



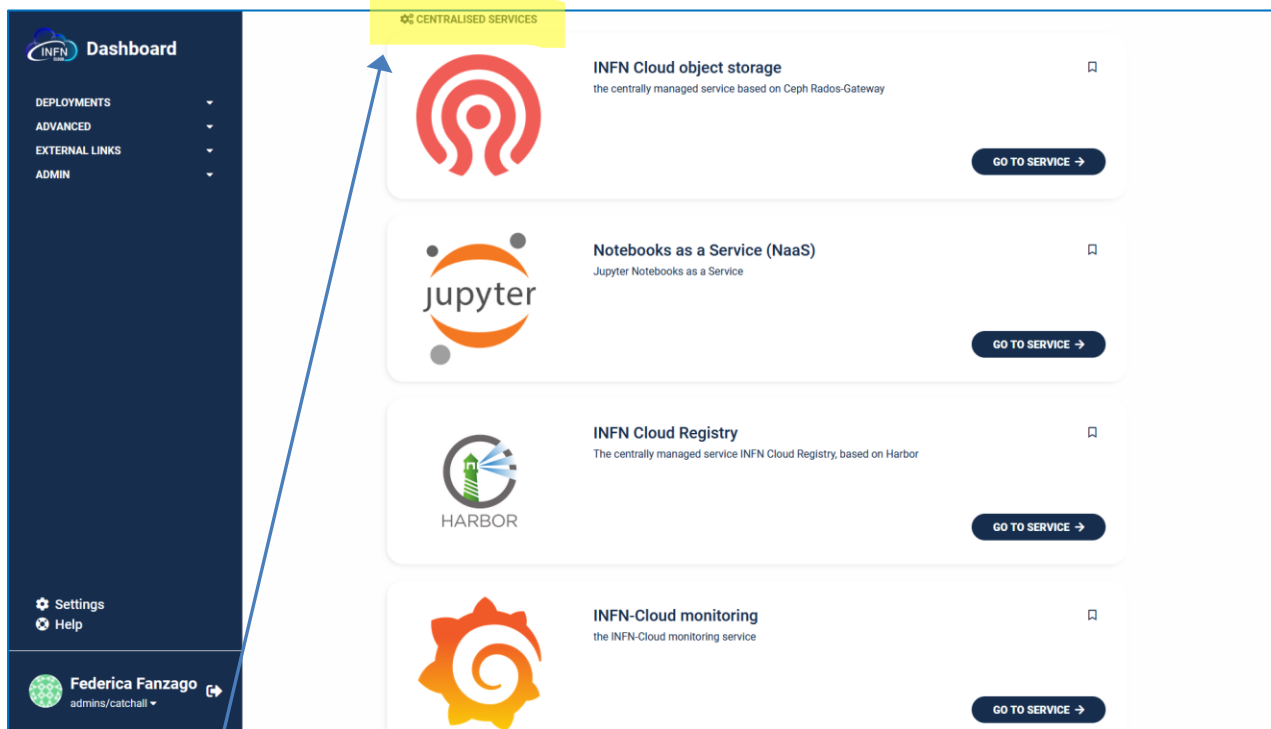
## INFN Cloud centralized services

Corso di formazione per neoassunti nelle attività di computing – INFN DATACloud  
17 Luglio 2024

# Outline

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

# Available services



**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 →

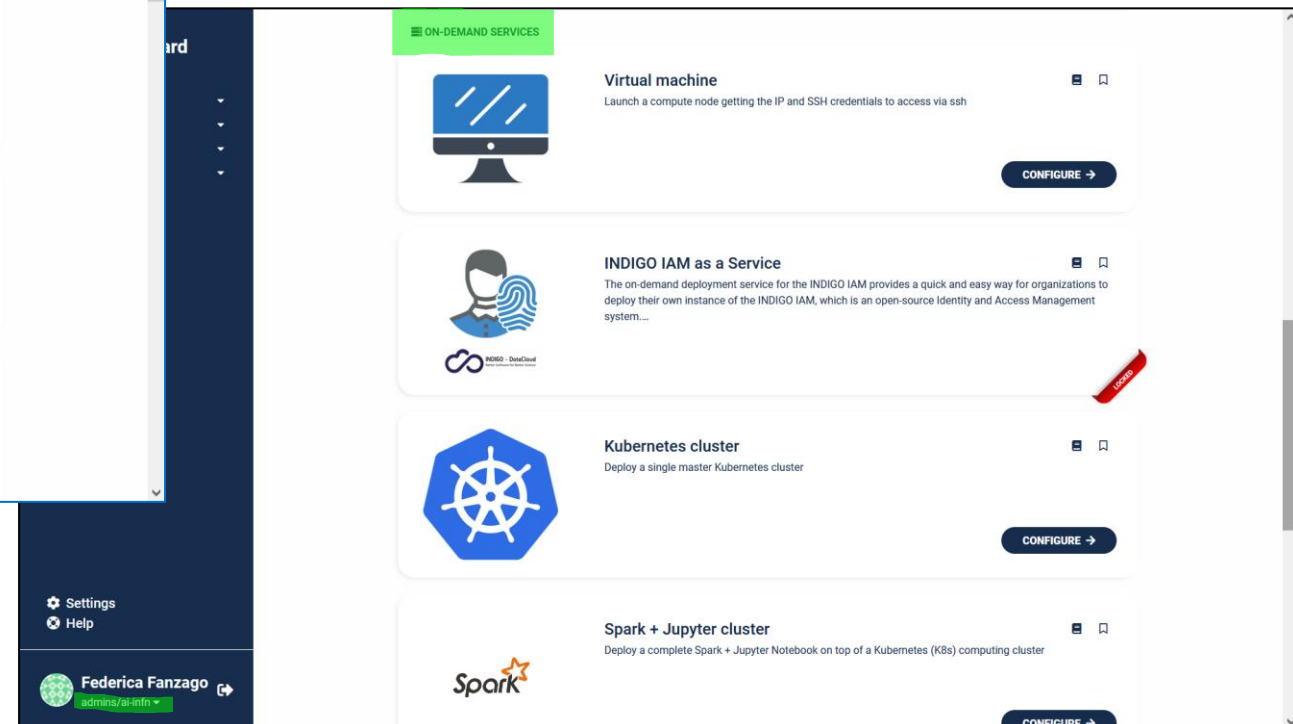
Dashboard  
DEPLOYMENTS  
ADVANCED  
EXTERNAL LINKS  
ADMIN

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...
- 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 →

Dashboard

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



- <https://s3webui.cloud.infn.it>
- [https://guides.cloud.infn.it/docs/users-guides/en/latest/users\\_guides/centralised/objectstorage.html](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



The image shows a sequence of three browser windows illustrating the login process:

- Window 1 (Left):** `https://s3webui.cloud.infn.it/login`. The main page features the INFN logo and a "Welcome" message. Below it are four login options: "Access Key Id", "Secret Access Key", "Login with local credentials", and "Login with INDIGO IAM". An arrow points from the "Login with INDIGO IAM" option to the next window.
- Window 2 (Middle):** `https://iam.cloud.infn.it/login`. The page displays the "INFN CLOUD" logo and the text "Welcome to infn-cloud". It includes a "Sign in with" section with the "INFN CCR - AAI" logo, a "Not a member?" link, and a green "Apply for an account" button.
- Window 3 (Right):** `https://idp.infn.it/module.php/wawa/login.php/plain?StateId=_22e311b5df78ba70ffa0596`. This is the OpenID provider login page, featuring the "INFN CCR - AAI" logo, language options "IT | EN", and input fields for "Username or e-mail" and "Password". Below the fields are buttons for "LOGIN", "REGISTER", "Change or Reset Password - Retrieve Username", "X.509 CERTIFICATE", and "KERBEROS - GSSAPI".

A separate box at the top right contains the URL `https://s3webui.cloud.infn.it`.

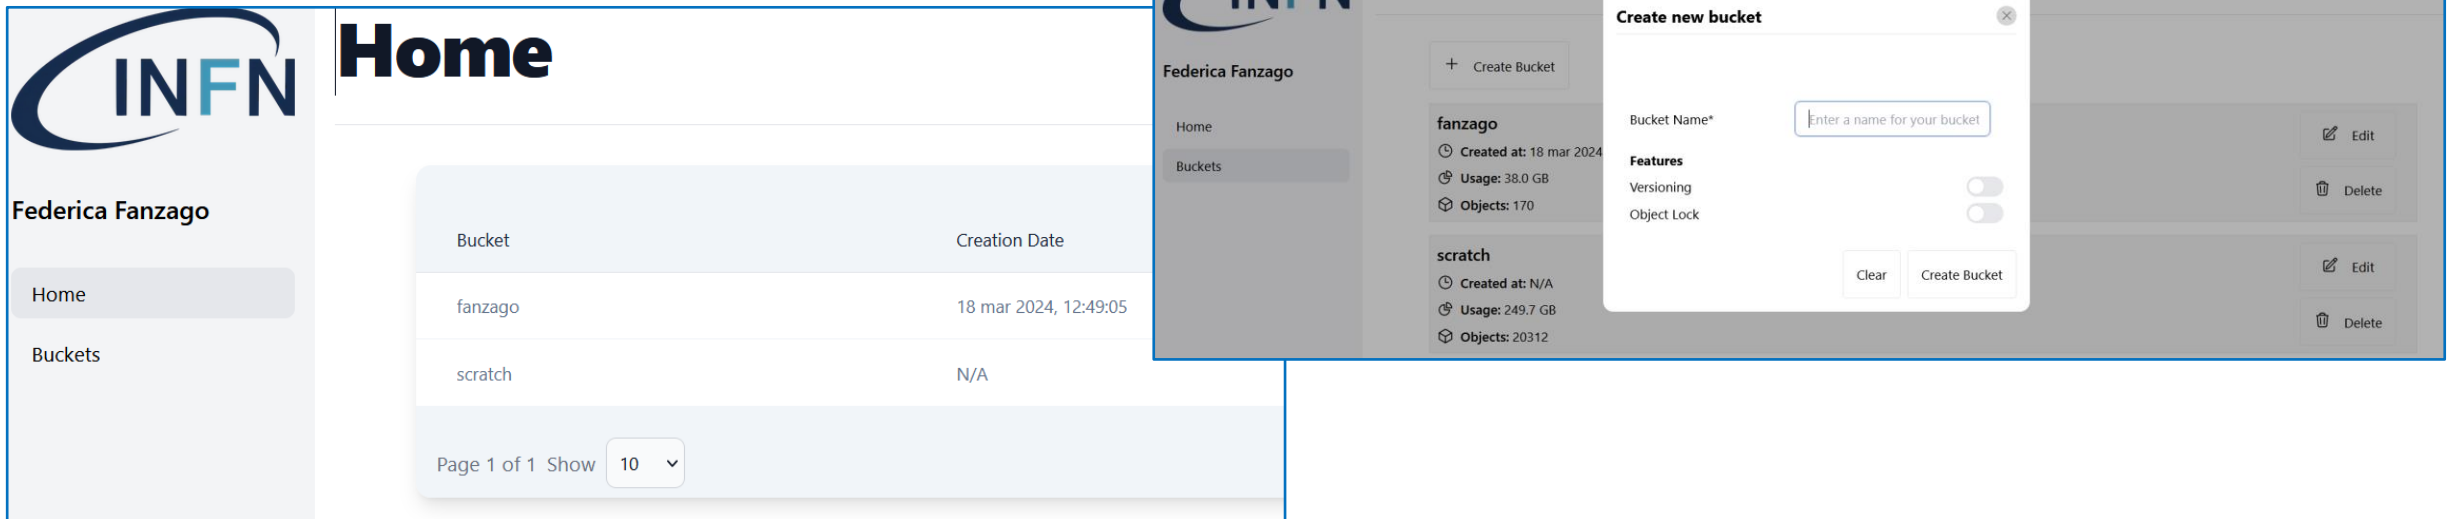
# Persistent data with object storage

- 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 be created by user accessing the Ceph RGW webui otherwise it is automatically created the first time user instantiates a Jupyter notebook via INFN Cloud.



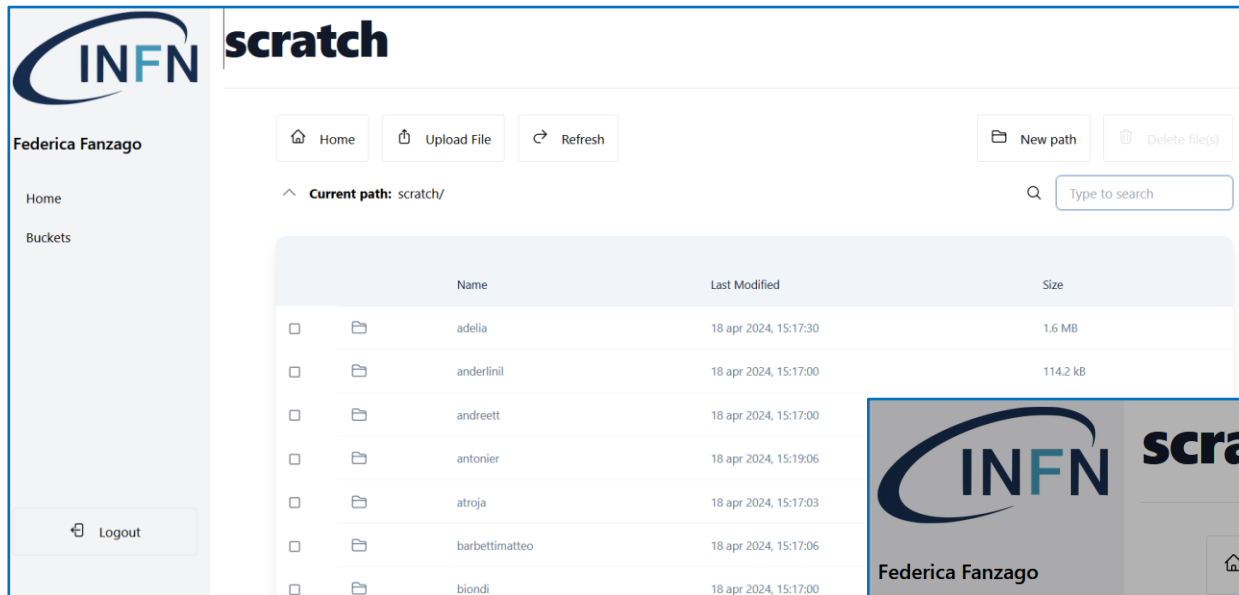
The screenshot displays the INFN Cloud web interface. On the left, a sidebar shows the user 'Federica Fanzago' and navigation options for 'Home' and 'Buckets'. The main content area is titled 'Home' and features a table of buckets:

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

Below the table, it shows 'Page 1 of 1' and a 'Show 10' dropdown menu. On the right, the 'Buckets' page is visible, showing details for 'fanzago' (Created at: 18 mar 2024, Usage: 38.0 GB, Objects: 170) and 'scratch' (Created at: N/A, Usage: 249.7 GB, Objects: 20312). A 'Create new bucket' dialog is open, prompting for a 'Bucket Name\*' and offering options for 'Versioning' and 'Object Lock'.

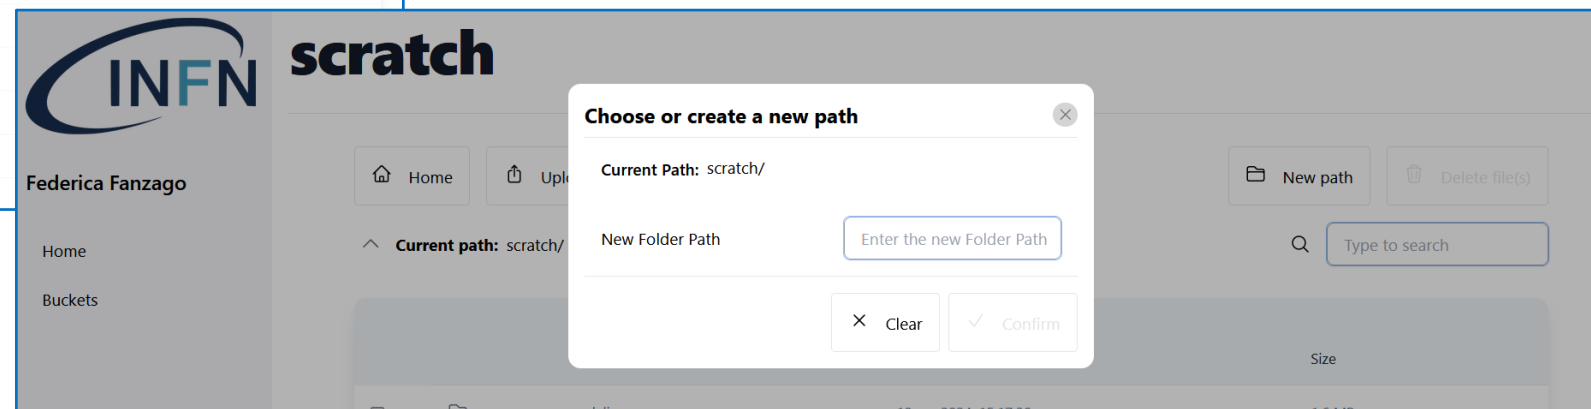
# Personal area under scratch

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



The screenshot shows the INFN scratch web interface. The user is logged in as Federica Fanzago. The current path is /scratch/. The interface includes a sidebar with 'Home' and 'Buckets' options, and a main area with a table of folders. The table has columns for Name, Last Modified, and Size.

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	



The screenshot shows the INFN scratch web interface with a modal dialog titled 'Choose or create a new path'. The dialog is open over the 'New path' button. The current path is /scratch/. The dialog has a text input field for 'New Folder Path' and two buttons: 'Clear' and 'Confirm'.

Choose or create a new path

Current Path: /scratch/

New Folder Path:

Clear Confirm

# 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 work in progress...
- Via Rclone the object storage is accesible as local file system
  - rclone ls profile:/<bucket>

```
s3client = boto3.client('s3',
    aws_access_key_id = response['Credentials']['AccessKeyId'],
    aws_secret_access_key = response['Credentials']['SecretAccessKey'],
    aws_session_token = response['Credentials']['SessionToken'],
    endpoint_url="https://rgw.cloud.infn.it",
    region_name='default')

# Show files in a bucket

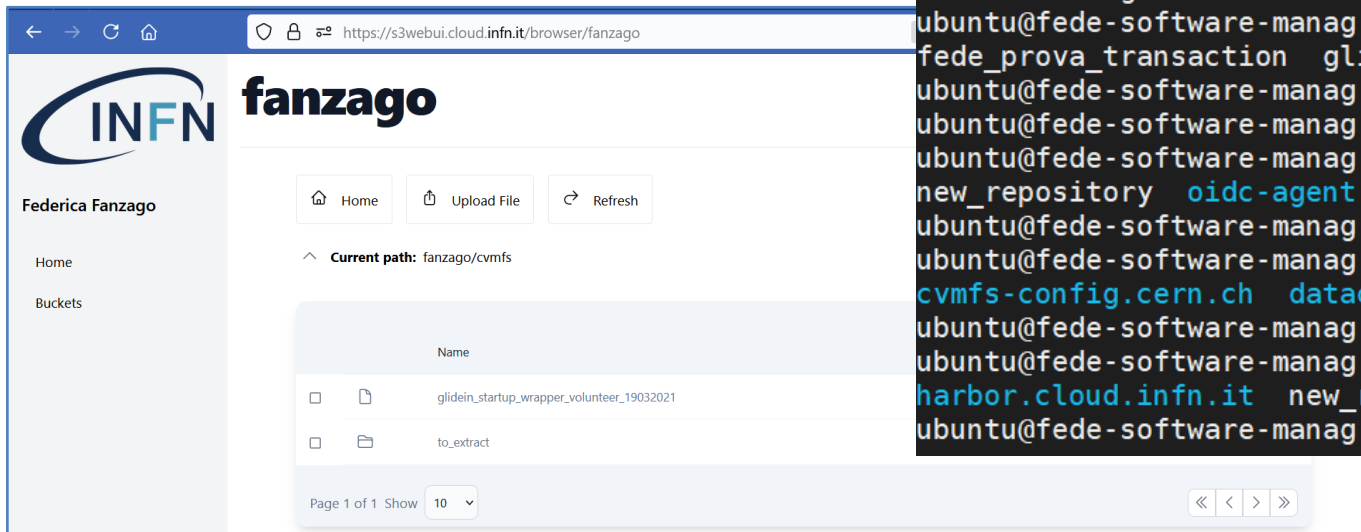
for key in s3client.list_objects(Bucket='scratch')['Contents']:
    print(key['Key'])
```

# Object storage also for



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

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

Harbor container registry




The screenshot shows a web browser window with the URL `https://s3webui.cloud.infn.it/browser/fanzago`. The page header includes the INFN logo and the name "fanzago". On the left, there is a sidebar with the name "Federica Fanzago" and navigation links for "Home" and "Buckets". The main content area shows a file browser interface with buttons for "Home", "Upload File", and "Refresh". Below these buttons, the "Current path" is indicated as "fanzago/cvmfs". A table lists files and folders:

	Name
<input type="checkbox"/>	 glidein_startup_wrapper_volunteer_19032021
<input type="checkbox"/>	 to_extract

At the bottom of the table, it says "Page 1 of 1 Show 10" with a dropdown menu.

```
/cvmfs/fanzago.infn.it
ubuntu@federe-software-manag:/cvmfs/fanzago.infn.it$ ls
federe_prova_transaction glidein_startup_wrapper_volunteer_19032021 mysoftware new_
ubuntu@federe-software-manag:/cvmfs/fanzago.infn.it$ cd ..
ubuntu@federe-software-manag:/cvmfs$ cd datacloud.infn.it
ubuntu@federe-software-manag:/cvmfs/datacloud.infn.it$ ls
new_repository oidc-agent prova-new-cvmfs-keysetup.txt repo sts-wire test
ubuntu@federe-software-manag:/cvmfs/datacloud.infn.it$ cd ..
ubuntu@federe-software-manag:/cvmfs$ ls
cvmfs-config.cern.ch datacloud.infn.it fanzago.infn.it
ubuntu@federe-software-manag:/cvmfs$ cd unpacked.infn.it
ubuntu@federe-software-manag:/cvmfs/unpacked.infn.it$ ls
harbor.cloud.infn.it new_repository
ubuntu@federe-software-manag:/cvmfs/unpacked.infn.it$
```

# INFN Cloud Registry



- <https://harbor.cloud.infn.it>
- [https://guides.cloud.infn.it/docs/users-guides/en/latest/users\\_guides/centralised/harbor.html](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

[LOGIN VIA OIDC PROVIDER](#)

OR

[LOGIN VIA LOCAL DB](#)

[More info...](#)

Search INFN Cloud... English Default **fanzago**

### Projects

Projects: Public 8    Repositories: Public 61

ACTION... All Projects

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

Page size 15 1 - 11 of 11 items

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

Limited guest, Guest, Developer, Master, Project Admin



Projects > beta-testers

Access Level: Private | Quota used: 14.08GiB of 50GiB

Summary | **Repositories** | Members | Labels | Scanner | Policy | Robot Accounts | Webhooks | Logs | Configuration

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

DELETE | PUSH COMMAND

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]
```

beta-testers > fanzago-alpine

Info | **Artifacts**

SCAN | STOP SCAN | ACTIONS

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

Manage Columns | Page size 15 | 1 - 2 of 2 items

sha256:0dff30ce

Tags

+ ADD TAG | REMOVE TAG

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

COPY PULL COMMAND

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



# 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”

## ## login

```
$ docker login harbor.cloud.infn.it
```

## ## tag the image

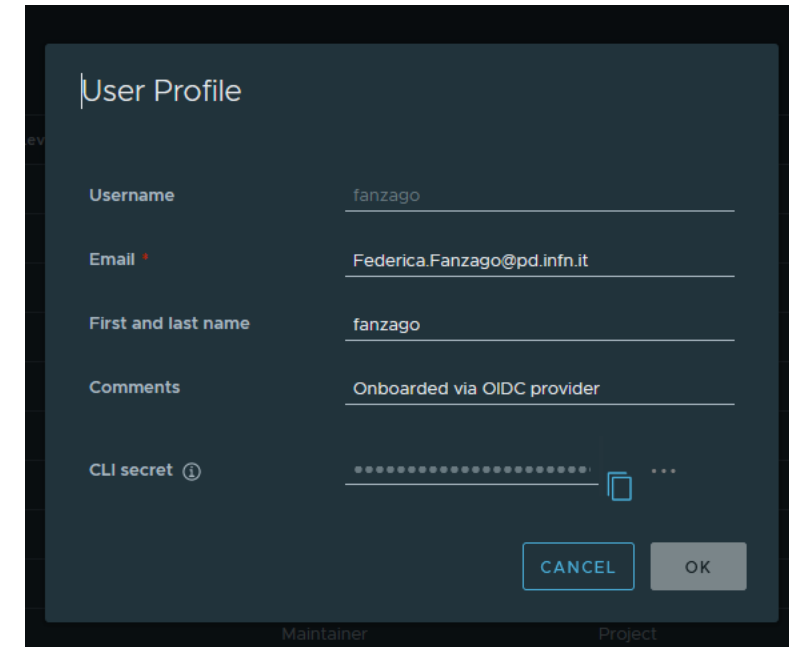
```
$ docker tag SOURCE_IMAGE[:TAG] harbor.cloud.infn.it/project/REPOSITORY[:TAG]
```

## ## push the image

```
$ docker push harbor.cloud.infn.it/project/REPOSITORY[:TAG]
```

## ## pull the image

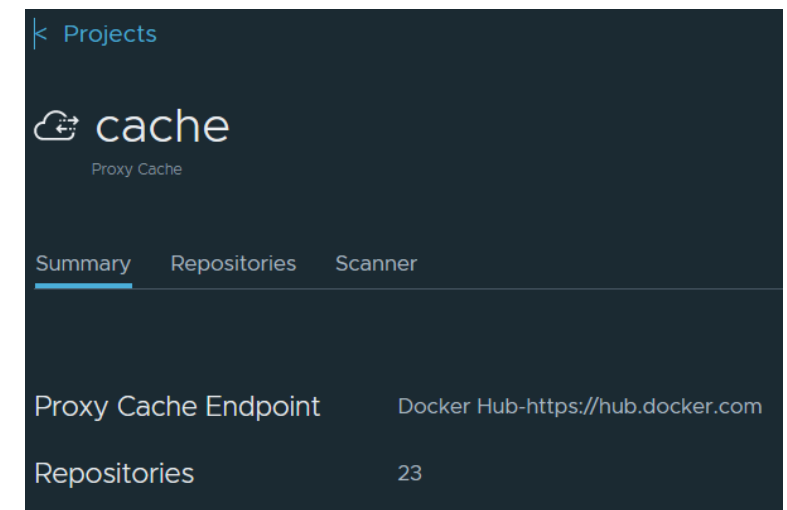
```
$ docker pull harbor.cloud.infn.it/project/REPOSITORY[:TAG]
```



# Harbor: other details

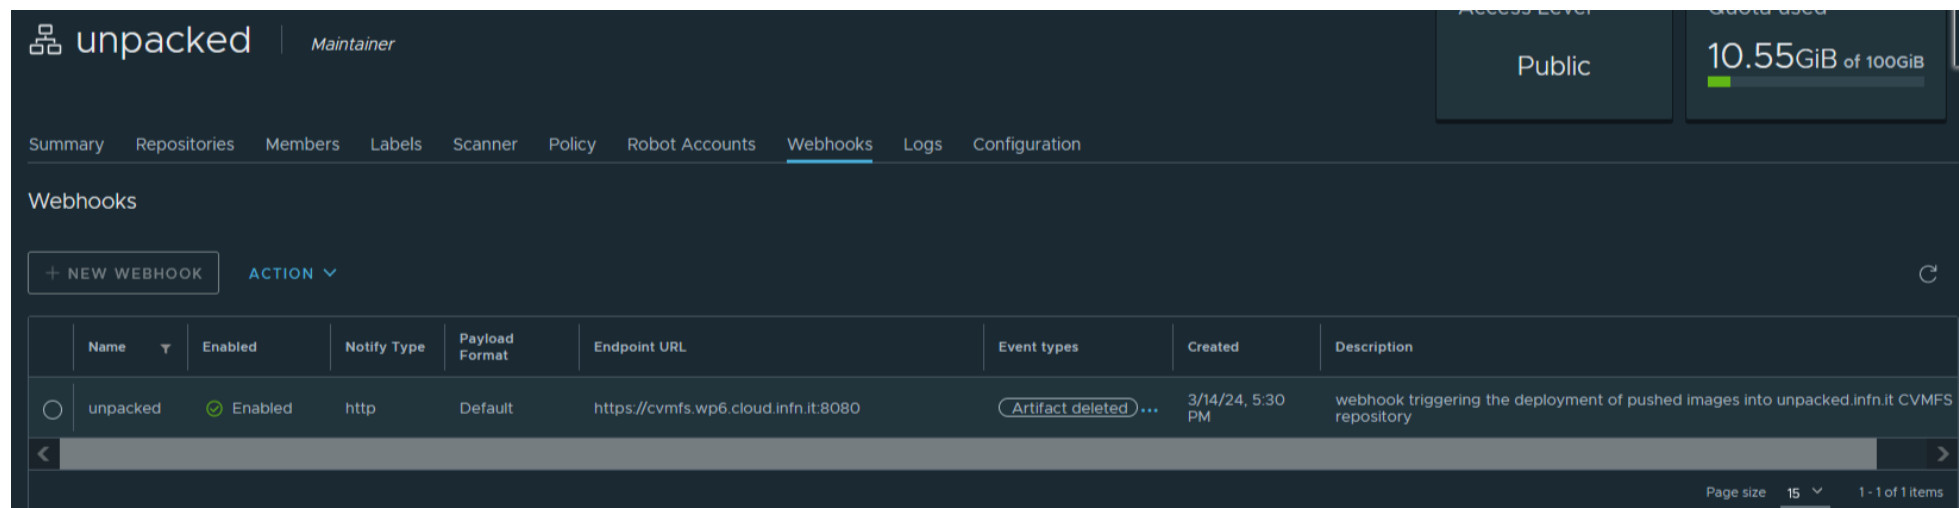
- Images are stored in the object storage of INFN Cloud, a quota can be guarantee for projects. The bucket is unique.
- Proxy cache configured:
  - when a pull request comes to a proxy cache project, if the image is not cached, Harbor pulls the image from the target registry and serves the pull command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

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



# Harbor and cvmfs

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



The screenshot shows the Harbor UI for the 'unpacked' repository. The 'Webhooks' tab is active, displaying a table with one configured webhook. The table has columns for Name, Enabled, Notify Type, Payload Format, Endpoint URL, Event types, Created, and Description. The webhook is named 'unpacked', is enabled, uses 'http' as the notify type, 'Default' as the payload format, and has an endpoint URL of 'https://cvmfs.wp6.cloud.infn.it:8080'. It is configured to trigger on 'Artifact deleted' events and was created on 3/14/24 at 5:30 PM. The description is 'webhook triggering the deployment of pushed images into unpacked.infn.it CVMFS repository'.

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
```

# Notebook as a service (NaaS)



- <https://hub.cloud.infn.it>
- [https://guides.cloud.infn.it/docs/users-guides/en/latest/users\\_guides/centralised/naas.html](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:1.2.0`
- 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

# 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



**Elasticsearch and Kibana**


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

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




**Analytics**

A collection of ad-hoc solutions for analytic purpose




**Spark + Jupyter cluster**

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




**Scientific Community Customizations**

Customized environments




**Machine Learning**

List of ready-to-use Machine Learning services




**Jupyter with persistence for Notebooks**

Run Jupyter on a single VM enabling Notebooks persistence




**Working Station for CYGNO experiment**

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



**Data Services**

Data management and storage services



**Centralized Notebooks as a Service**

Use the INFN Cloud centrally managed Jupyter Notebooks as a Service Solution  
Endpoint: <https://hub.cloud.infn.it>



**Machine Learning**

List of ready-to-use Machine Learning services

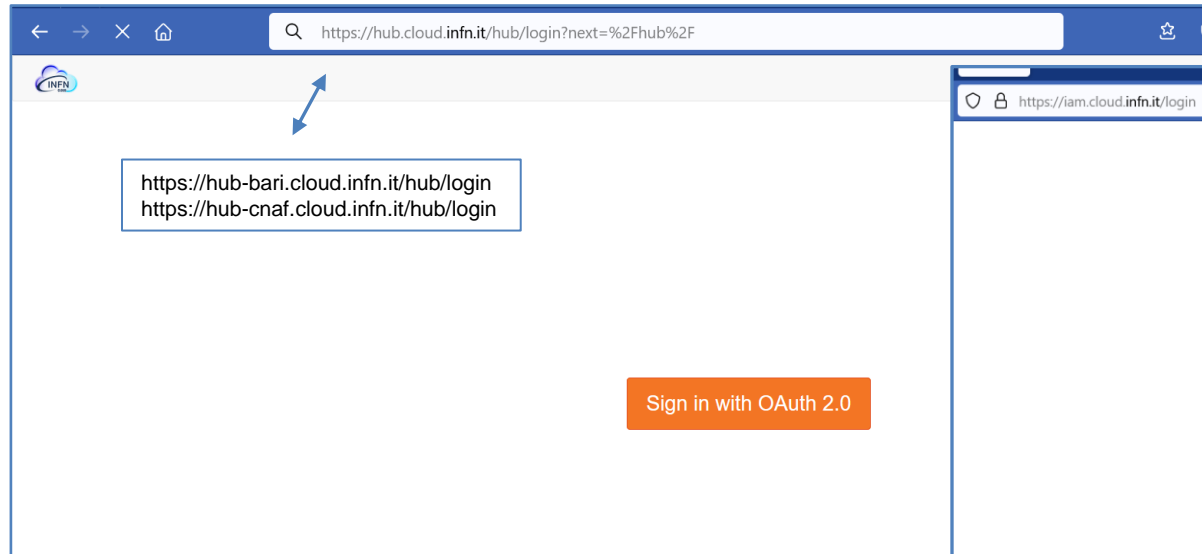


**Working Station for Machine Learning INFN (ML-INFN)**

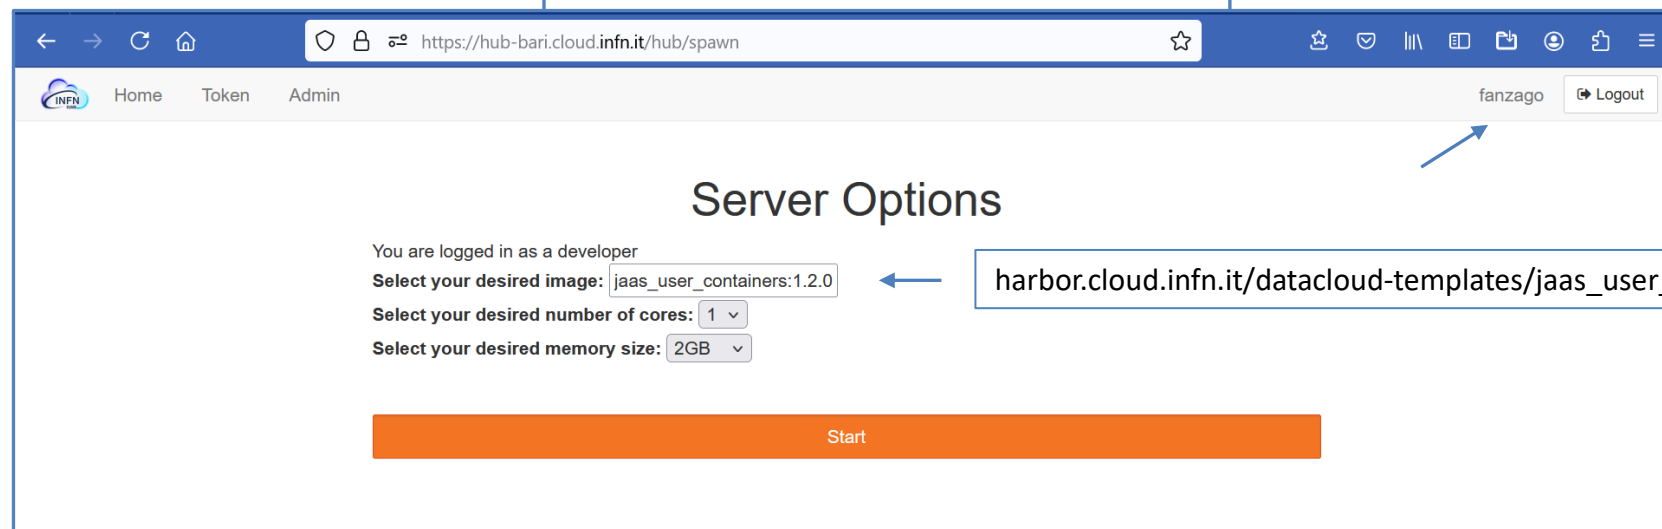
Run a single VM with all the ML-INFN environment exposing both ssh access and Jupyter



# NaaS Hub Interface

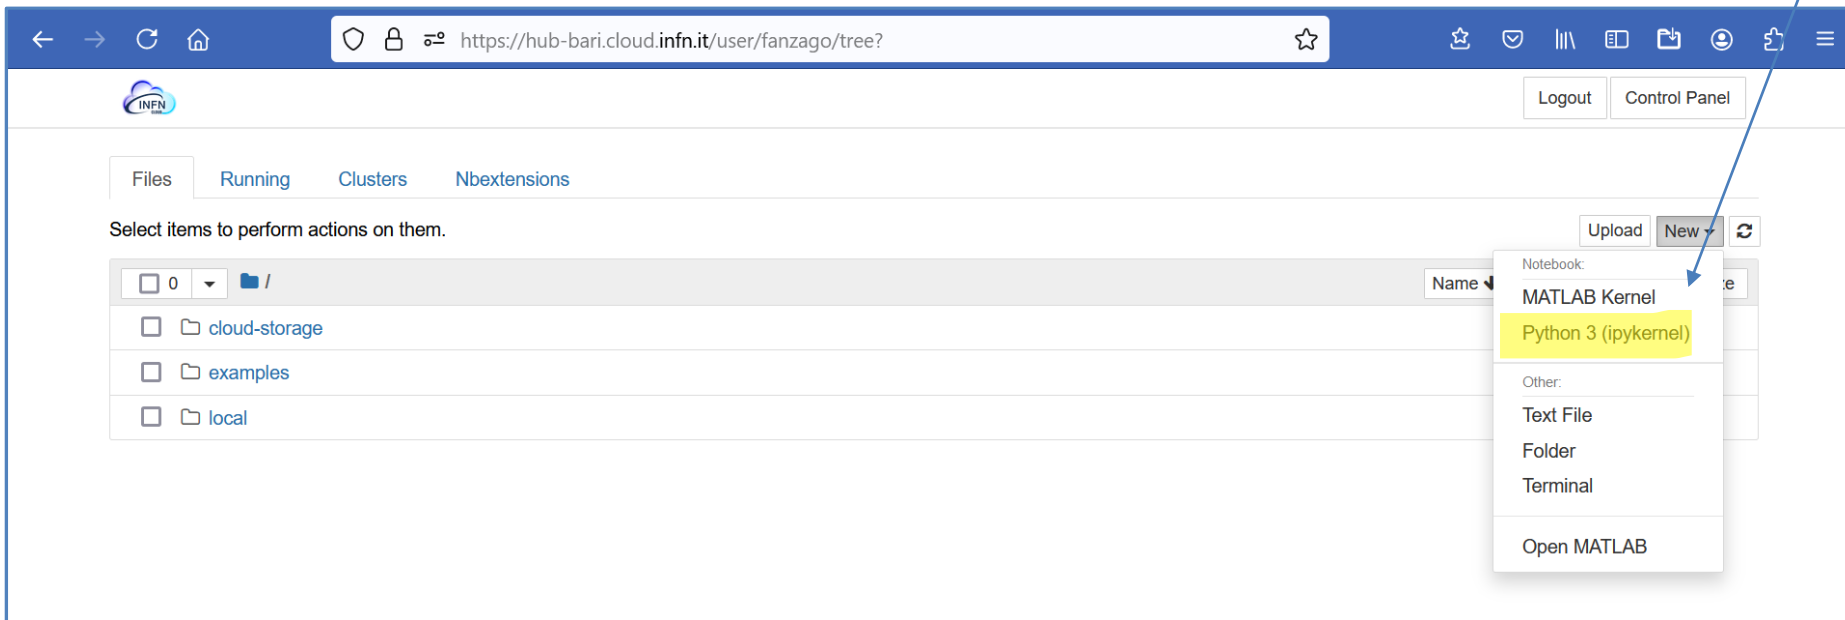


<https://hub.cloud.infn.it>



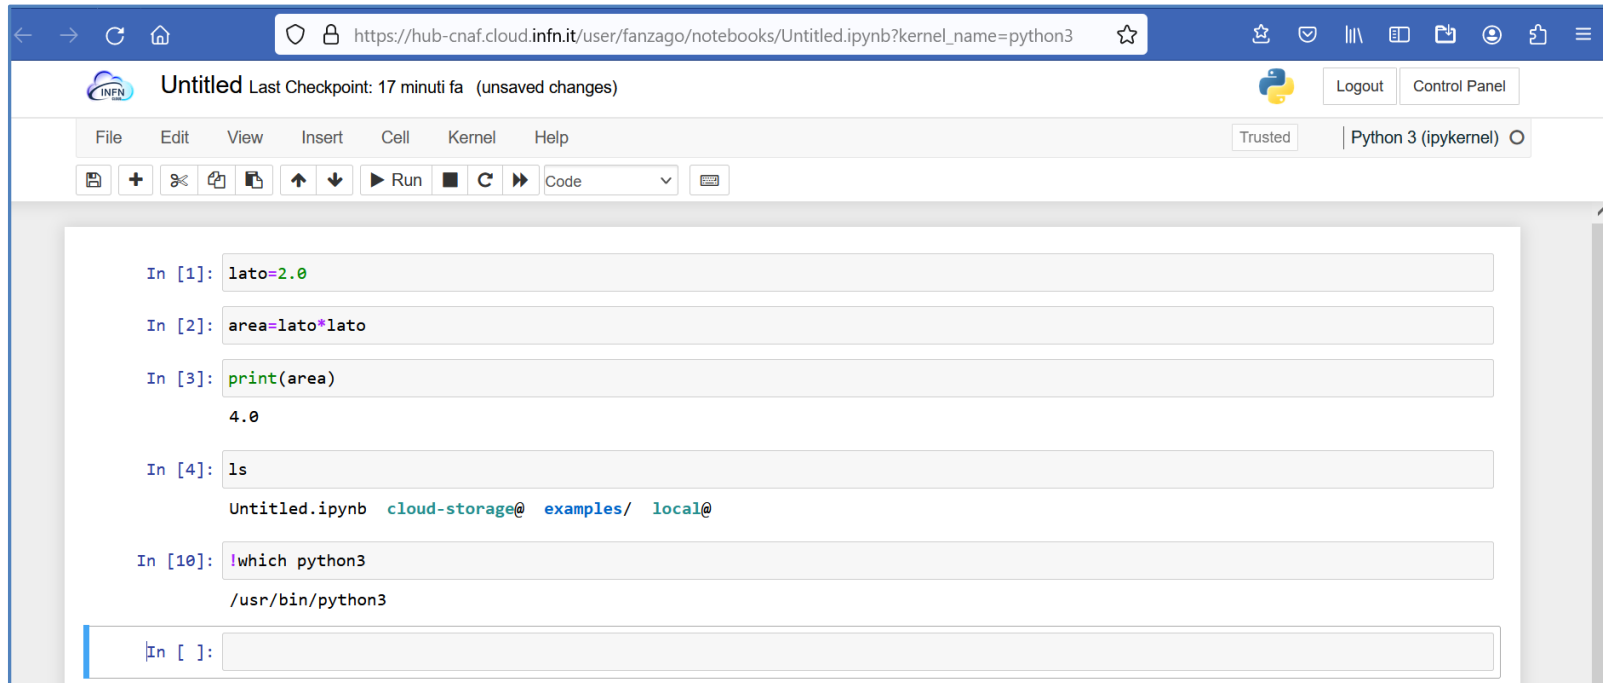
# Jupyter notebook interface

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



# 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].



```
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

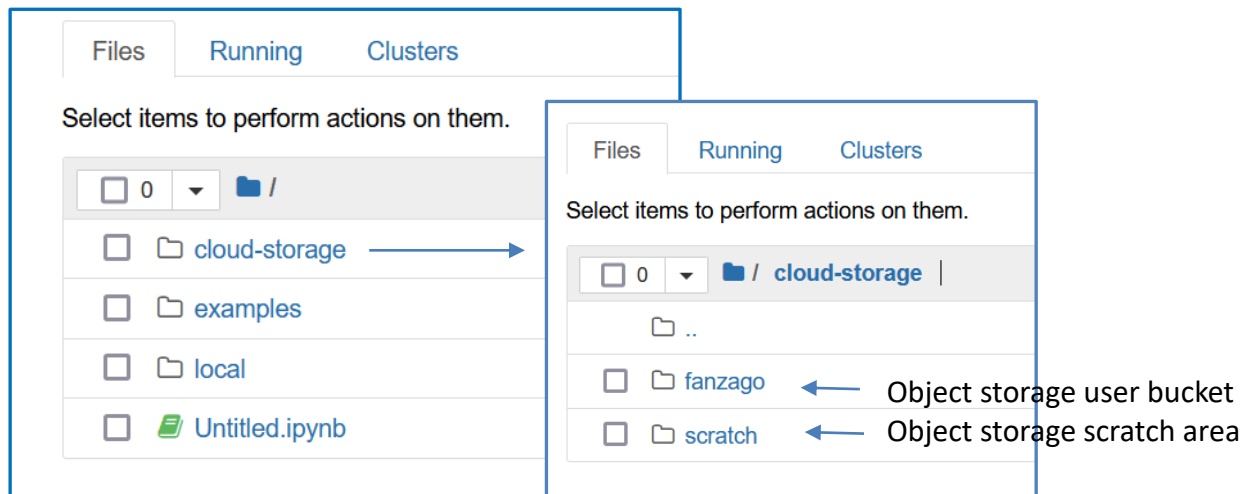
%lsmagic  
%run <script.py>  
%load <script.py>  
%who - list of python variables  
%pinfo <var> - description of variable  
%env - environment variables

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

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



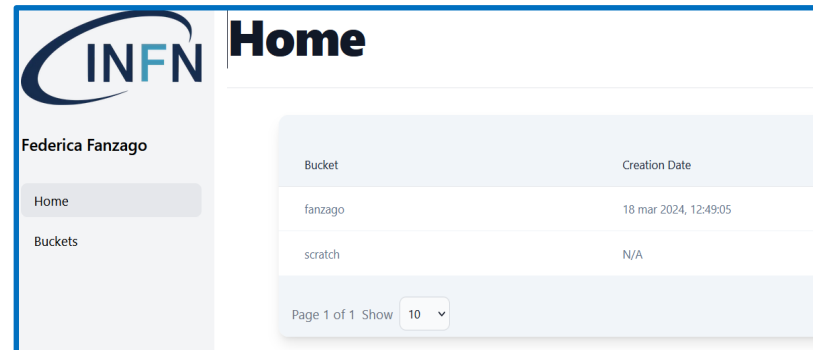
The screenshot shows two panels of the JupyterLab file browser. The left panel shows the root directory with subdirectories: cloud-storage, examples, local, and Untitled.ipynb. An arrow points from 'cloud-storage' to the right panel. The right panel shows the contents of the 'cloud-storage' directory, including '..' and subdirectories 'fanzago' and 'scratch'. Arrows point from 'fanzago' and 'scratch' to labels: 'Object storage user bucket' and 'Object storage scratch area' respectively.

```
# /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#
```

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

Using sts-wire, fuse and Rclone



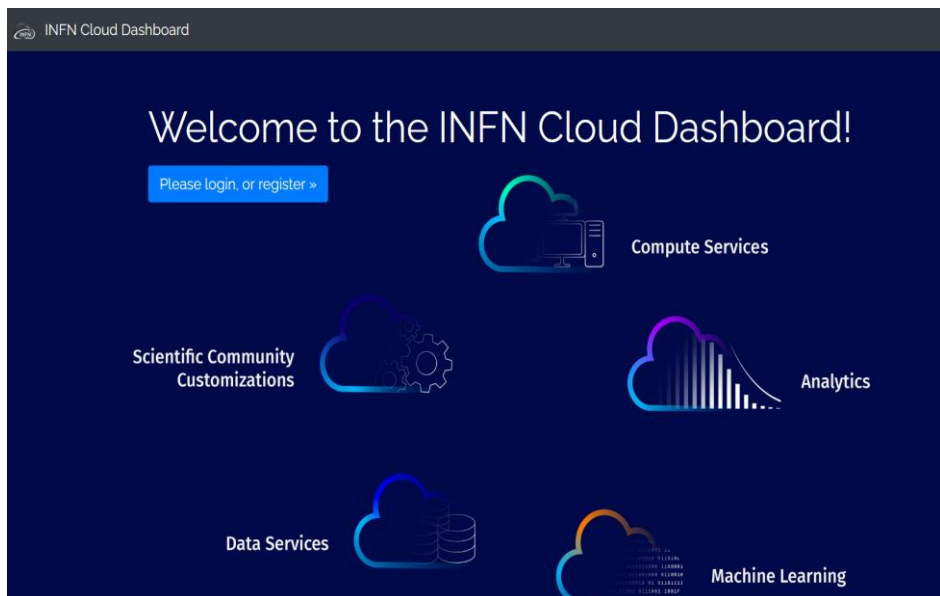
The screenshot shows the Ceph-RGW web interface for user Federica Fanzago. The page title is 'Home'. A table lists the buckets created for the user:

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

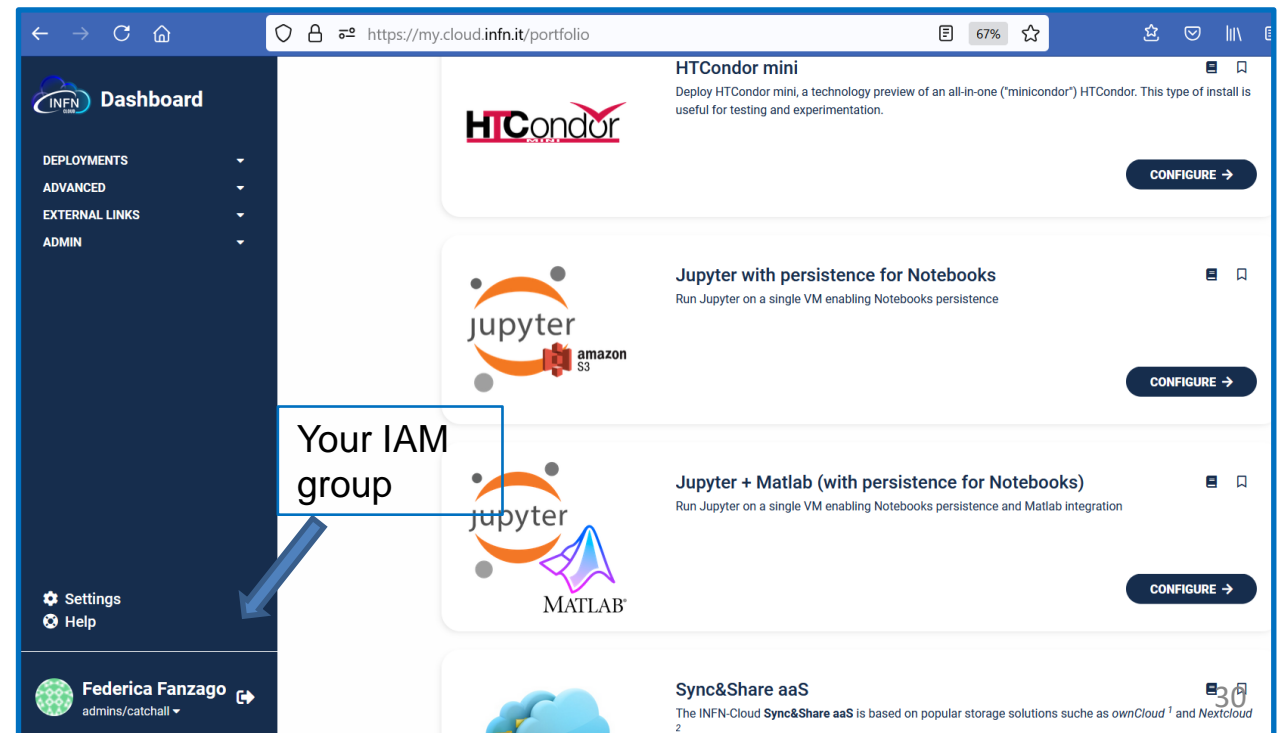
At the bottom of the table, it shows 'Page 1 of 1 Show 10' with a dropdown arrow.

# 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.



Authentication and authorization via IAM



# Filling the form...

### Jupyter with persistence for Notebooks

STEP 1/2

DEPLOYMENT DESCRIPTION (14/50)

Mandatory field

GENERAL AUTHORIZATIONS ADVANCED

NUM CPUS

Number of virtual cpus for the VM

MEM SIZE

 GB

Amount of memory for the VM

ENABLE MONITORING

Enable/disable monitoring

JUPYTER IMAGES

Hub and Lab image

Default image

JUPYTERLAB COLLABORATIVE

enable the jupyter collaborative service

JUPYTERLAB COLLABORATIVE IMAGE

Default image for jupyter collaborative service

CONTACT EMAIL

Email address of certificate management administrator

PORTS

+ Add rule

The necessary ports for Jupyter are already configured

PORTS to open on the VM

CERTIFICATE TYPE

letsencrypt-prod

CANCEL CONTINUE

### Jupyter with persistence for Notebooks

STEP 1/2

DEPLOYMENT DESCRIPTION (14/50)

GENERAL AUTHORIZATIONS ADVANCED

IAM GROUPS

Authorized group of users

IAM groups for authorization management

IAM ADMIN GROUPS

IAM groups for JupyterHub ADMIN authorization management

CANCEL CONTINUE

User guide:

[https://guides.cloud.infn.it/docs/users-guides/en/latest/users\\_guides/sysadmin/compute/jh\\_with\\_persistence.html](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

GENERAL AUTHORIZATIONS **ADVANCED**

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

**CANCEL** **CONTINUE**

### My deployments

Show 10 entries

Refresh + New deployment

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 🔄 Refresh + New deployment

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
fed_e_jup_2_new	11ef3e93-1027-9a13-a163-76b2587994cf	CREATE_COMPLETE	2024-07-10 08:04:00	CLOUD-CNAF-T1

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

Remember your ssh key to access the Vm

JupyterHub url

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

**Description:** fed\_e\_jup\_2\_new

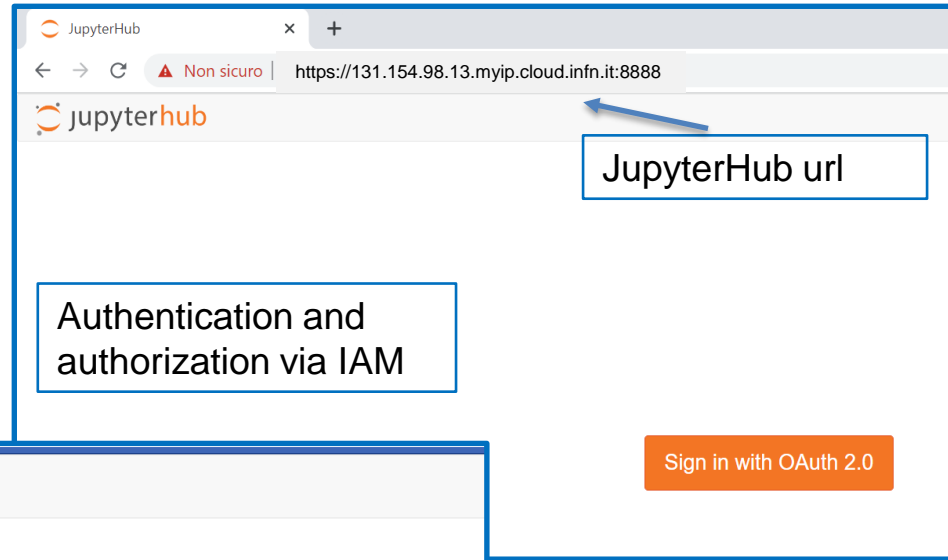
OVERVIEW
INPUT VALUES
OUTPUT VALUES

**node\_ip:** 131.154.98.13

**jupyter\_endpoint:** <https://131.154.98.13.myip.cloud.infn.it:8888>

**ssh\_account:** fanzago

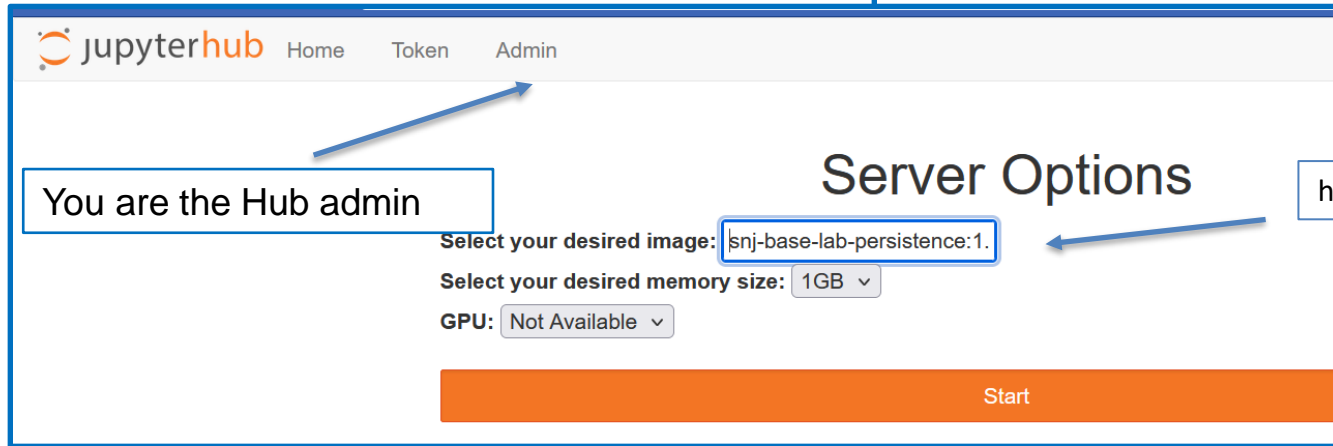
# Your JupyterHub



JupyterHub url

Authentication and authorization via IAM

Sign in with OAuth 2.0



You are the Hub admin

## Server Options

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

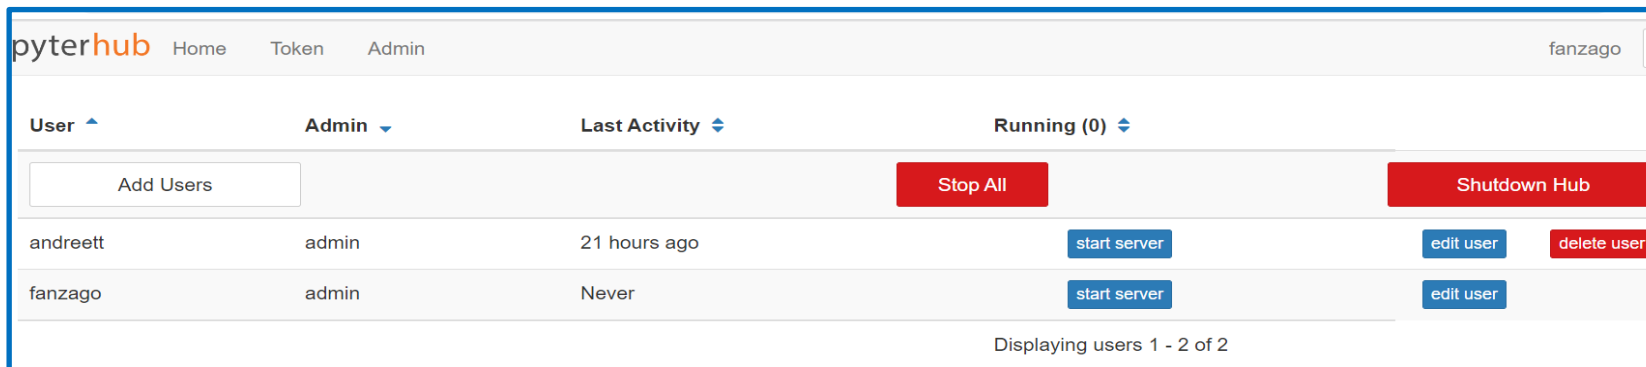
Select your desired memory size: 1GB

GPU: Not Available

Start

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



pyterhub Home Token Admin fanzago

User	Admin	Last Activity	Running (0)
Add Users			
andreett	admin	21 hours ago	start server
fanzago	admin	Never	start server

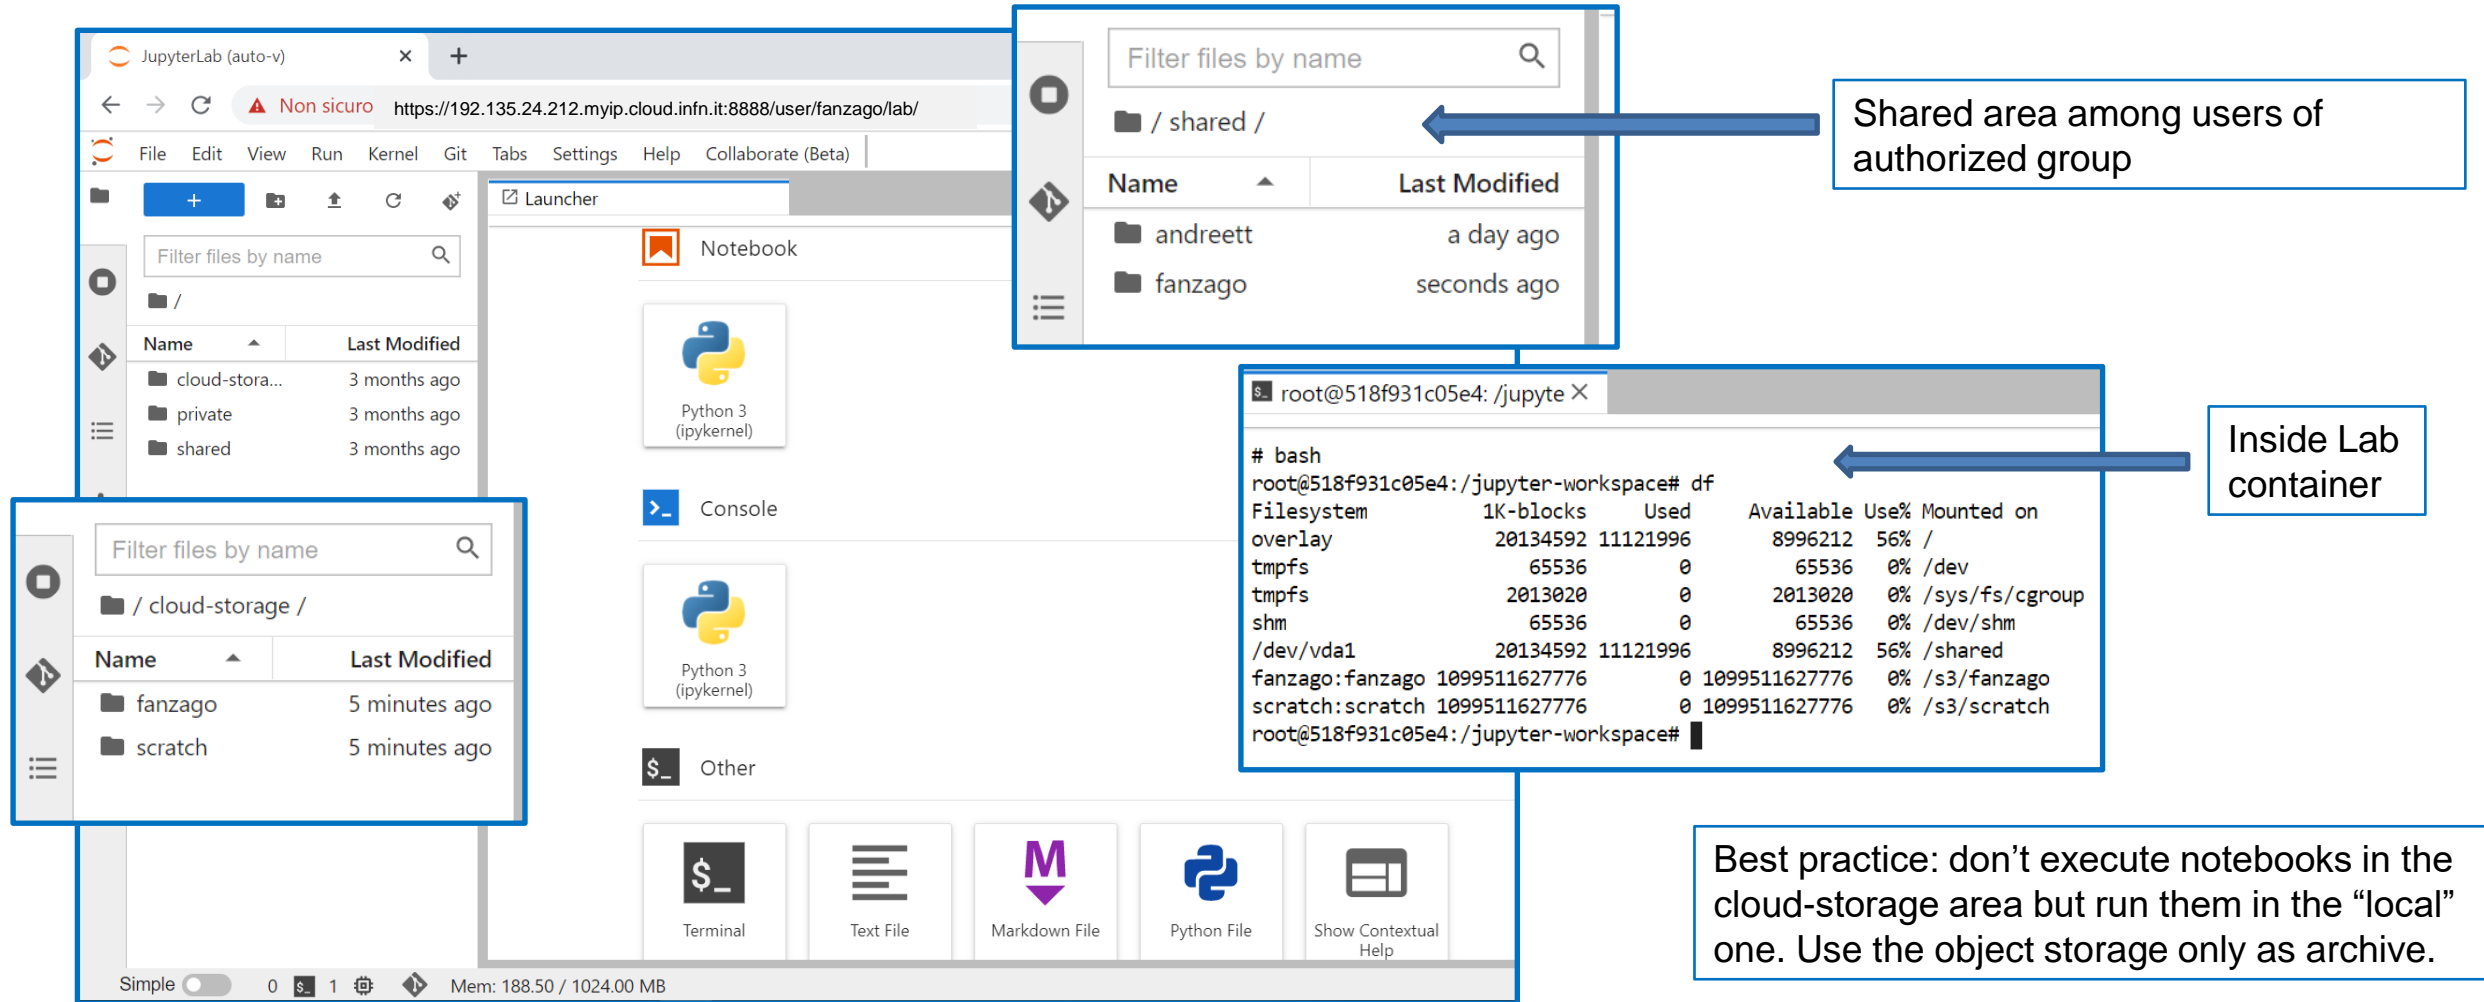
Stop All Shutdown Hub

edit user delete user edit user

Displaying users 1 - 2 of 2

Hub's admin can control users and notebooks

# And your JupyterLab



The screenshot shows the JupyterLab interface with several callouts:

- Shared area among users of authorized group:** Points to the file browser showing a `/ shared /` directory containing `andrett` and `fanzago` folders.
- Inside Lab container:** Points to a terminal window showing the output of the `df` command, which lists the filesystems and their usage.
- Best practice:** A text box stating: "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."

The terminal output of the `df` command is as follows:

```

root@518f931c05e4: /jupyter X
# 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#

```

# Containers in the VM

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

Container Lab

```
root@vnode-0:~# docker ps
CONTAINER ID   IMAGE
a3c03a94c48c  harbor.cloud.infn.it/datacloud-templates/snj-base-lab-persistence:1.2.0
p
4c0121547797  harbor.cloud.infn.it/datacloud-templates/snj-base-jhub:1.2.0
p
930e5a55145b  jupyterhub/configurable-http-proxy
:8001->8001/tcp, :::8001->8001/tcp, 8000/tcp, 0.0.0.0:8888->8888/tcp, :::8888->8888/tcp
40e1cd49b0b7  harbor.cloud.infn.it/cache/library/traefik:latest
:80->80/tcp, :::80->80/tcp, 0.0.0.0:443->443/tcp, :::443->443/tcp
root@vnode-0:~#
```

CONTAINER ID	IMAGE	COMMAND NAMES	CREATED	STATUS	PORTS
a3c03a94c48c	harbor.cloud.infn.it/datacloud-templates/snj-base-lab-persistence:1.2.0	"tini -s -- jupyterh..." jupyter-fanzago	2 days ago	Up 2 days	8889/tc
4c0121547797	harbor.cloud.infn.it/datacloud-templates/snj-base-jhub:1.2.0	"/usr/bin/python3 /u..." jupyterhub-jupyterhub-1	2 days ago	Up 2 days	8000/tc
930e5a55145b	jupyterhub/configurable-http-proxy	"/bin/sh -c 'while ...'" jupyterhub-http_proxy-1	2 days ago	Up 2 days	0.0.0.0
40e1cd49b0b7	harbor.cloud.infn.it/cache/library/traefik:latest	"/entrypoint.sh --lo..." proxy	2 days ago	Up 2 days	0.0.0.0

Container Hub

# 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>