Rosetta

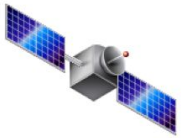

A container-centric science platform for interactive workloads

Rosetta in a nutshell

- Web-based science platform supporting GUI-based, interactive computing
- Batteries included, runs entirely on Docker containers
- Executes workloads in Docker or Singularity containers on a number of computing elements:
    - local CPUs;
    - remote servers (including commercial Clouds);
    - HPC Clusters (i.e. Slurm).
- Requires minimum effort to get Rosetta-compatible (just extend base containers)
- Allows users to setup and run their own containers, autonomously.

Example

- Task list and remote desktop container connection via automatic tunneling and Web VNC

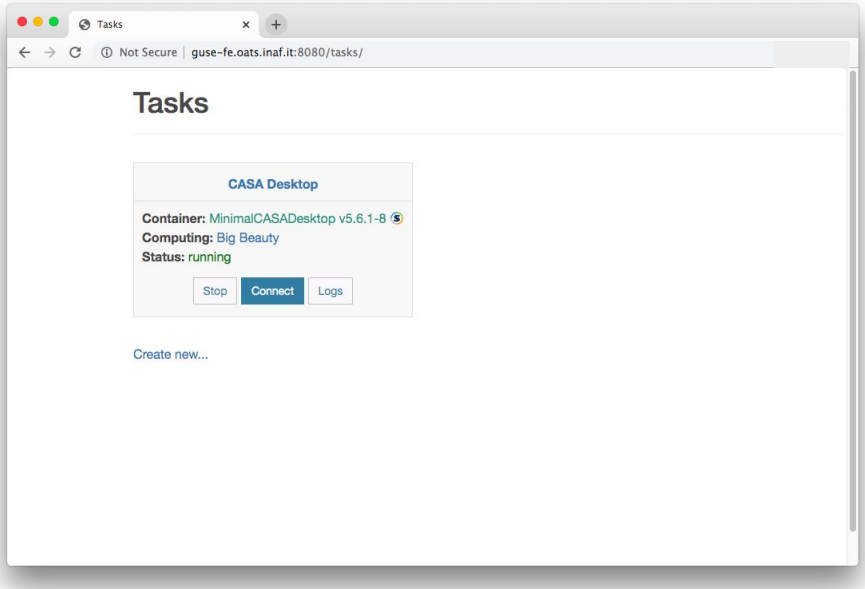

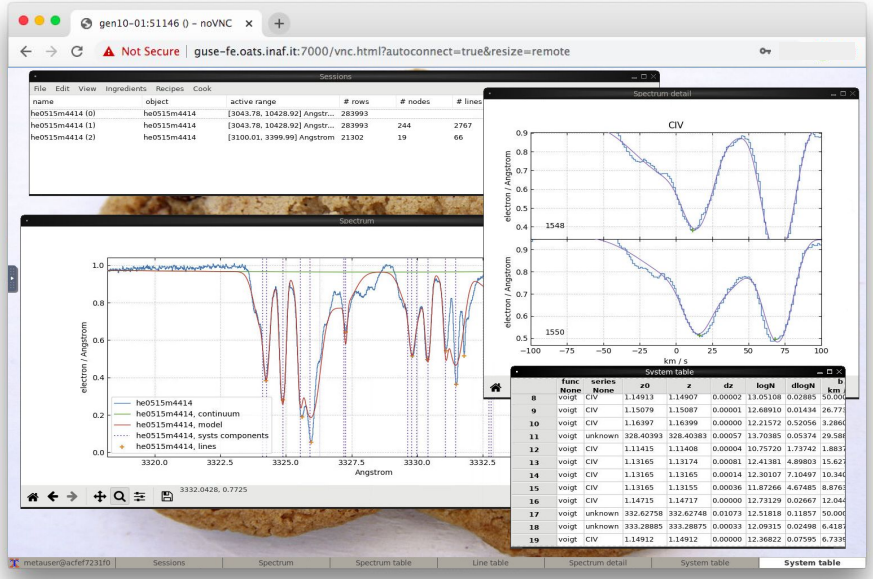

A step back: what container actually are?

> “Containers are a solution to the problem of how to get software to run reliably when moved from one computing environment to another.”

cio.com

A step back: what container actually are?

“Containers are a solution to the problem of how to get software to run reliably when moved from one computing environment to another.”

cio.com

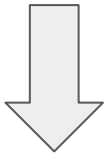

A.K.A. the dependency hell problem

The “dependency hell” problem

Mike wants to use a new software.

Mike cannot find a precompiled version that works with his OS and/or libraries.

Mike asks/Google for help and gets some basic instructions - like “compile it”.

Mike starts downloading all the development environment, and soon realizes that he needs to upgrade (or downgrade!) some parts of his main Operating Systems.

During this process, something goes wrong.

Mikes spends an afternoon fixing his own OS, and all the next day in trying to compile the software. Which at the end turns out not to do what he wanted.

The “dependency hell” problem: solutions spectrum

The “dependency hell” problem: solutions spectrum

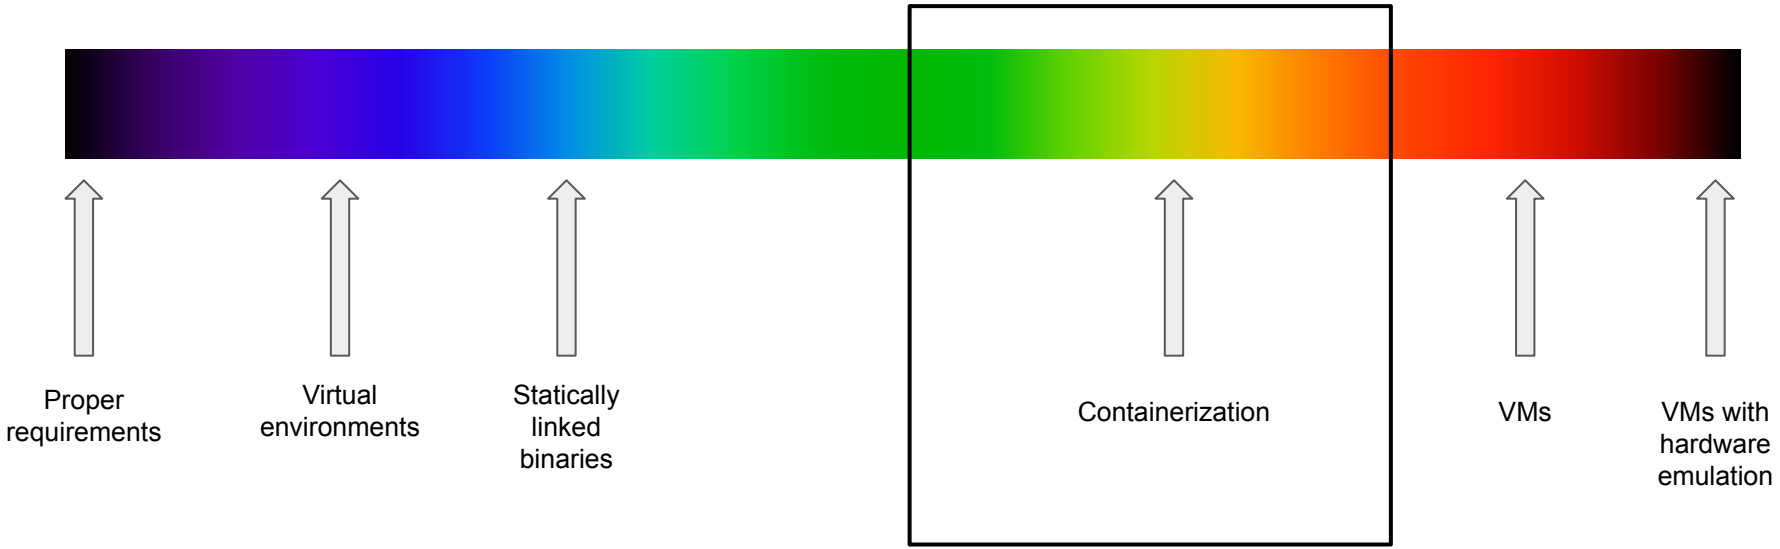

The “dependency hell” problem: solutions spectrum

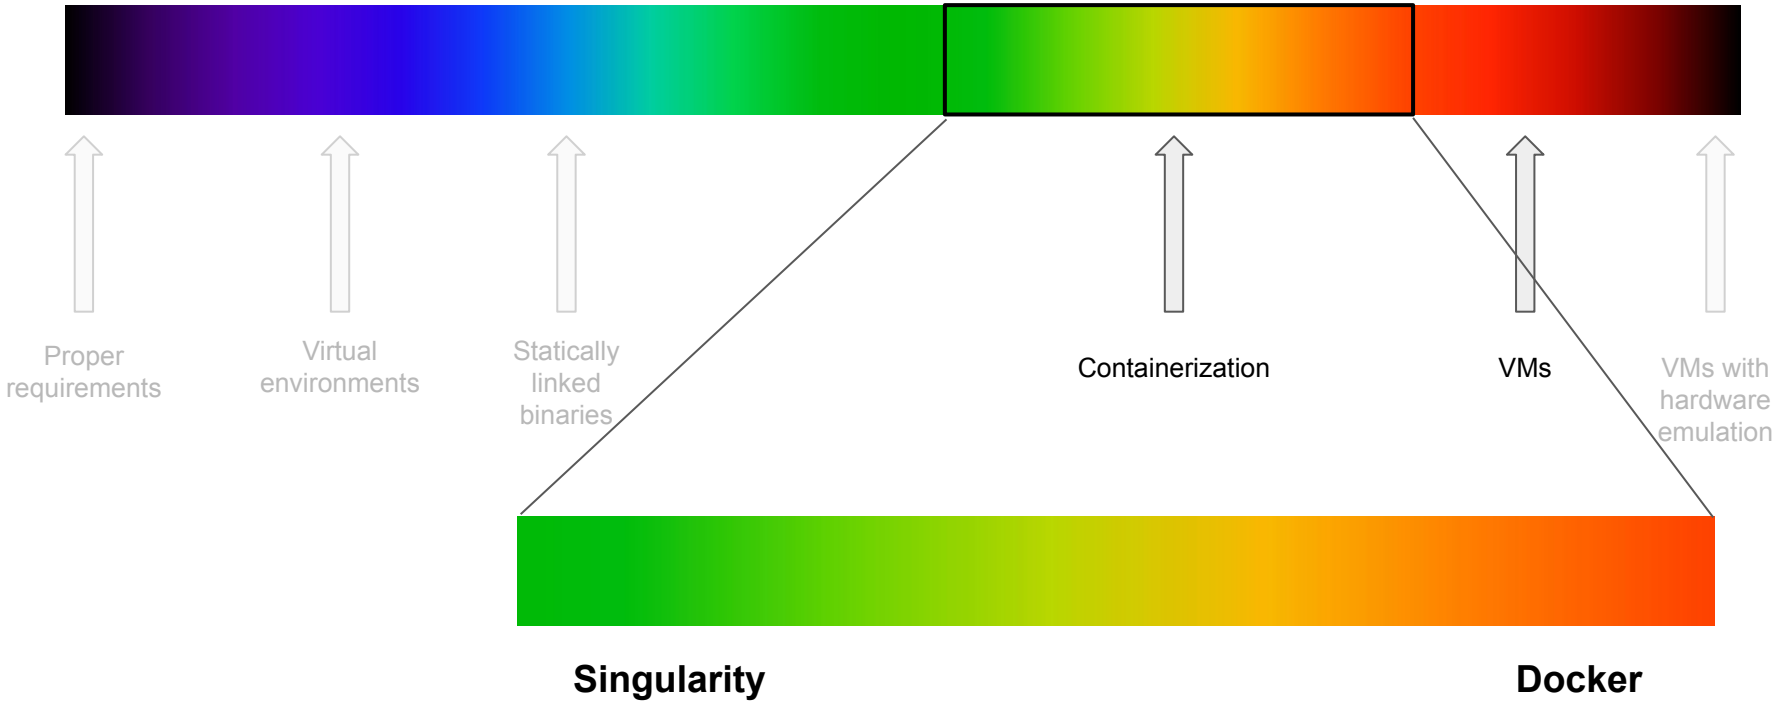

Singularity vs Docker

Singularity	Docker
Scientific computing	IT industry standard
Running container are seen as processes	Running containers are seen as (micro)services
Build as root, run as user	Need near-root access or proper orchestrators
Limited or no support for networking	Extensive support for networking

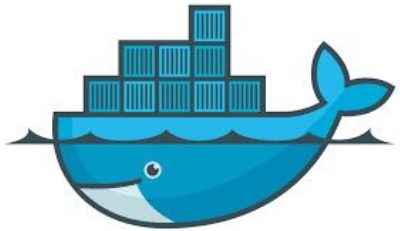

Singularity vs Docker

Singularity	Docker
Filesystem: only partially isolated,directories as $HOME, /tmp, /proc, /sys, and /dev are all binded by default.	**Filesystem:** completely isolated by default, volume or folder binds must be explicitly set
Environment: from the host	**Environment:** from scratch
Network: from the host	**Network:** dedicated subnet

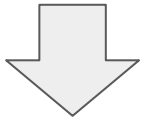

More similar to an environment

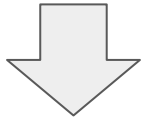

More similar to a virtual machine

Hard truth:

Environments-like containerisation solutions like Singularity do not not solve the dependency hell problem and how to reliably move software around, which in turn **does not guarantee reproducibility**!!

Hard truth:

Environments-like containerisation solutions like Singularity do not not solve the dependency hell problem and how to reliably move software around, which in turn **does not guarantee reproducibility**!!

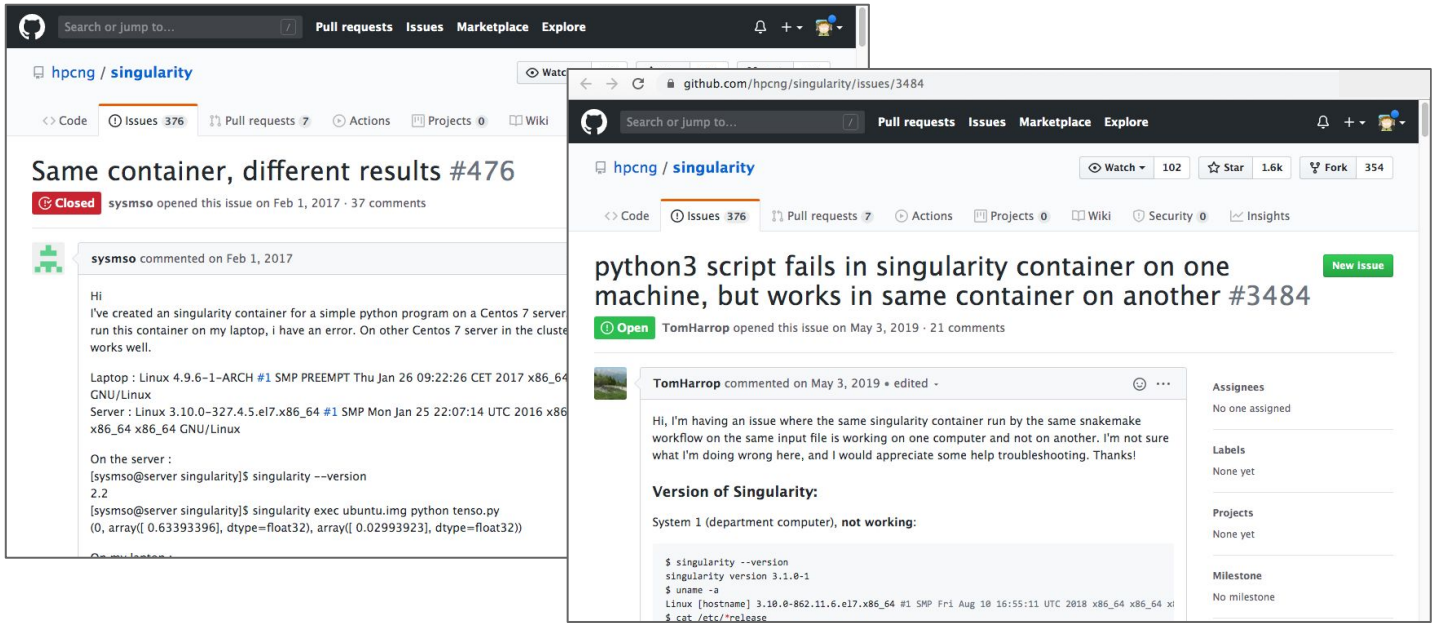

Singularity in Rosetta

Singularity is still very interesting as it can be used autonomously by users without asking much effort from sysadmins or to change the scheduling / workload management.

To overcome its environment-like limitations, in Rosetta Singularity is always forced to run in a specific way which allow to achieve an almost container-like behaviors, as Docker.

This requires a few specific command line switched (set by Rosetta itself) and some extra orchestration (for sandboxing)

```
mkdir -p /tmp/ee3f2776-1b9e-4dd4-a250-f80813dc9fe0_data/tmp && \
mkdir -p /tmp/ee3f2776-1b9e-4dd4-a250-f80813dc9fe0_data/home && \
chmod 700 /tmp/ee3f2776-1b9e-4dd4-a250-f80813dc9fe0_data && \
singularity run --pid --no-home --containall --cleanenv --writable-tmpfs
--workdir /tmp/ee3f2776-1b9e-4dd4-a250-f80813dc9fe0_data/tmp --home=/home/metauser \
-B/tmp/ee3f2776-1b9e-4dd4-a250-f80813dc9fe0_data/home:/home/metauser
```

Container interfaces & graphical applications

One missing bit in using containers (regardless of Docker or Singularity) is that the software you use to interact with your container ***is*** a dependency.

Usually, this dependency is a command-line, or a terminal, which is stable from the 80's.

However, with GUI applications, how to ensure the most common possible dependency?

Container interfaces & graphical applications

Possible options:

1. X11
2. VNC
3. Web-based VNC

Container interfaces & graphical applications

Possible options:

1. X11
2. VNC
3. Web-based VNC

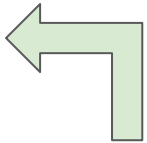

This one. Why? ..because like it or not, web browsers and Javascript are the most standard technologies you can find in your usebase

..therefore:

The interface is the web browser

Reproducible scientific stuff within Singularity

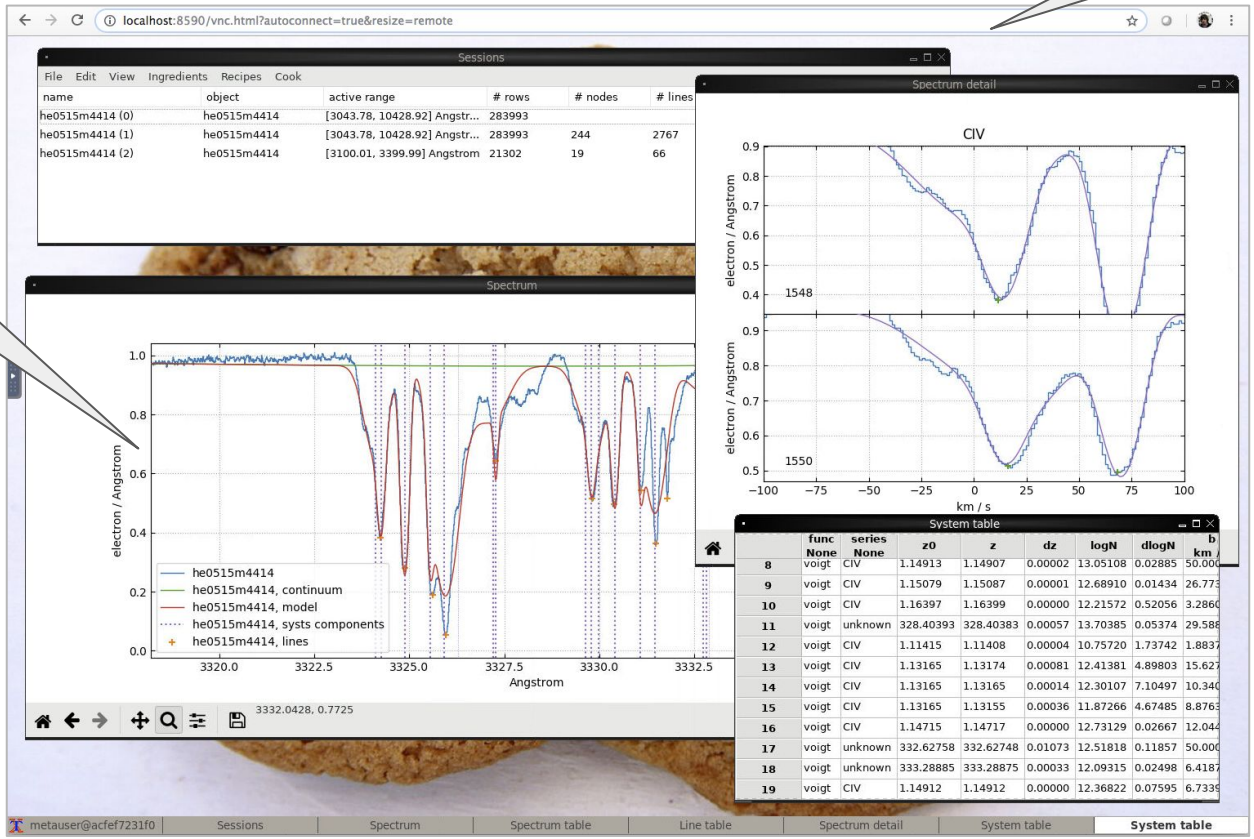

Rosetta quickstart & demo (for users)

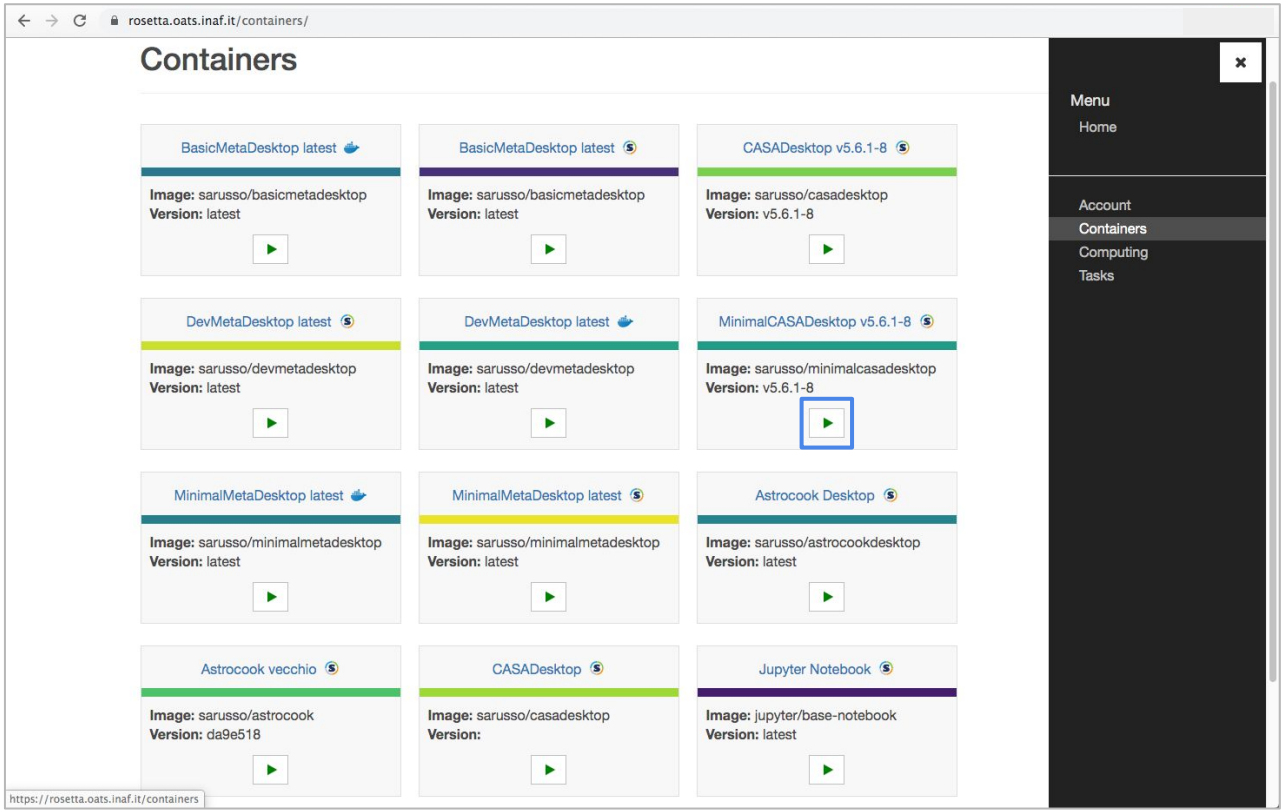

Rosetta quickstart & demo (for users)

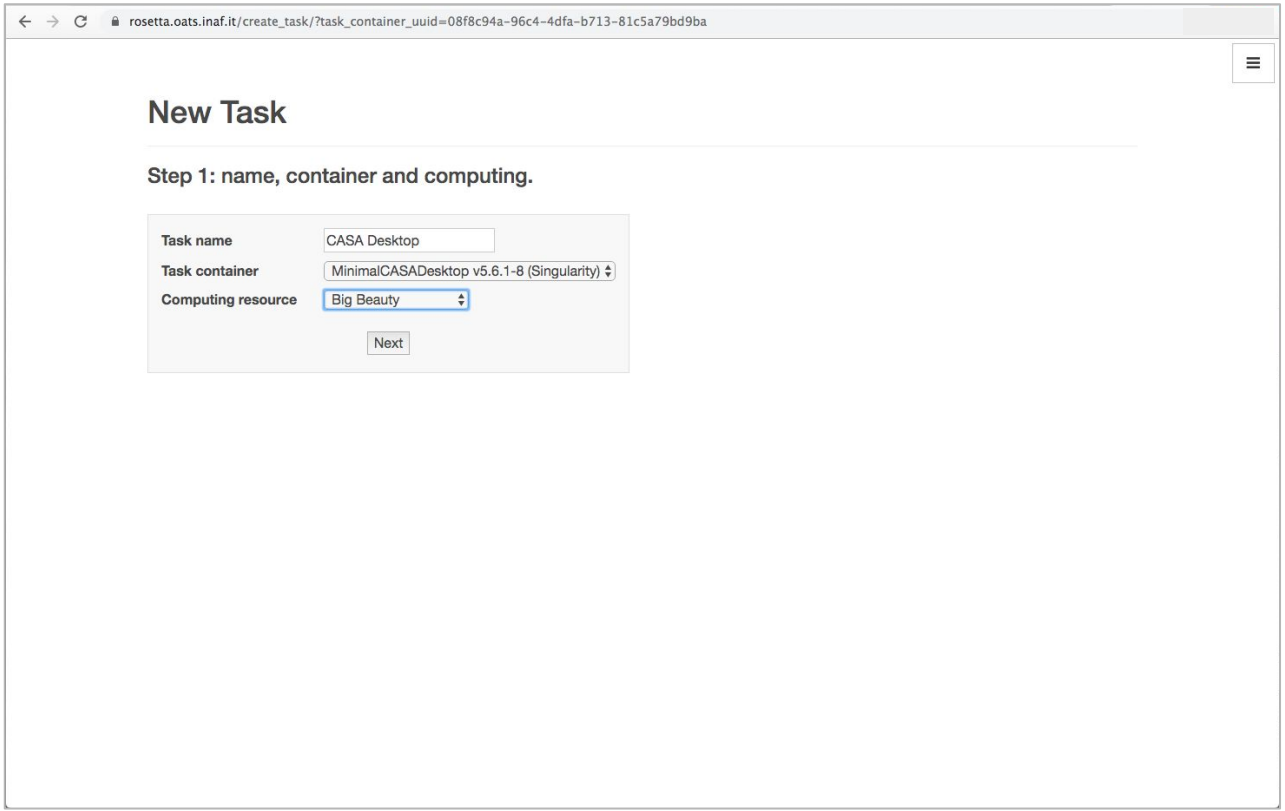

Rosetta quickstart & demo (for users)

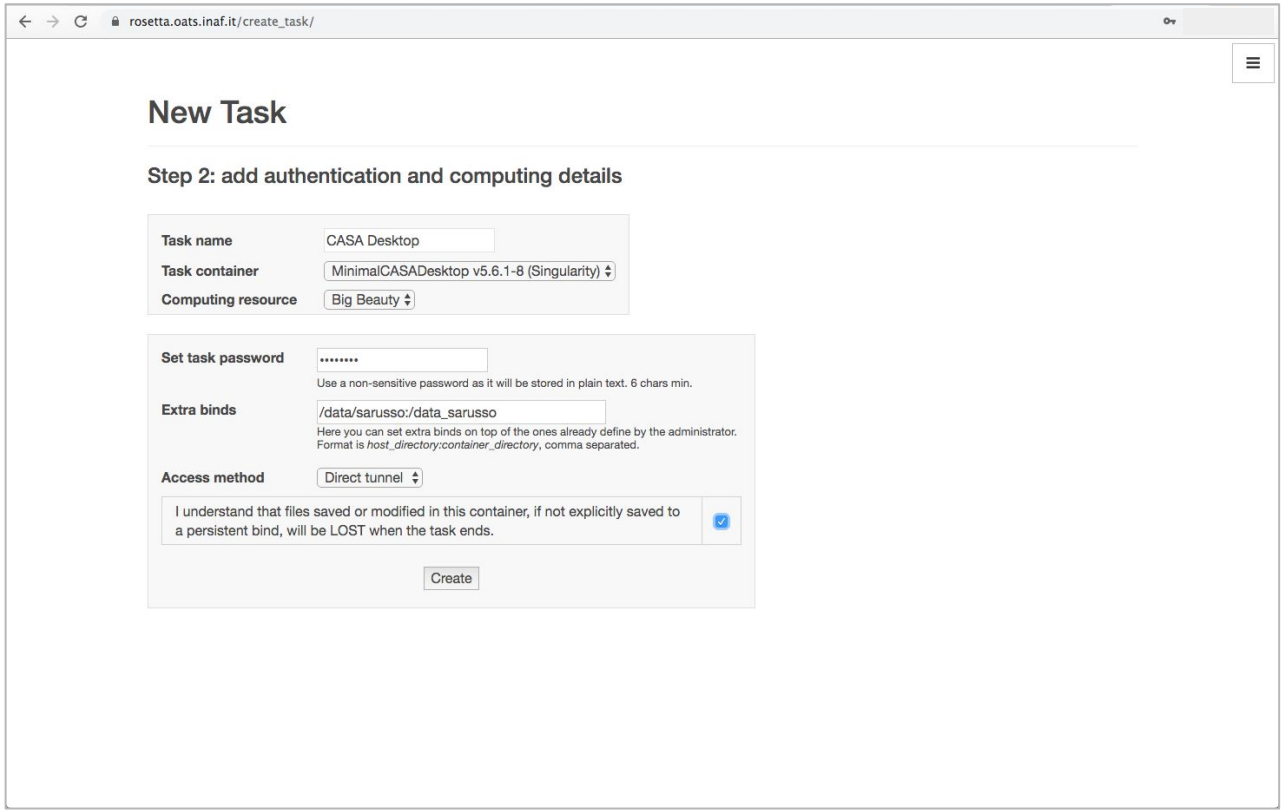

Rosetta quickstart & demo (for users)

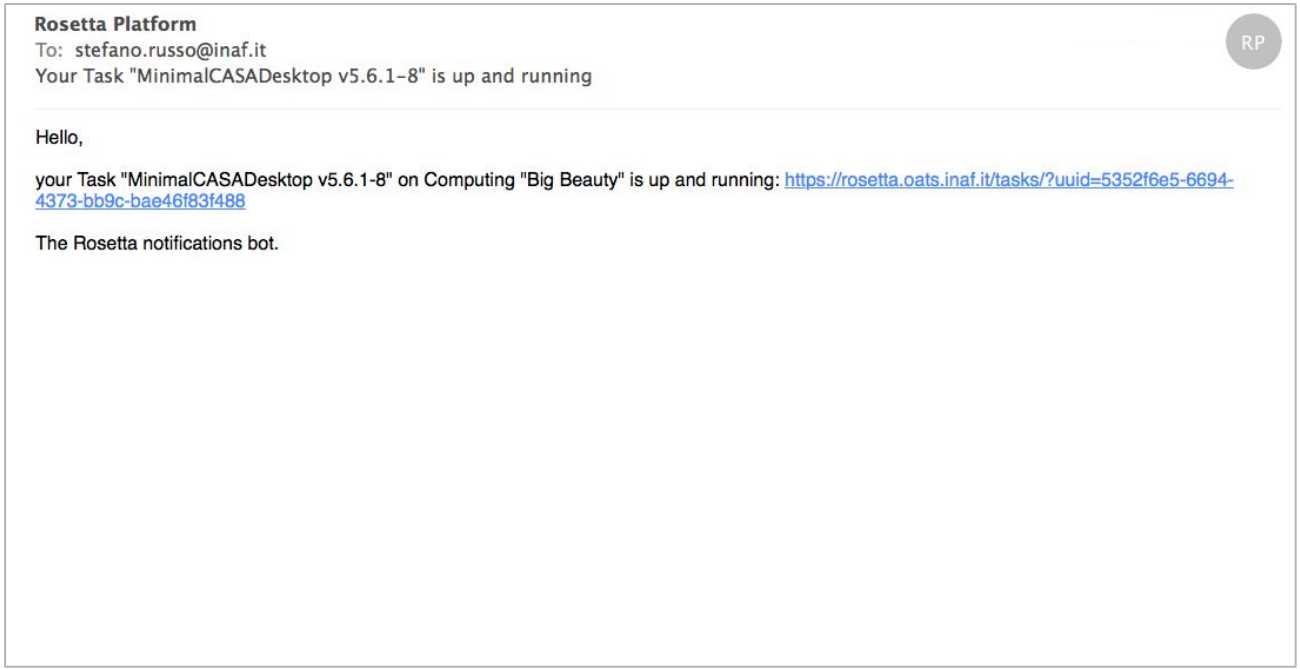

Rosetta Platform
To: stefano.russo@inaf.it
Your Task "MinimalCASADesktop v5.6.1-8" is up and running

RP

Hello,

your Task "MinimalCASADesktop v5.6.1-8" on Computing "Big Beauty" is up and running: https://rosetta.oats.inaf.it/tasks/?uuid=5352f6e5-6694-4373-bb9c-bae46f83f488

The Rosetta notifications bot.

Rosetta quickstart & demo (for users)

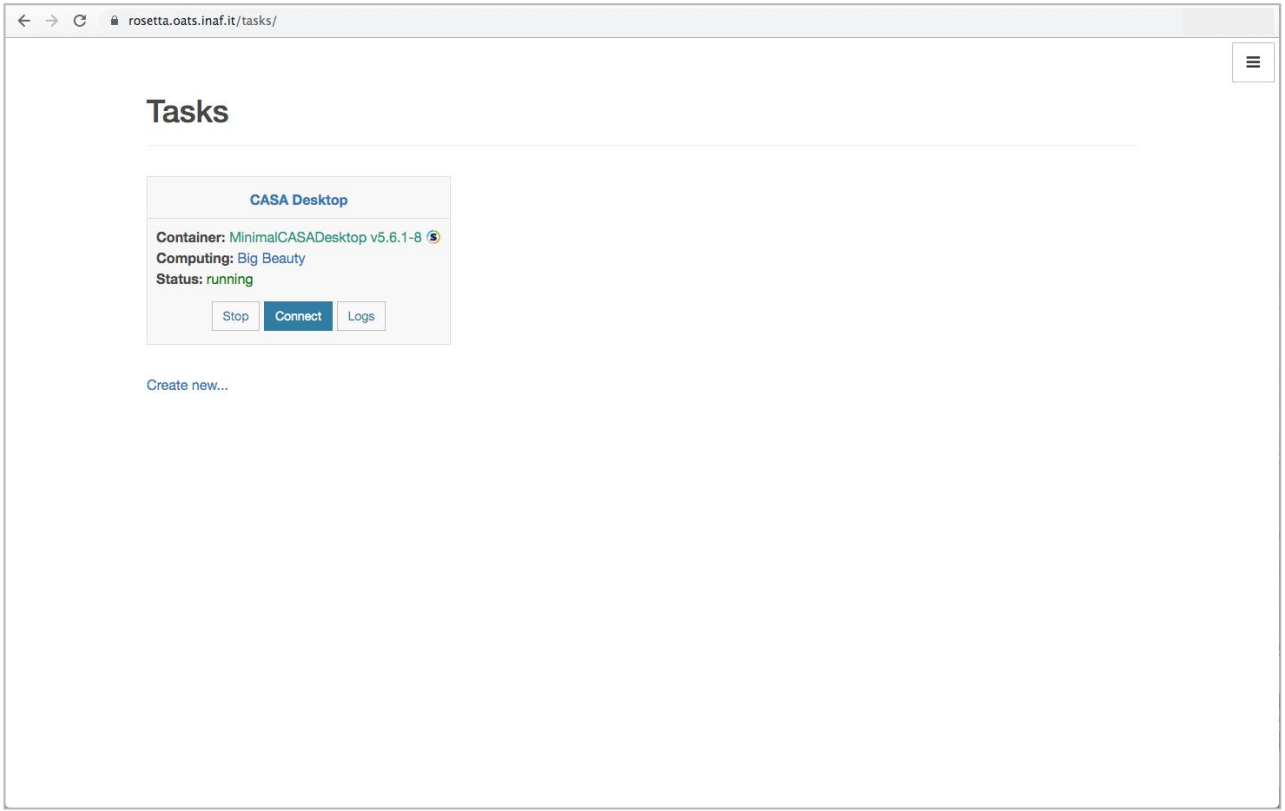

Rosetta quickstart & demo (for users)

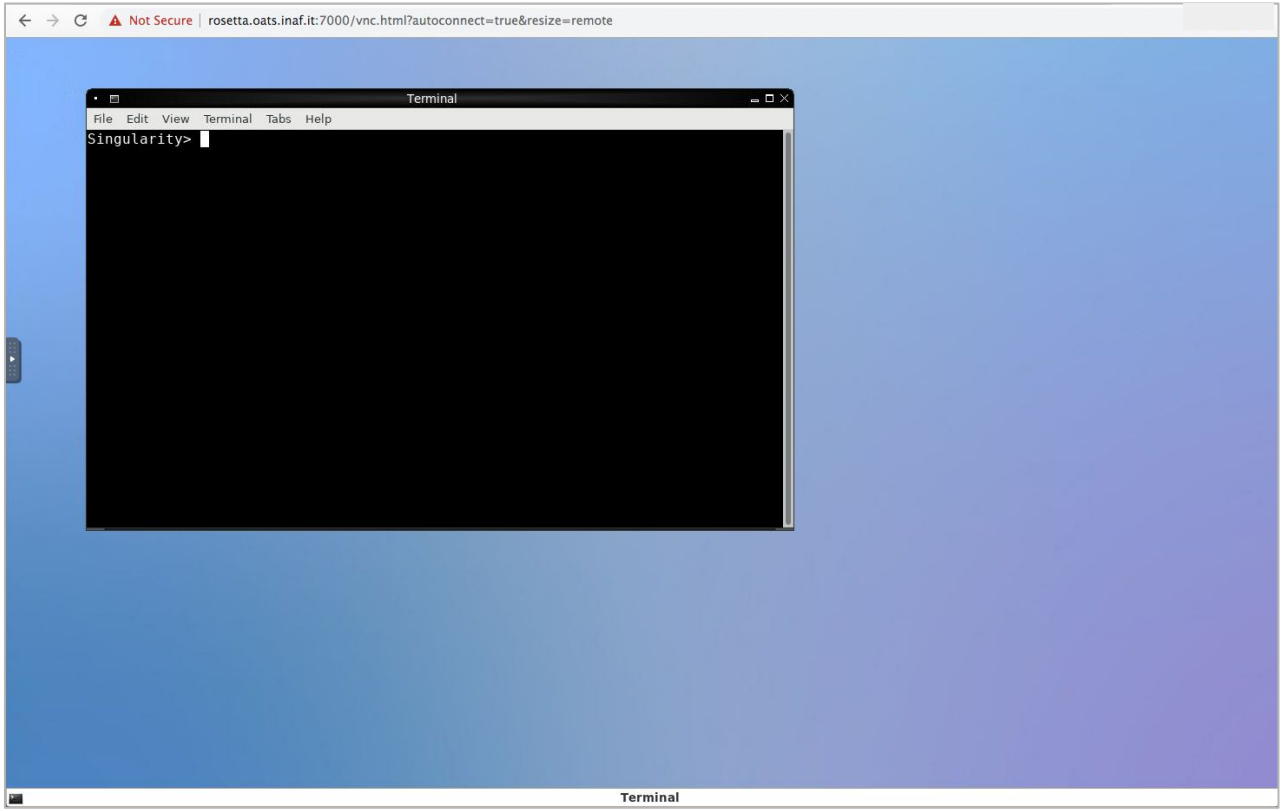

Rosetta quickstart & demo (for users)

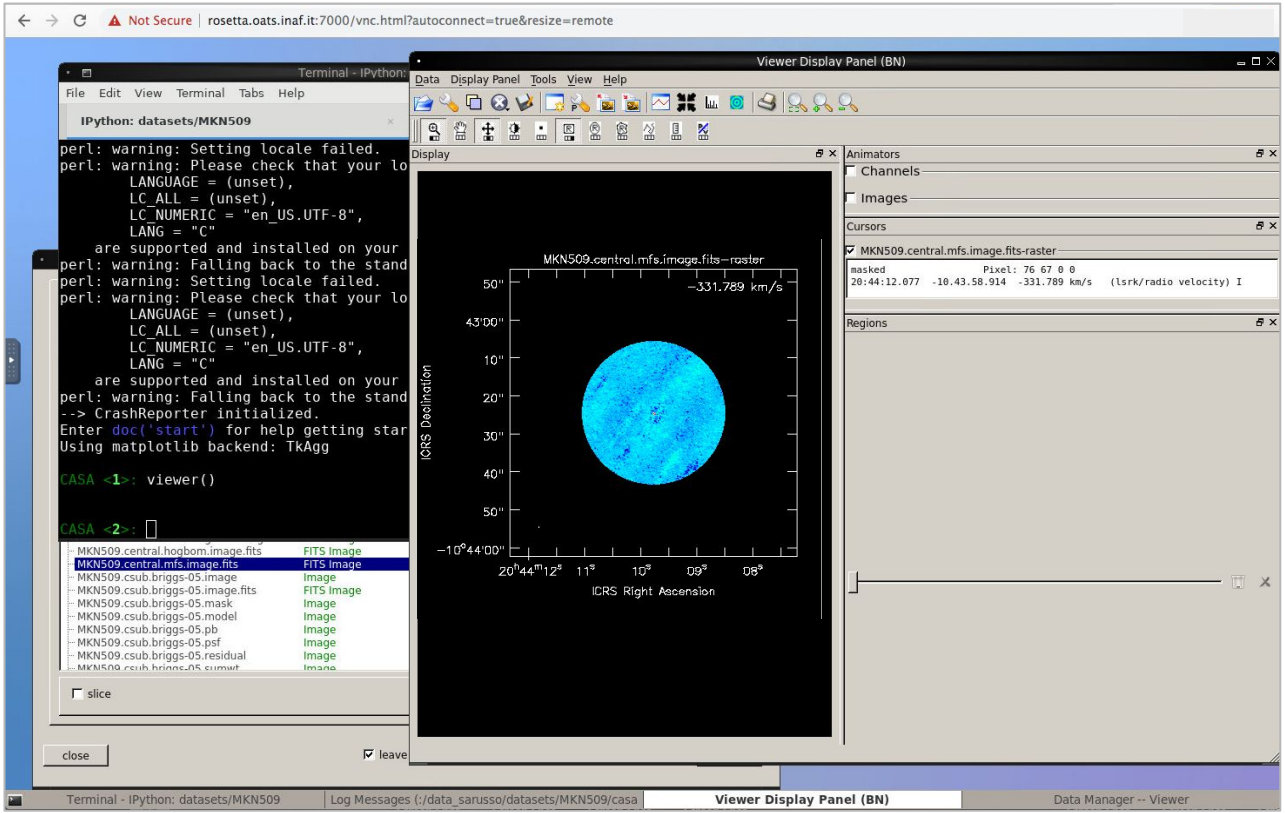

Rosetta quickstart & demo (for users)

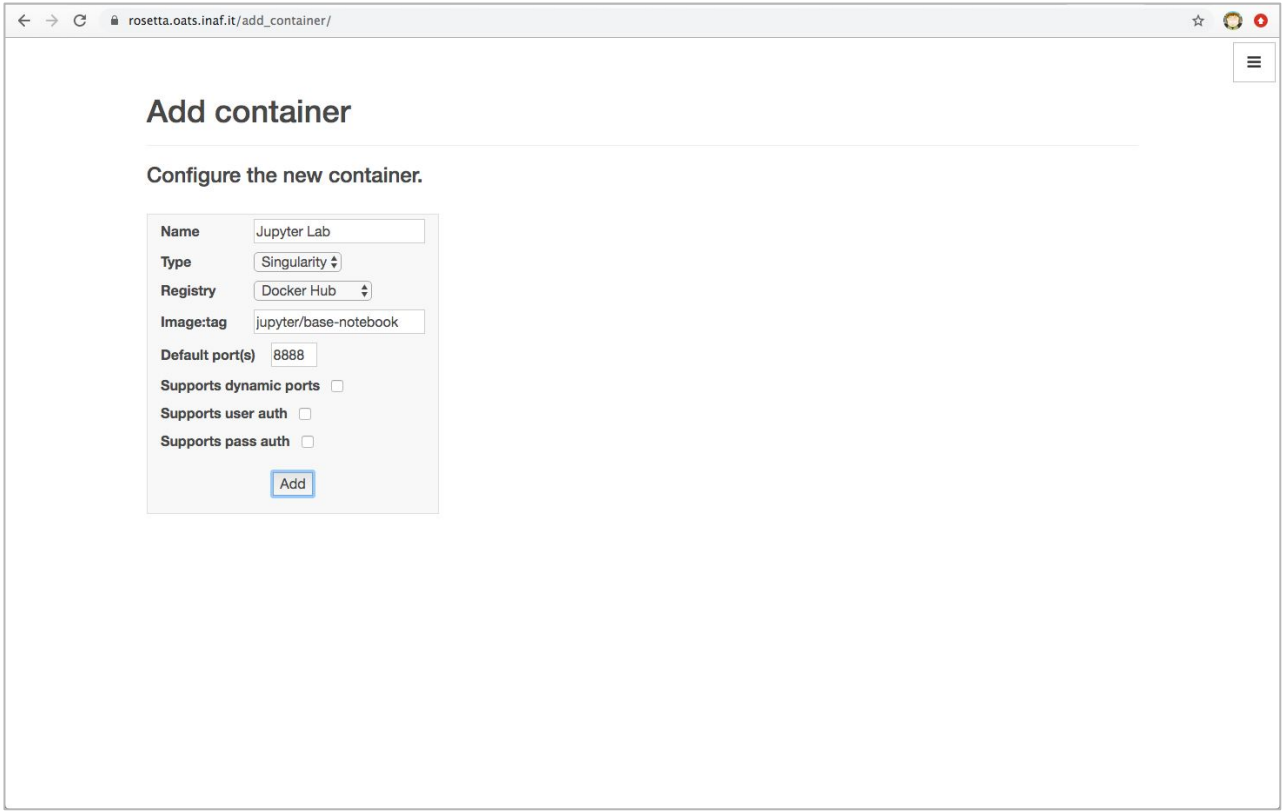

Rosetta quickstart & demo (for users)

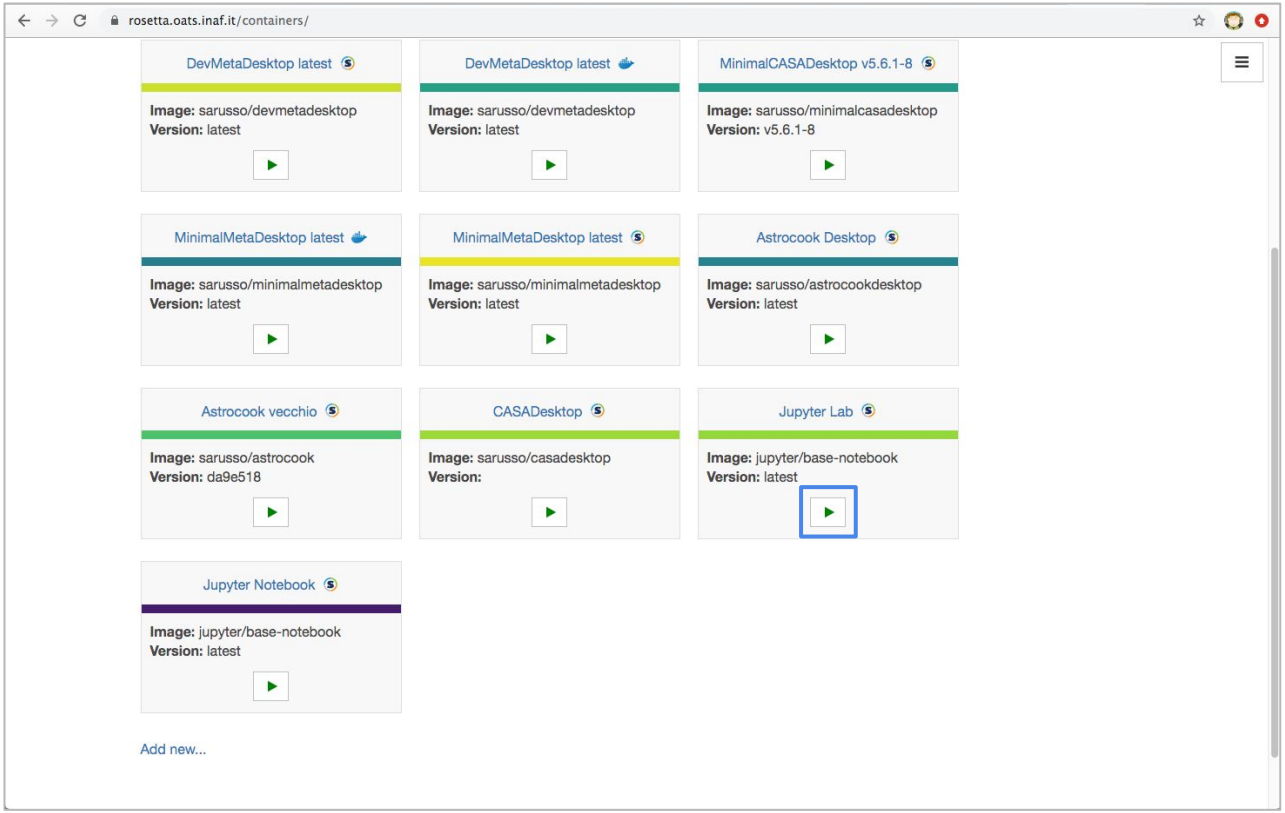

Rosetta quickstart & demo (for users)

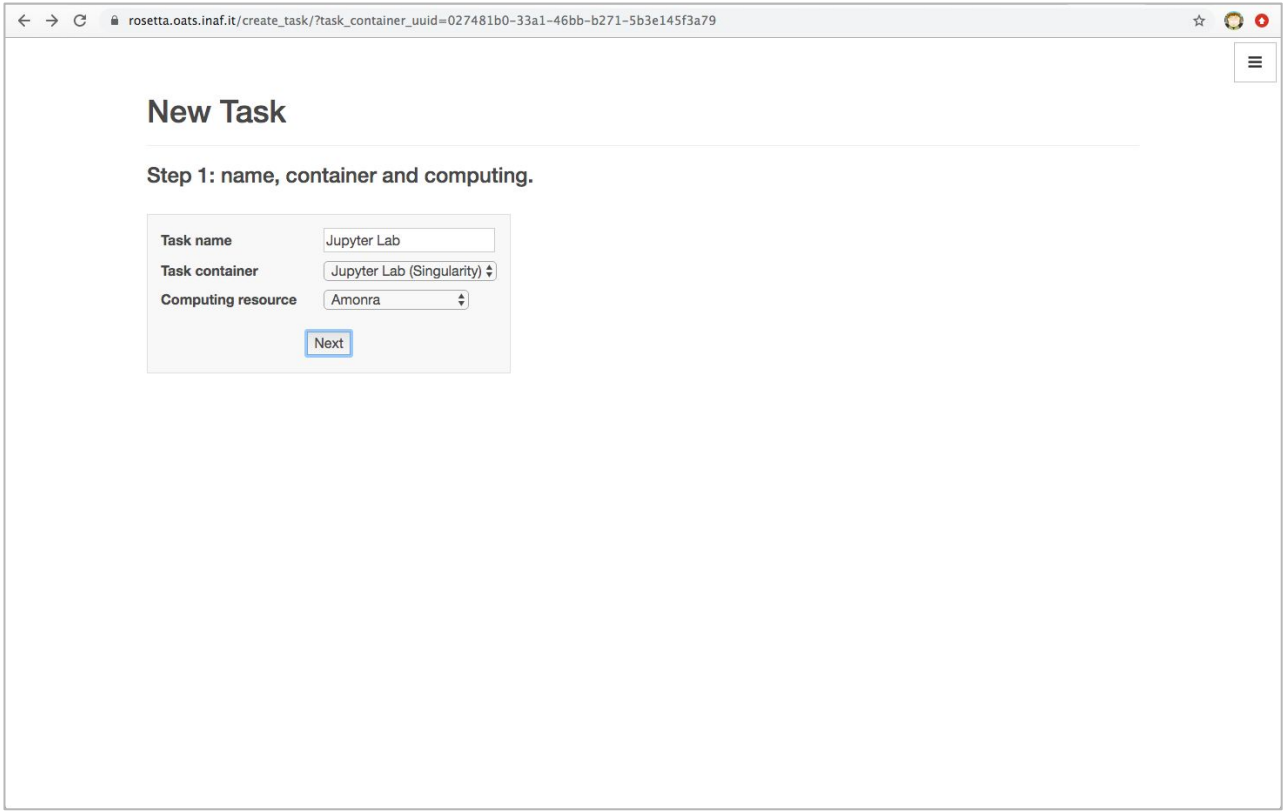

Rosetta quickstart & demo (for users)

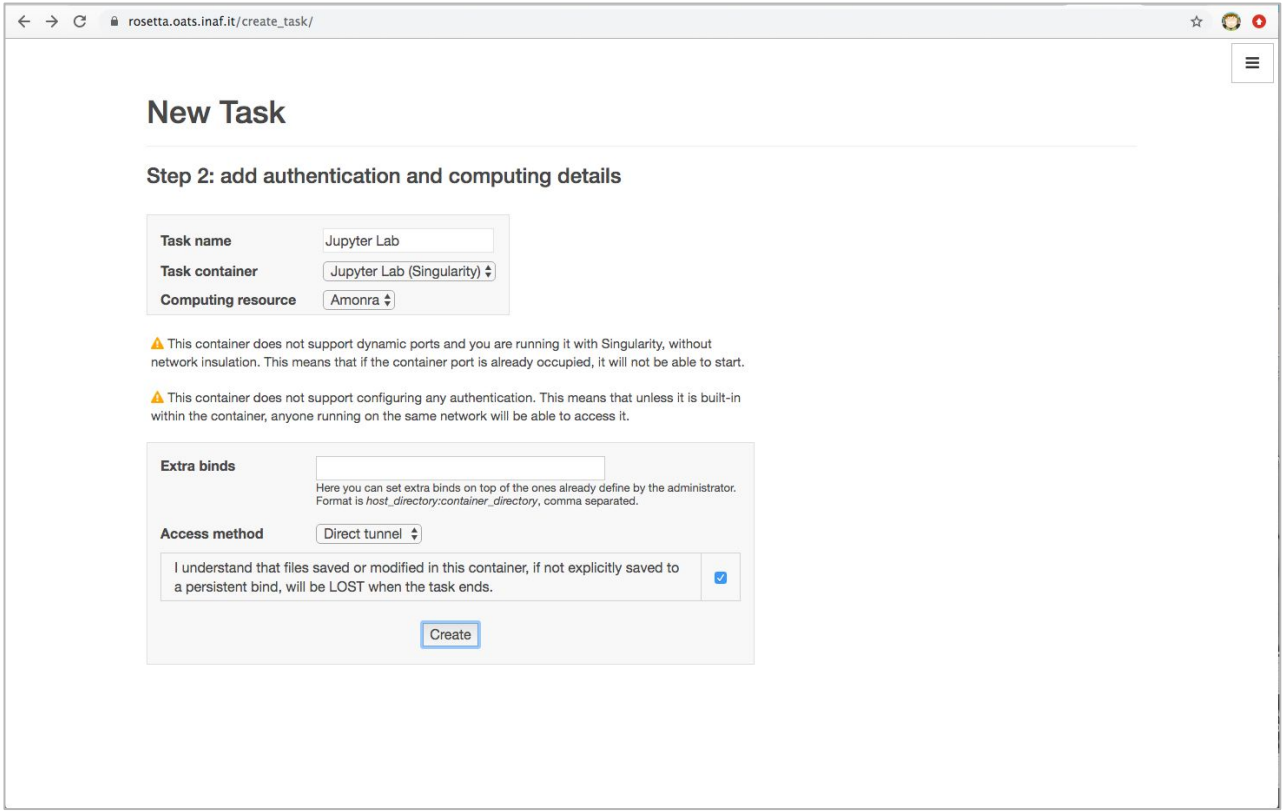

Rosetta quickstart & demo (for users)

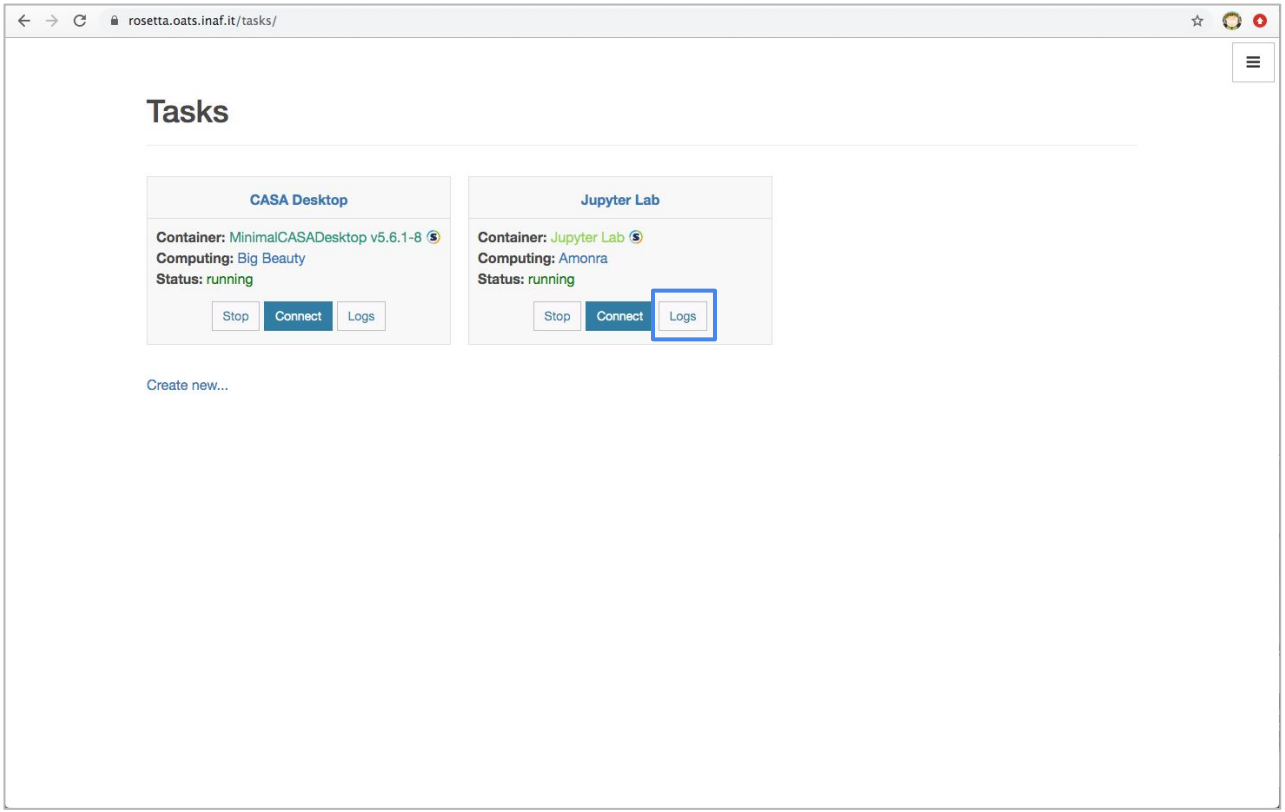

Rosetta quickstart & demo (for users)

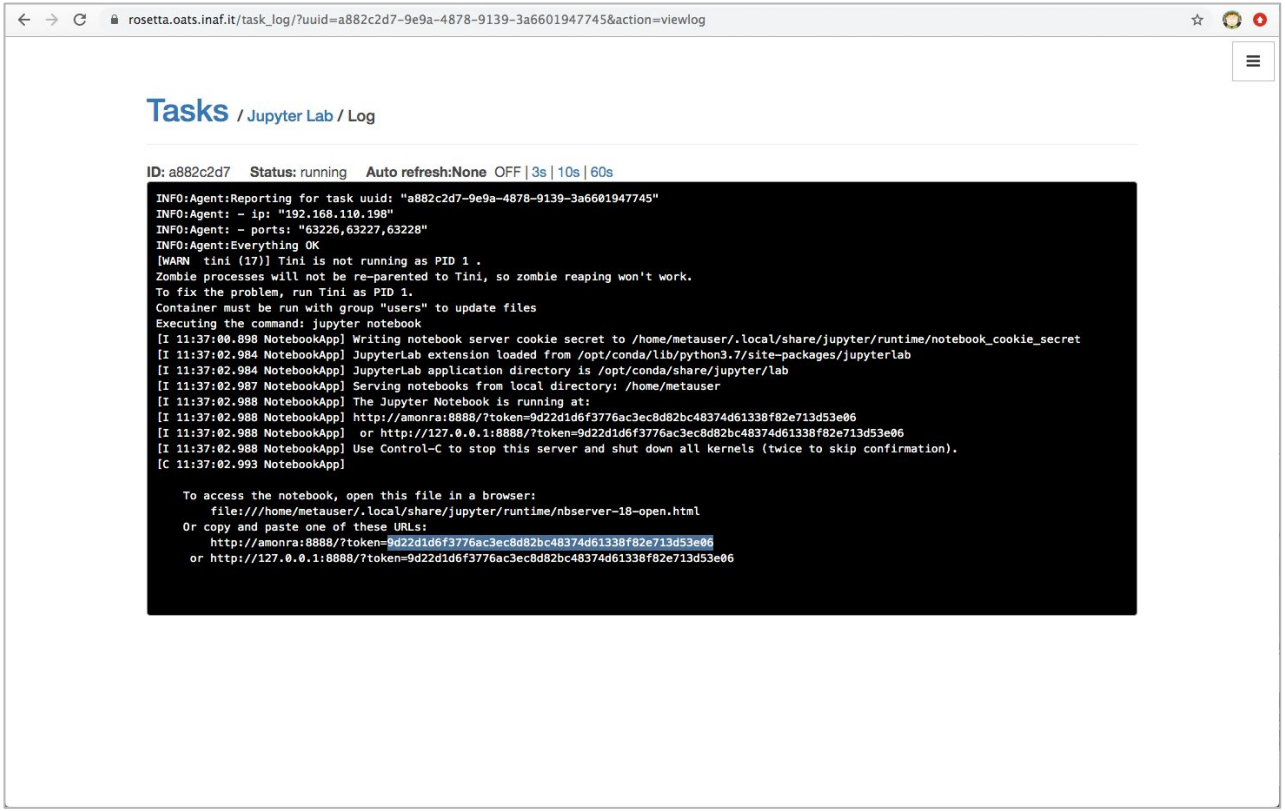

Rosetta quickstart & demo (for users)

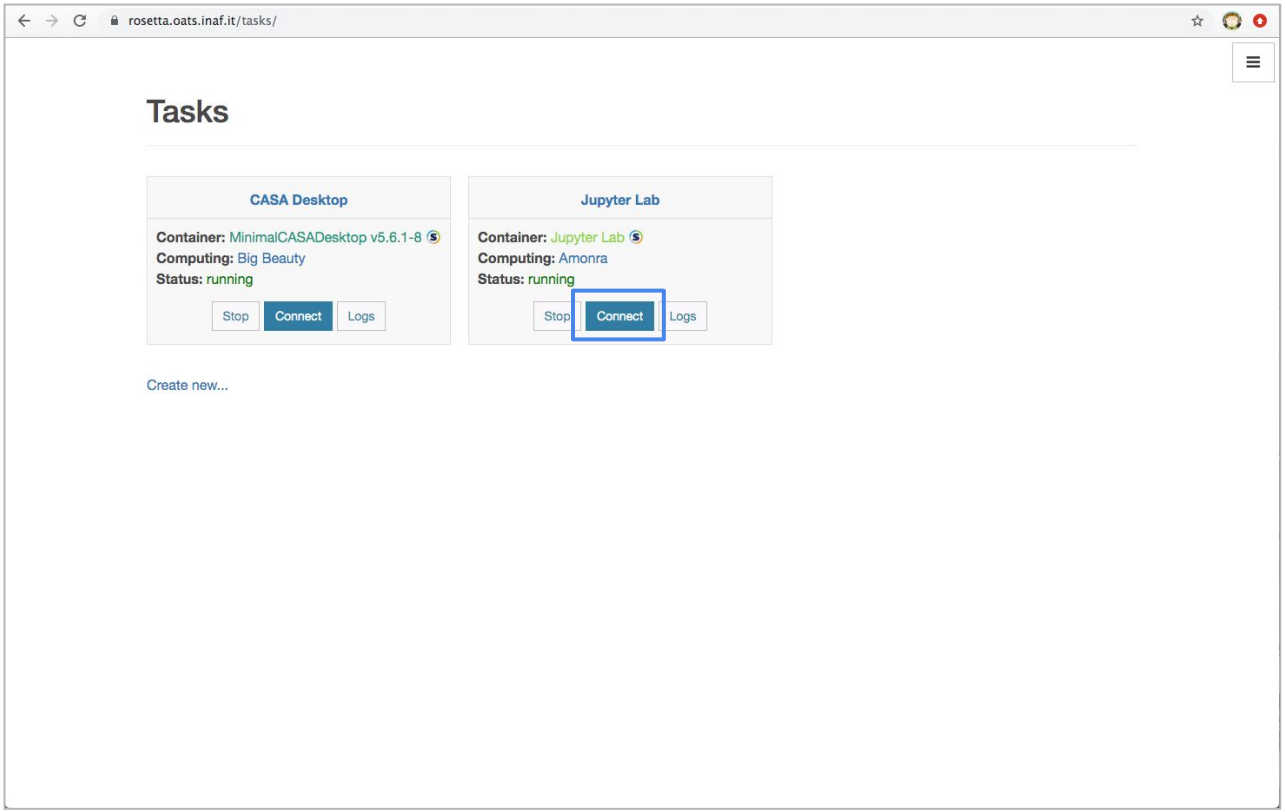

Rosetta quickstart & demo (for users)

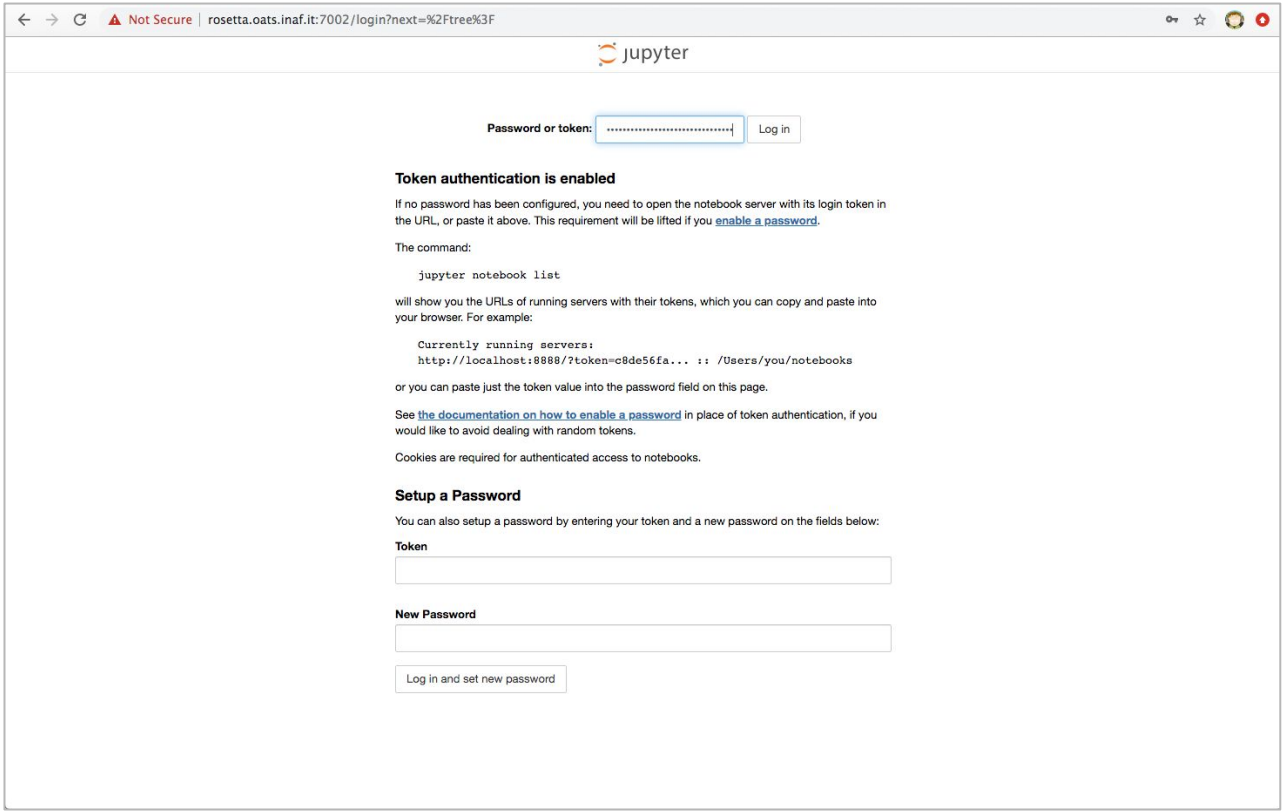

Rosetta quickstart & demo (for users)

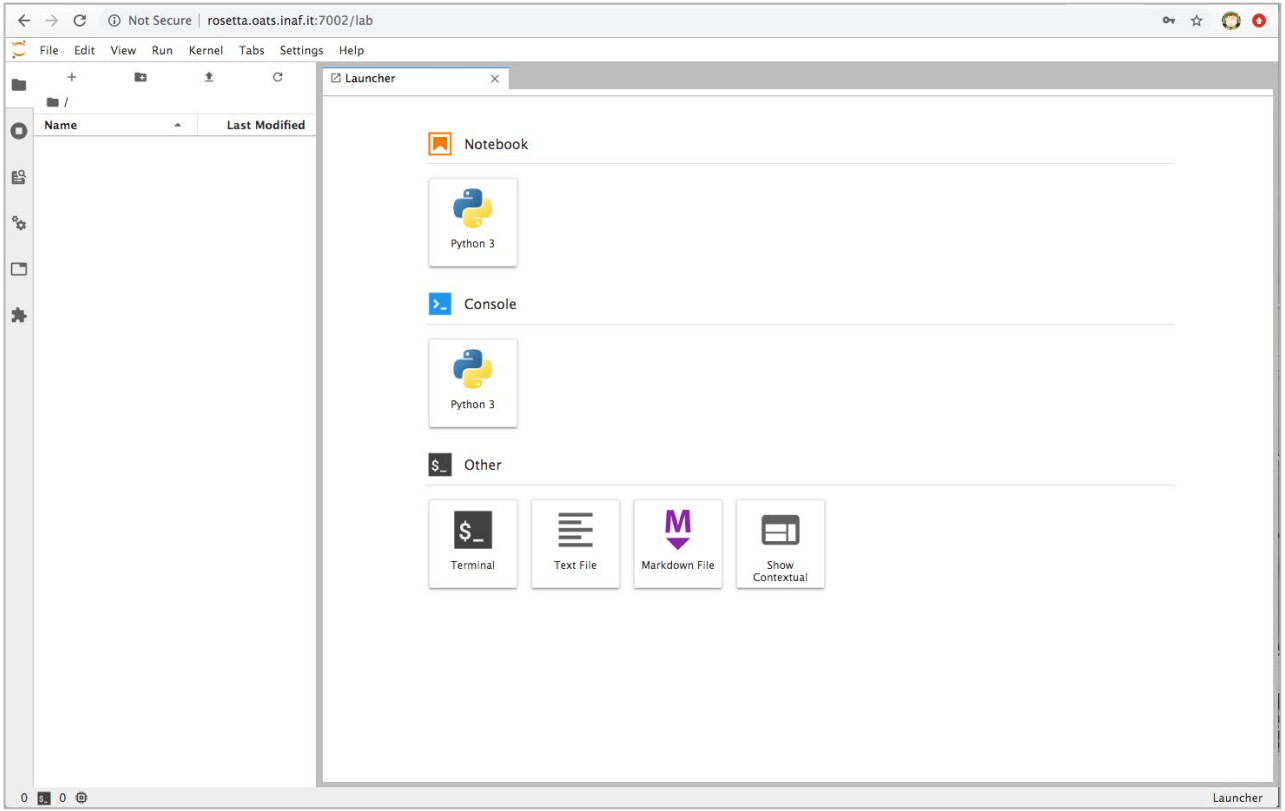

Rosetta Architecture

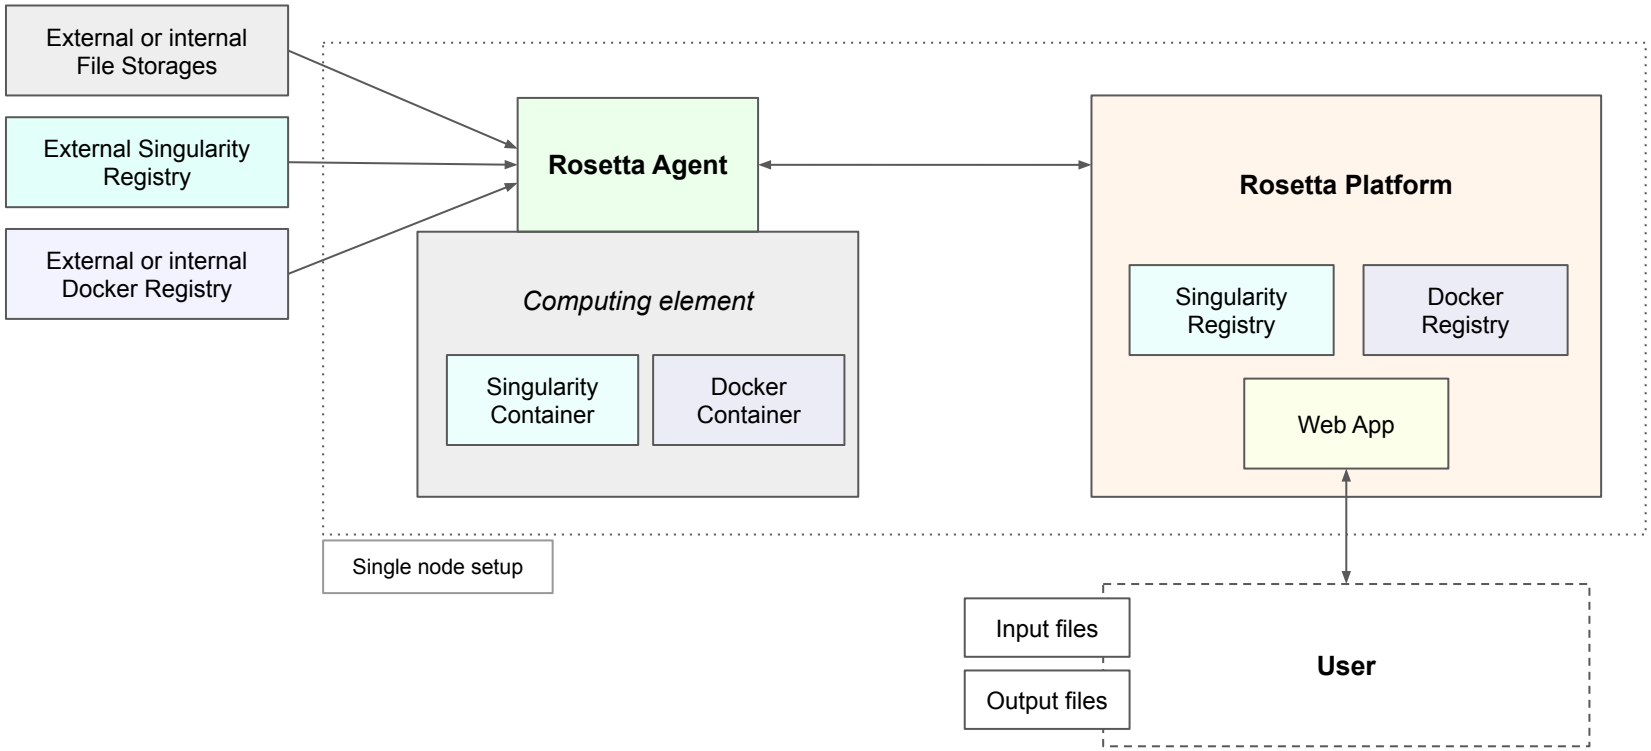

Rosetta technology stack

1. Docker + docker compose for orchestrating Rosetta itself
2. Python + Django for the orchestration
3. HTML for the Webapp
4. REST APIs (in development)
5. Postgres (or SQLite) database

Rosetta quickstart (for devs)

Quickstart

Requirements:

```
Bash, Git and Docker. Runs on Linux, Mac or Windows*.
```

*Windows not supported in development mode due to lack of support for symbolic links.

Setup

```
$ rosetta/setup
```

Build

```
$ rosetta/build
```

Run

```
$ rosetta/run
```

Play

```
You can now point your browser to http://localhost:8080.
```

Rosetta status

- Deployed on rosetta.inaf.oats.it, invite-based only for power-users beta tester
- A few user from other institutes using it autonomously on their own laptops for its orchestration capabilities

  → (i.e. to easily access remote desktops when locked out from offices in the pandemic.)
- Supports Singularity containers with preliminary support for Docker.
- Ships “batteries included”, with a few containers preconfigured and a test Slurm cluster instance up & running

That’s it :)

Questions?

Stefano Alberto Russo

stefano.russo@inaf.it

Rosetta repository: https://github.com/sarusso/Rosetta