM

Healthchecks interface

- Fill the form with your parameters

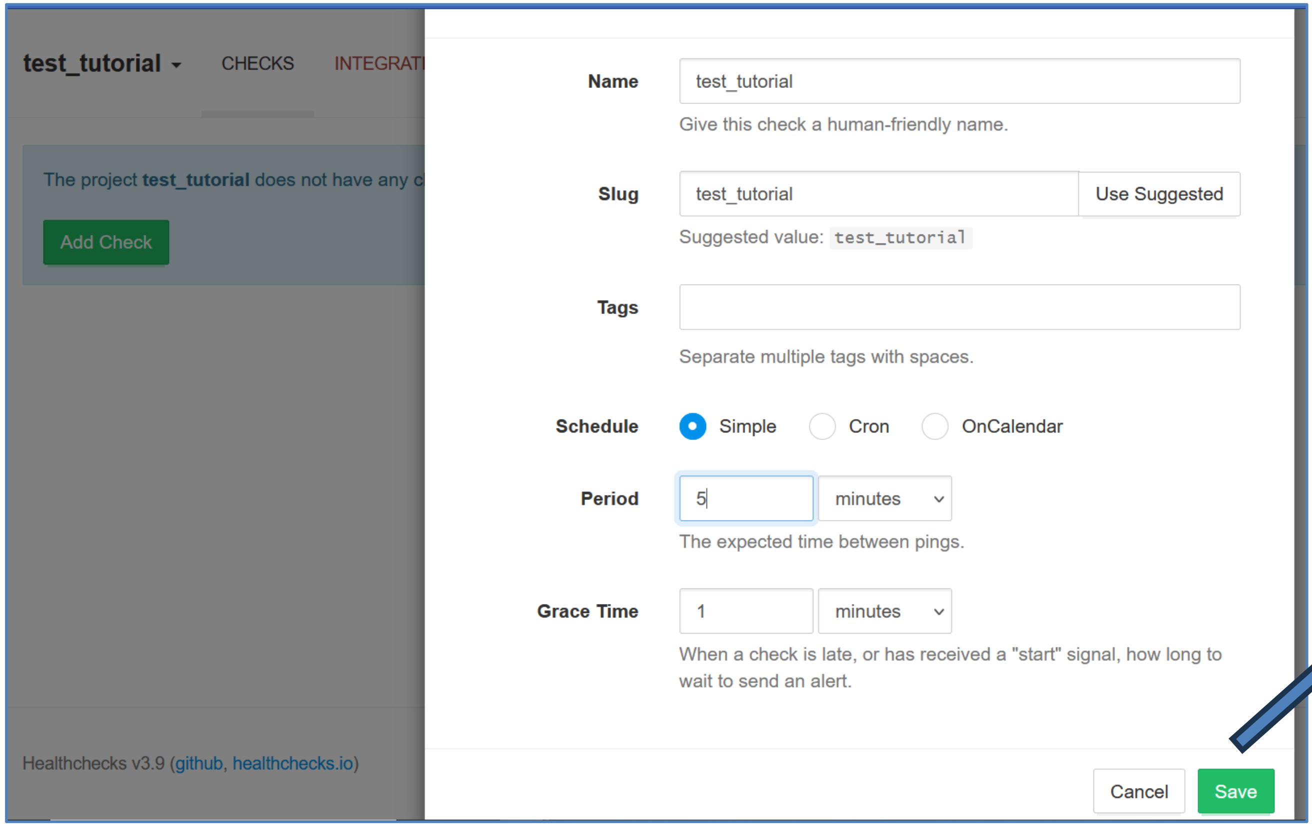

test_tutorial ▾ CHECKS INTEGRATI

The project **test_tutorial** does not have any c

Add Check

Healthchecks v3.9 (github, healthchecks.io)

Name	test_tutorial Give this check a human-friendly name.
Slug	test_tutorial [Use Suggested] Suggested value: `test_tutorial`
Tags	 Separate multiple tags with spaces.
Schedule	● Simple ○ Cron ○ OnCalendar
Period	5 minutes The expected time between pings.
Grace Time	1 minutes When a check is late, or has received a "start" signal, how long to wait to send an alert.

Cancel Save

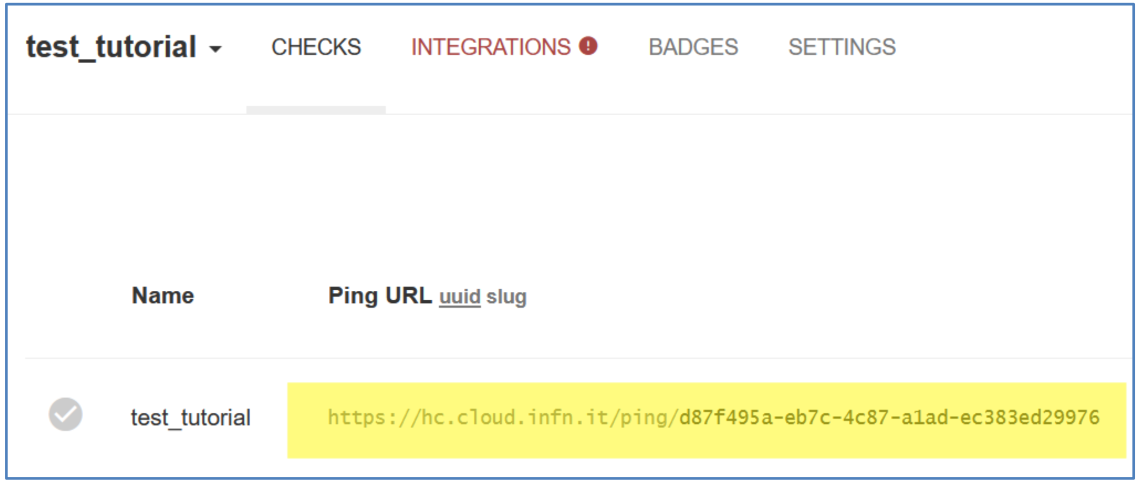

test_tutorial ▾ CHECKS INTEGRATIONS ❗ BADGES SETTINGS

Name	Ping URL uuid slug
test_tutorial	https://hc.cloud.infn.it/ping/d87f495a-eb7c-4c87-a1ad-ec383ed29976

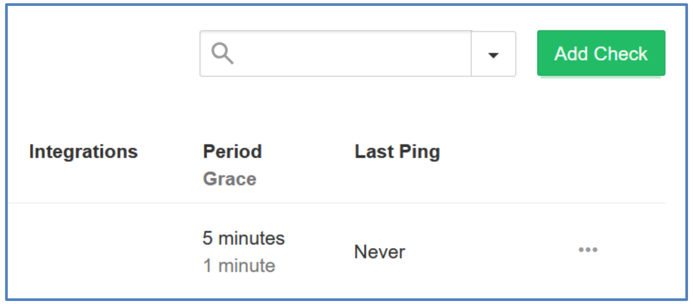

Add Check

Integrations	Period Grace	Last Ping
	5 minutes 1 minute	Never

Healthchecks interface

- Your new project

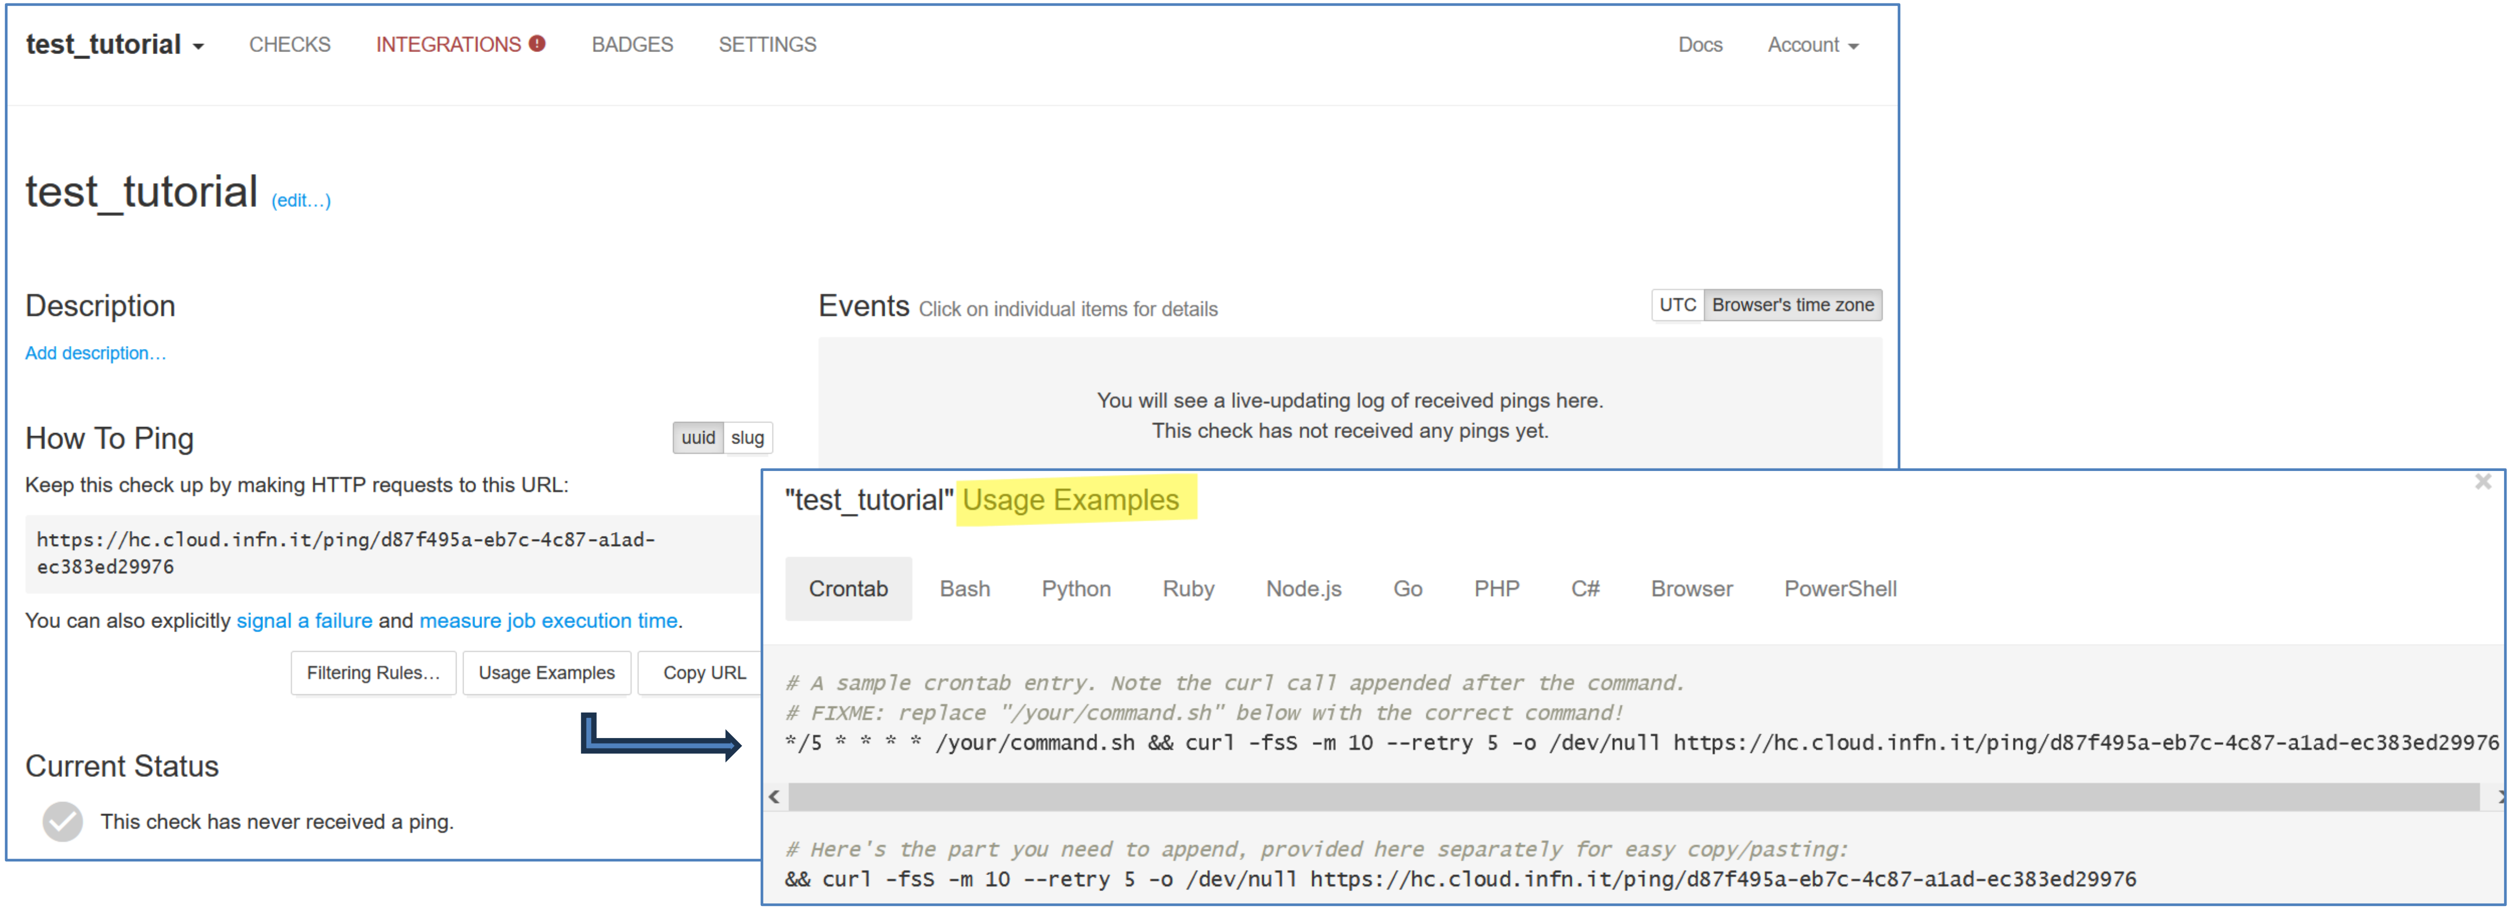

test_tutorial (edit…)

Description

Add description…

How To Ping

Keep this check up by making HTTP requests to this URL:

```
https://hc.cloud.infn.it/ping/d87f495a-eb7c-4c87-a1ad-ec383ed29976
```

You can also explicitly signal a failure and measure job execution time.

Current Status

This check has never received a ping.

Events Click on individual items for details

You will see a live-updating log of received pings here.
This check has not received any pings yet.

"test_tutorial" Usage Examples

```
# A sample crontab entry. Note the curl call appended after the command.
# FIXME: replace "/your/command.sh" below with the correct command!
*/5 * * * * /your/command.sh && curl -fsS -m 10 --retry 5 -o /dev/null https://hc.cloud.infn.it/ping/d87f495a-eb7c-4c87-a1ad-ec383ed29976
```

```
# Here's the part you need to append, provided here separately for easy copy/pasting:
&& curl -fsS -m 10 --retry 5 -o /dev/null https://hc.cloud.infn.it/ping/d87f495a-eb7c-4c87-a1ad-ec383ed29976
```

Healthchecks interface

- Check details and results

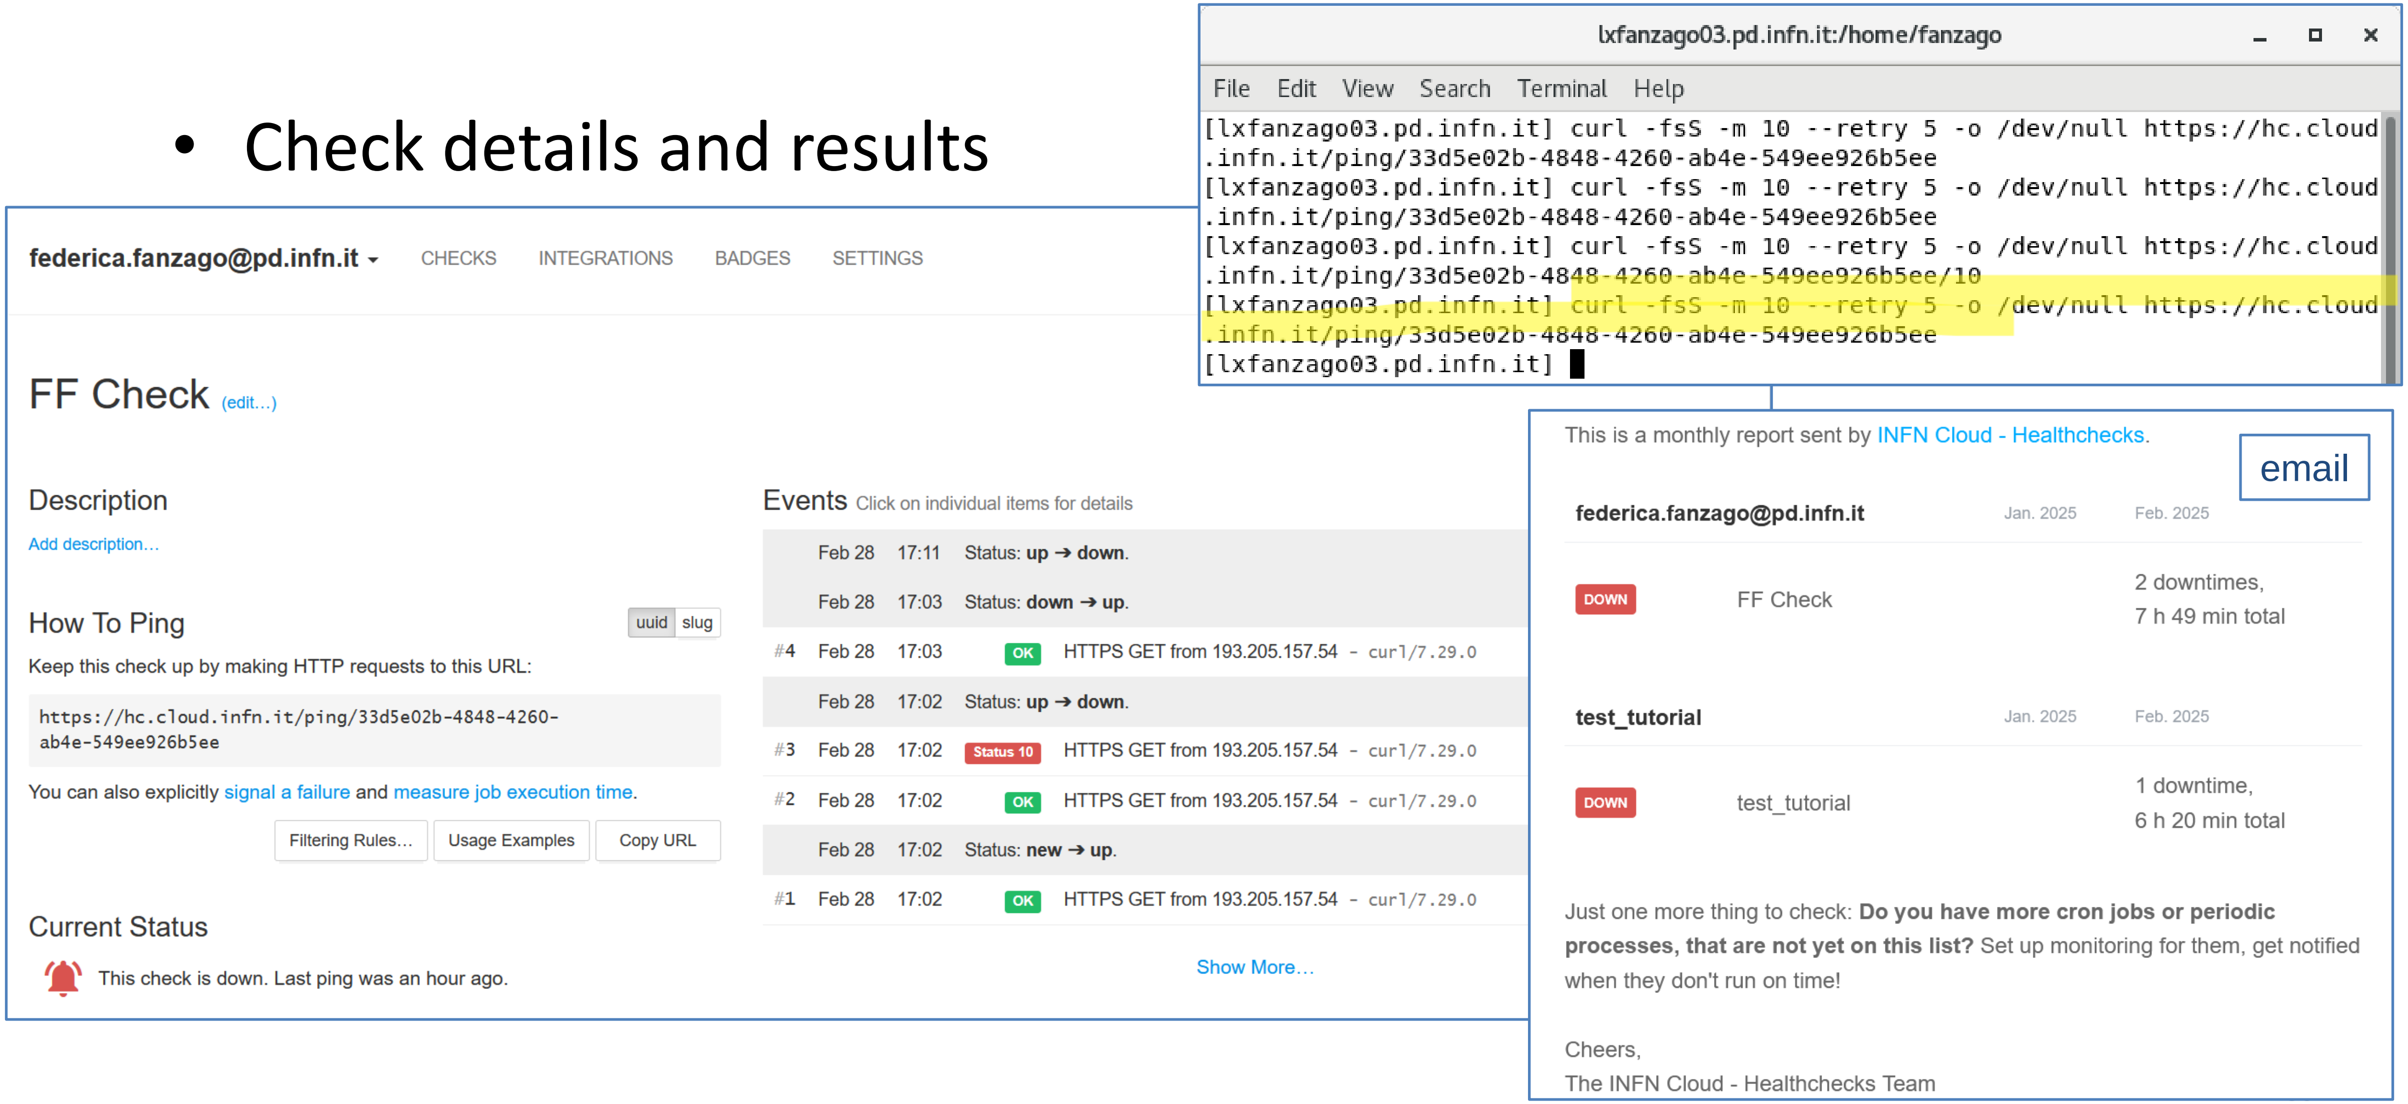

lxfanzago03.pd.infn.it:/home/fanzago

```
[lxfanzago03.pd.infn.it] curl -fsS -m 10 --retry 5 -o /dev/null https://hc.cloud
.infn.it/ping/33d5e02b-4848-4260-ab4e-549ee926b5ee
[lxfanzago03.pd.infn.it] curl -fsS -m 10 --retry 5 -o /dev/null https://hc.cloud
.infn.it/ping/33d5e02b-4848-4260-ab4e-549ee926b5ee
[lxfanzago03.pd.infn.it] curl -fsS -m 10 --retry 5 -o /dev/null https://hc.cloud
.infn.it/ping/33d5e02b-4848-4260-ab4e-549ee926b5ee/10
[lxfanzago03.pd.infn.it] curl -fsS -m 10 --retry 5 -o /dev/null https://hc.cloud
.infn.it/ping/33d5e02b-4848-4260-ab4e-549ee926b5ee
[lxfanzago03.pd.infn.it] █
```

FF Check (edit…)

Description

Add description…

How To Ping

Keep this check up by making HTTP requests to this URL:

```
https://hc.cloud.infn.it/ping/33d5e02b-4848-4260-
ab4e-549ee926b5ee
```

You can also explicitly signal a failure and measure job execution time.

Current Status

This check is down. Last ping was an hour ago.

Events Click on individual items for details

	Feb 28	17:11	Status: **up** ➔ **down**.	
	Feb 28	17:03	Status: **down** ➔ **up**.	
#4	Feb 28	17:03	OK	HTTPS GET from 193.205.157.54 - curl/7.29.0
	Feb 28	17:02	Status: **up** ➔ **down**.	
#3	Feb 28	17:02	Status 10	HTTPS GET from 193.205.157.54 - curl/7.29.0
#2	Feb 28	17:02	OK	HTTPS GET from 193.205.157.54 - curl/7.29.0
	Feb 28	17:02	Status: **new** ➔ **up**.	
#1	Feb 28	17:02	OK	HTTPS GET from 193.205.157.54 - curl/7.29.0

Show More…

email

This is a monthly report sent by INFN Cloud - Healthchecks.

federica.fanzago@pd.infn.it		Jan. 2025	Feb. 2025
DOWN	FF Check		2 downtimes, 7 h 49 min total

test_tutorial		Jan. 2025	Feb. 2025
DOWN	test_tutorial		1 downtime, 6 h 20 min total

Just one more thing to check: **Do you have more cron jobs or periodic processes, that are not yet on this list?** Set up monitoring for them, get notified when they don't run on time!

Cheers,
The INFN Cloud - Healthchecks Team

References

- INFN Cloud homepage: https://www.cloud.infn.it/
- User guides: https://guides.cloud.infn.it/docs/users-guides/en/latest/
- Service catalogue: https://www.cloud.infn.it/service-catalogue/
- INFN Cloud dashboard: https://my.cloud.infn.it/
- INFN Cloud support: https://servicedesk.cloud.infn/it

BACKUP

Harbor and cvmfs

- CernVM-File system (cvmfs) is aread-only file system designed to deliver scientific software onto virtual machines and physical worker nodes in a fast, scalable, and reliable way (LHC experiments). It can be used to distribute container images pushed in Harbor (cvmfs webhook) under /cmvfs/unpacked.infn.it/harbor.cloud.infn.it/unpacked area

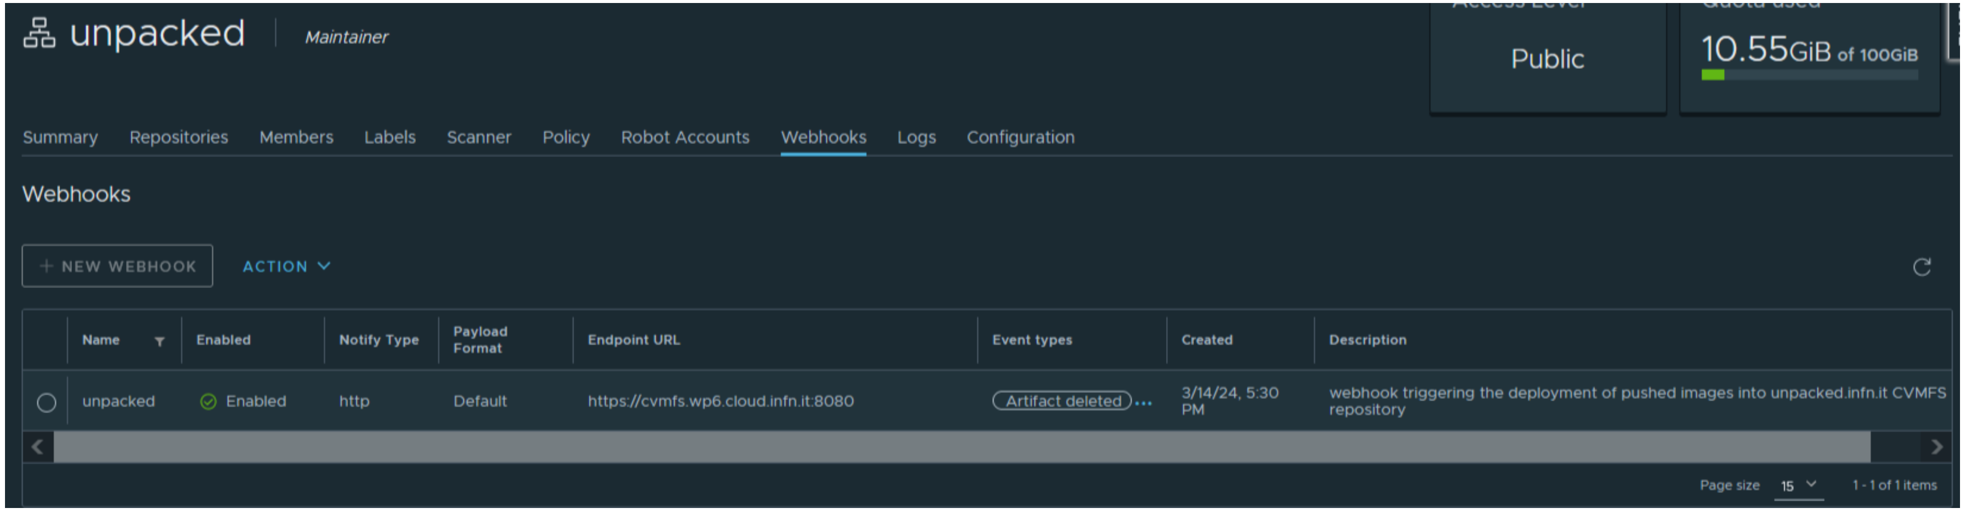

	Name	Enabled	Notify Type	Payload Format	Endpoint URL	Event types	Created	Description
○	unpacked	✓ Enabled	http	Default	https://cvmfs.wp6.cloud.infn.it:8080	Artifact deleted ...	3/14/24, 5:30 PM	webhook triggering the deployment of pushed images into unpacked.infn.it CVMFS repository

Harbor and apptainer

- Images are readily available to run with apptainer (formerly singularity)
  - Singularity containers are preferred when running applications in HPC systems. Singularity containers can be run without sudo.
- Push in “cvmfs harbor” area

```
## in the terminal where docker is installed and running
$ docker login harbor.cloud.infn.it
## tag the image
$ docker tag my-image:1.0  harbor.cloud.infn.it/unpacked/my-image:1.0
## push the image
$ docker push  harbor.cloud.infn.it/unpacked/my-image:1.0
## In the cvmfs client where apptainer is installed
$ apptainer exec '/cvmfs/unpacked.infn.it/harbor.cloud.infn.it/unpacked/my-image:1.0' /bin/sh
```