



**FENIX**  
RESEARCH INFRASTRUCTURE

## Interactive computing services in Fenix

**May 10th, 2023**

**Fabio Pitari (CINECA)**

[f.pitari@cineca.it](mailto:f.pitari@ Cineca.it)



The ICEI project has received funding from the European Union's Horizon 2020 research and innovation programme under the grant agreement No 800858.

# Overview

- What's interactive computing (IAC)
- ICEI interactive computing implementation (Cineca)
- Additional interactive computing services in Fenix sites:
  - IAC at CSCS
  - IAC at JSC
  - IAC at BSC
  - IAC at CEA

# Traditional approach to HPC resources

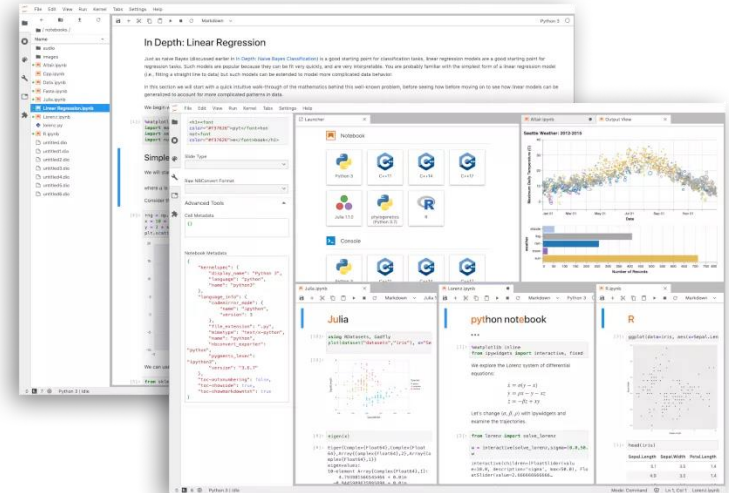
- Command line interface on remote machine
- Scheduler for jobs running the core of the computation
  - Queue-based system
  - Output is obtained at the end of the job
- Advantages:
  - Optimal sharing of resources among users
  - Large jobs can be scheduled
- Disadvantages:
  - Schedule time unknown, thus job results are usually checked at the end of the workflow
  - Visualization of results can't be obtained efficiently via command line
  - Command line is not user-friendly

# Interactive computing approach

- Interactive Computing service is an alternative approach to the traditional access to HPC resources
- Main idea: interaction on the fly with the workflow during its execution
  - Resources are allocated to the user despite the code is running or not (near-immediate access in some implementations)
  - User can employ those resources on the fly (and not queueing them)
- Advantages:
  - User can easily check intermediate results and change the workflow accordingly
- Disadvantages:
  - Not the optimal efficiency for resources occupancy (idle resources might occur)

# Jupyter

- The most common approach for interactive computing is via Jupyter (<https://jupyter.org/>) which grants additional features to the advantages already listed above:
  - Divide workflow into cells to be executed (=> execution can be driven on the fly)
  - Web interface (=> user-friendly, visualization tools)
  - Easily extensible (=> additional features)
- Other approaches can make use of ssh interactive sessions or remote desktop environments



# Current IAC services in Fenix

- One service developed in the context of the ICEI project:
  - Currently up and running at CINECA
  - Developed for all the Fenix sites in case they want to implement the same solution (Ansible playbooks for deployment)
  - Not publicly released yet (public pre-production expected before end of May 2023)
  - Based on Jupyter + Slurm Spawner, on Galileo 100 cluster
- Four pre-existing services in Fenix sites:
  - JSC: Jupyter + Unicore spawning on multiple clusters/cloud solutions
  - BSC: Jupyter via ssh, on HSM cluster
  - CSCS: Jupyter + Slurm Spawner, on Piz Daint cluster
  - CEA: NiceDCV remote desktop session via ssh
- **Please note: the only requirements for accessing to IAC services for each site is the membership in an active computing project at that site**



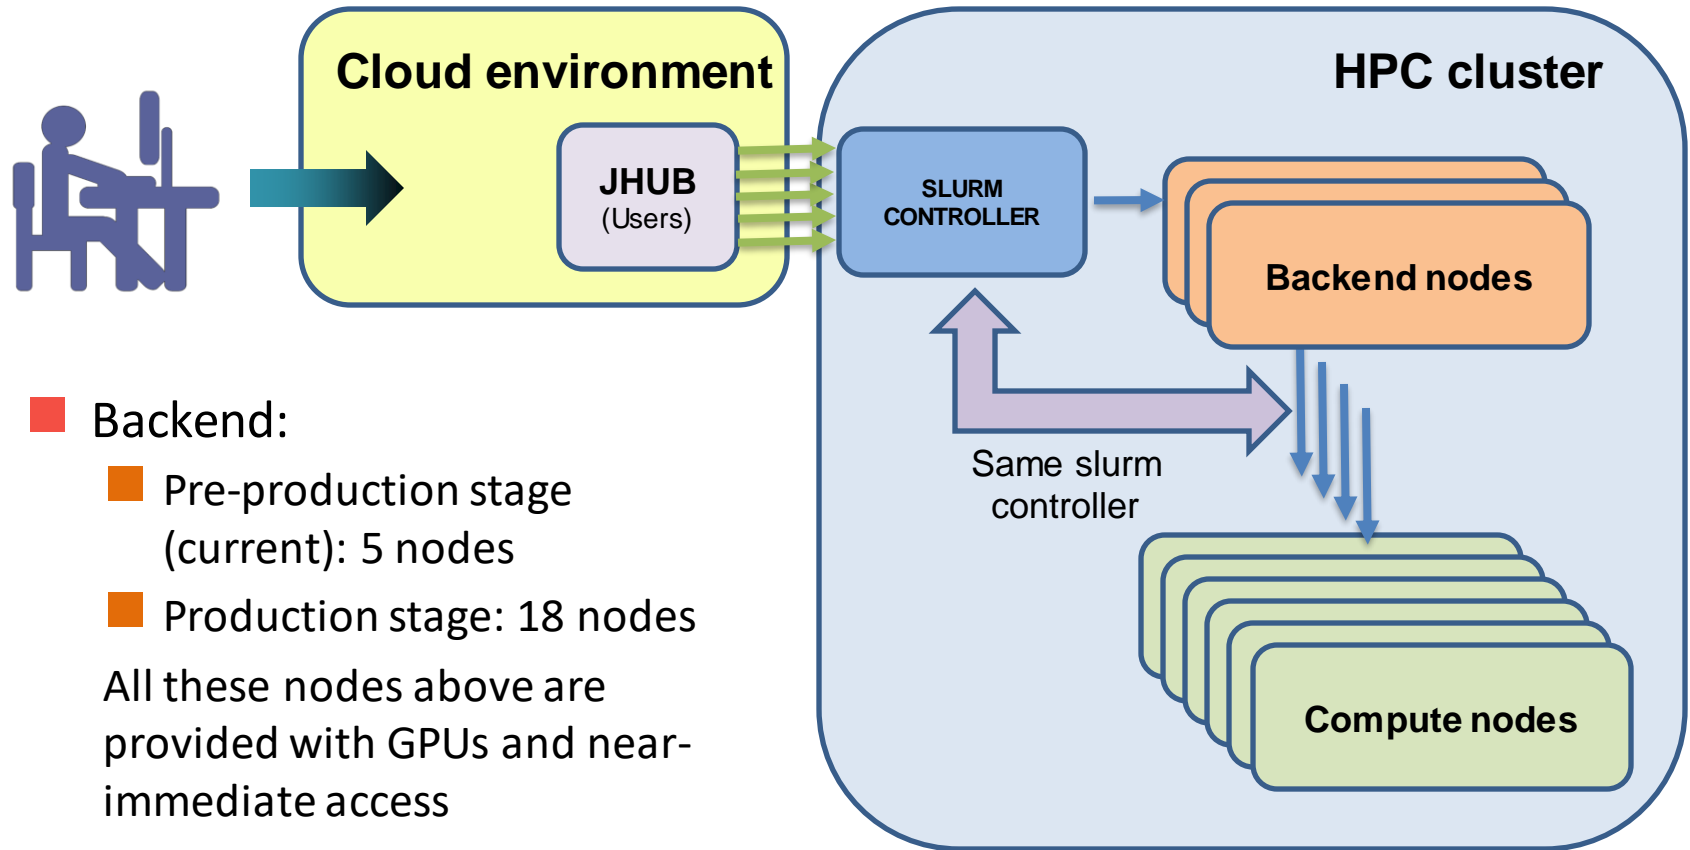
# **ICEI IAC IMPLEMENTATION**

# ICEI Interactive Computing implementation

- Current implementation at Cineca is deployed by E4 Computer Engineering (<https://www.e4company.com/>) thanks to ICEI project
- Deployment (via Ansible) is available for all the Fenix sites
- The framework involves three main components:
  - a frontend layer, based on Jupyterlab
  - a backend layer, running on the HPC cluster
  - Interaction between these two layers happens behind the scenes using Slurm scheduler



# Production infrastructure design



- Backend:
  - Pre-production stage (current): 5 nodes
  - Production stage: 18 nodesAll these nodes above are provided with GPUs and near-immediate access
- Compute:
  - 460 nodes without GPUS, batch jobs

# Frontend

## ■ Jupyterlab FrontEnd

- A (login) node with a public network interface (to expose the service)
- Can be hosted on a bare-metal server or virtual resource
- The latter has been chosen so far at Cineca, for several reason, e.g.:
  - Increased security
  - Being restricted, additional privileges can be given to E4 developers
  - Easy to export
  - Easy to expand
  - Different solutions can be tested, even during the production phase
- Easy to switch to bare-metal server if needed (being delivered via ansible playbooks)
- Interfaced with authentication services
- Resources can be easily extended with high availability and load balancing

# Frontend

- User access to a web service and is prompted with an authentication form
  - Note: this is compliant with cineca-hpc-idp (keycloak)
- Once authenticated, user is prompted to choose among a series of computational environment
  - Current Python-based envs: DASK, RAPIDS, Tensorflow, PyTorch, Life science packages
  - E4 solution integrates some other tools to the User interface
    - Visual Studio Code
    - Slurm batch interface for compute nodes
    - Remote desktop applications, I.e. VNC/Xpra (based on Jupyter remote desktop proxy (<https://github.com/jupyterhub/jupyter-remote-desktop-proxy>))

# Requirements

Login sessions for the users needs to be restored after the session is closed, i.e. session stays active as long as possible and can be resumed from the browser

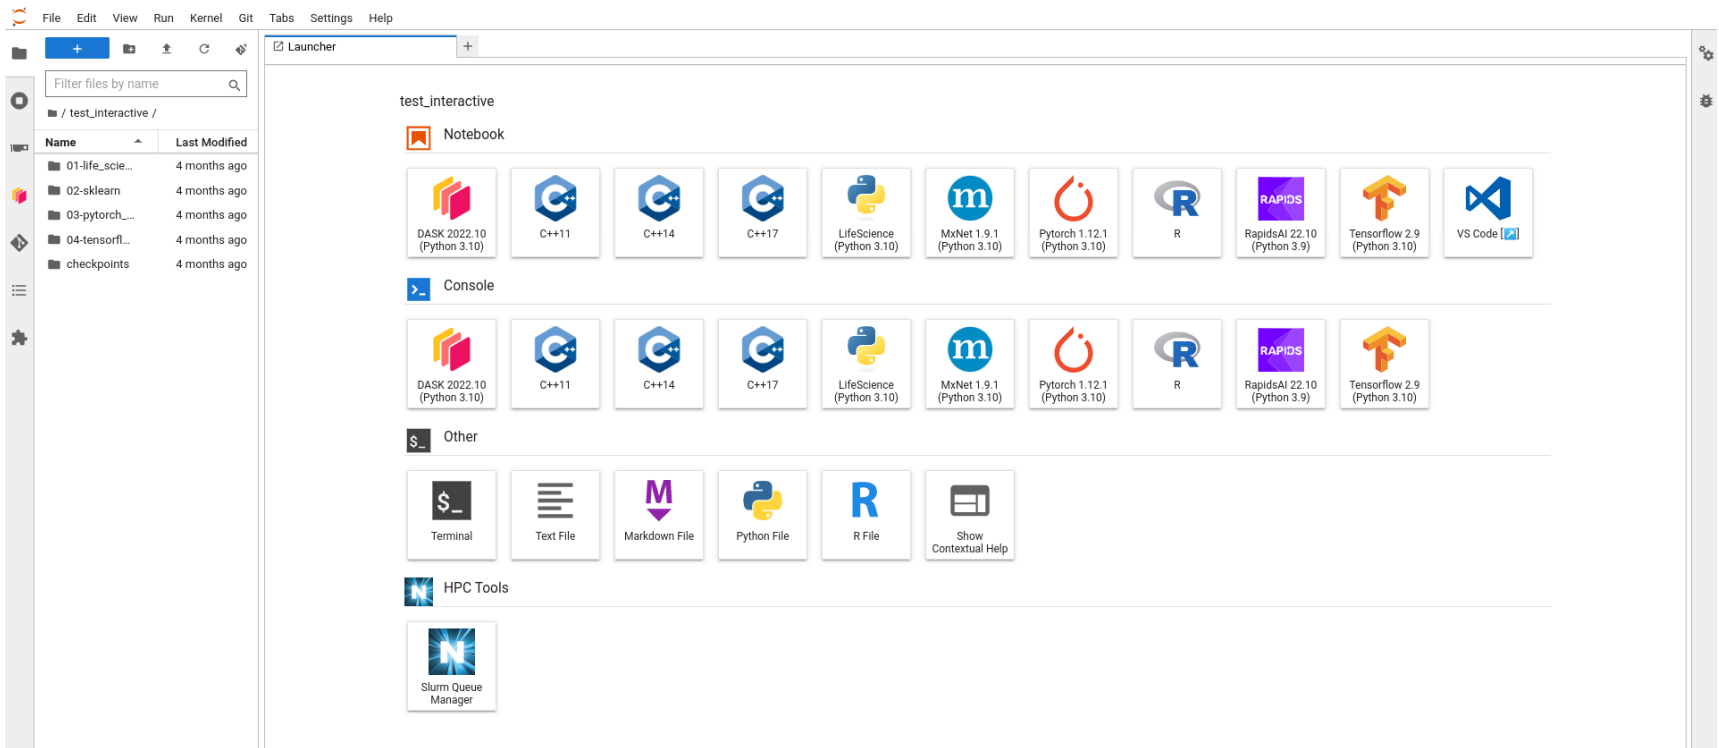


This is achieved with a slurm partition with oversubscription:

It allows to minimize queues for users' sessions (expected only in case of huge overload) and exploits the backend nodes at most

It is defined on the dedicated partition

# Frontend: layout



# Backend



Backend nodes: actual computing nodes drained from the cluster



Communication via frontend (VMs) and backend (bare metal) is achieved on a shared network

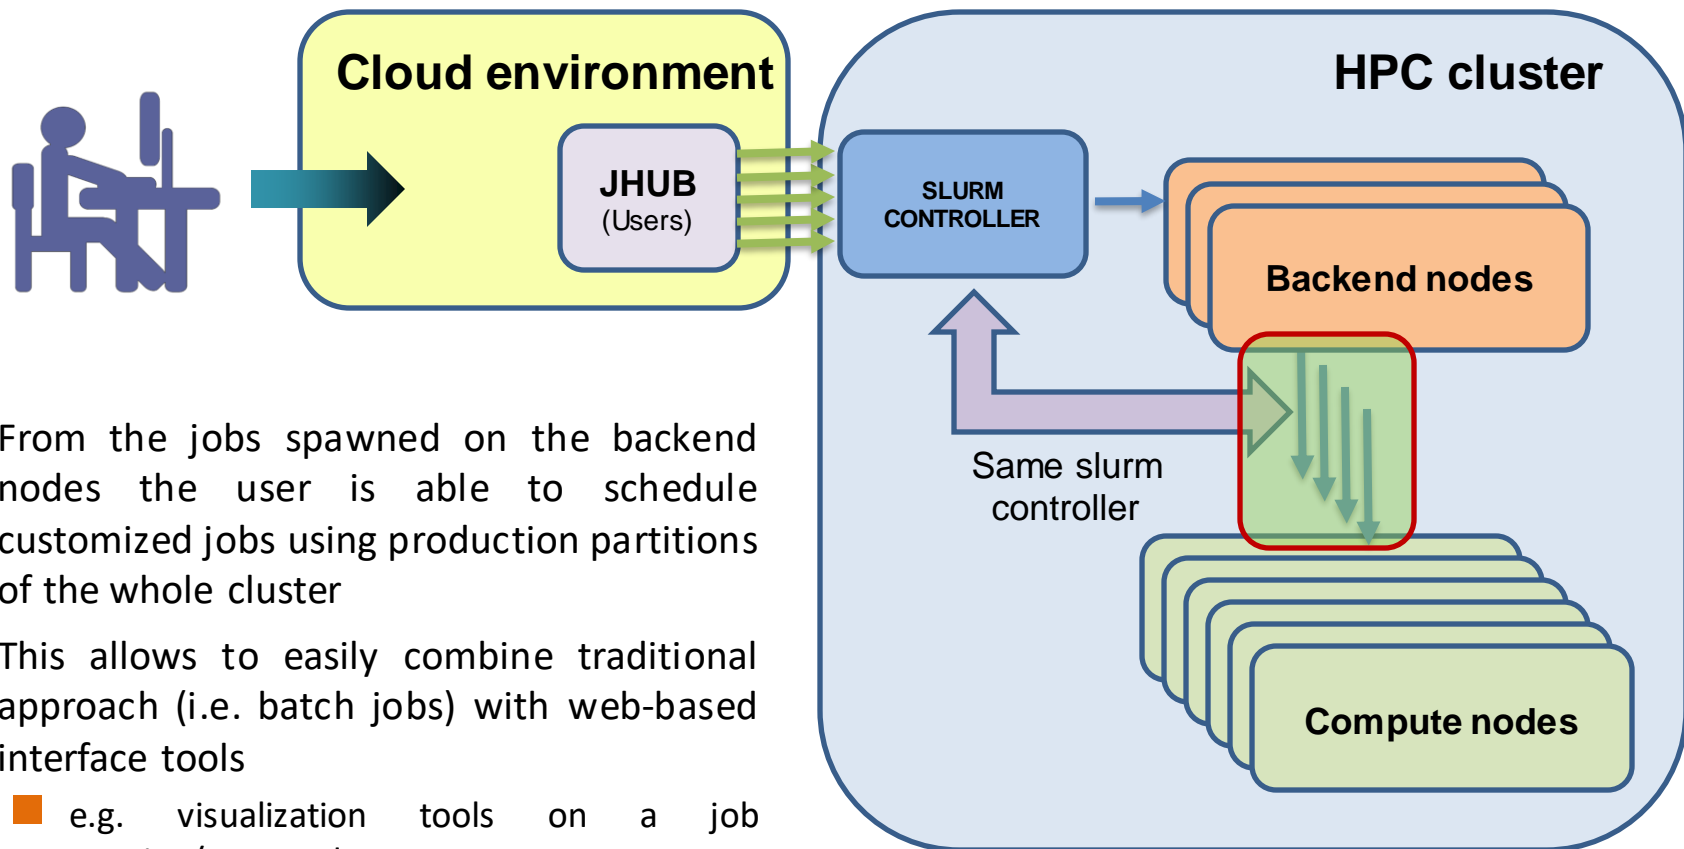


User sessions are handled on the backend nodes, whereas parallel tasks can be handled via compute nodes through a dedicated web interface



ssh and additional frontends for admin purpose can be given via different VMs

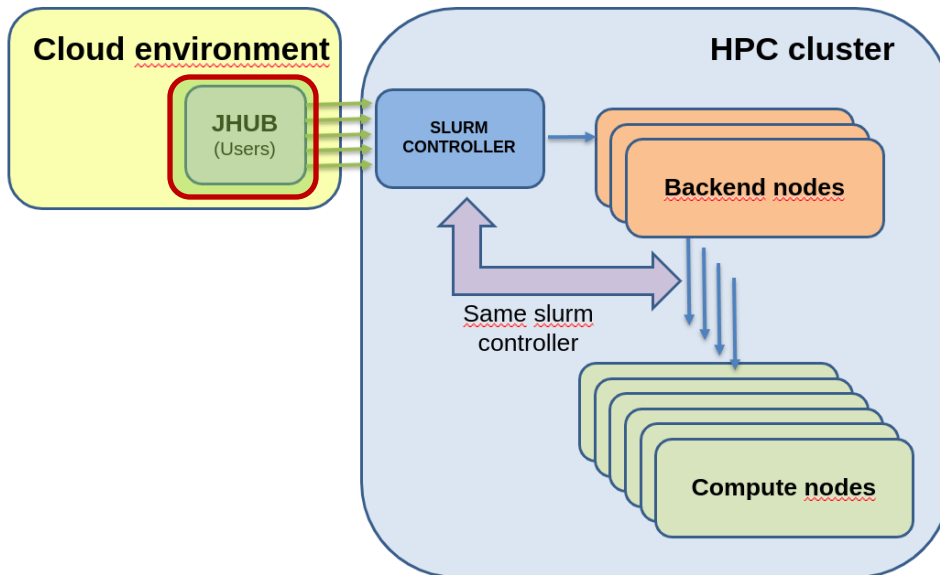
# Interaction with compute nodes



- From the jobs spawned on the backend nodes the user is able to schedule customized jobs using production partitions of the whole cluster
- This allows to easily combine traditional approach (i.e. batch jobs) with web-based interface tools
  - e.g. visualization tools on a job running/executed

# Allocation sequence

- Login interface can be connected with idp authentication (e.g. keycloak)
- Allowed to every user of the cluster, or restricted to a specific subset

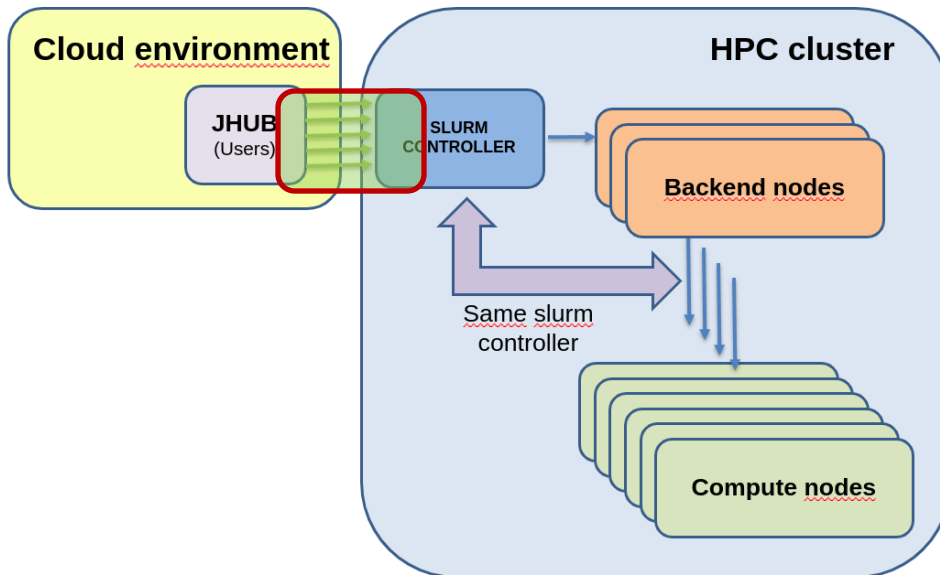


The screenshot shows the ICE4HPC login interface. At the top, there is a logo consisting of a green hexagon with a white 'Y' shape inside. Below the logo, the text "ICE4HPC" is displayed in a large, bold, teal font. Underneath, it says "INTERACTIVE COMPUTING ENVIRONMENT FOR HPC CLUSTERS" in a smaller, teal font. At the bottom of this section, it says "BY E4 ANALYTICS" in a small, teal font. Below this is a dark teal header with the text "ICE4HPC login" in white. The main area contains a "Username:" label followed by a text input field containing "fpitari1". Below that is a "Password:" label followed by a password input field with masked characters. At the bottom, there is a dark teal "Sign In" button.



# Allocation sequence

- A form is displayed in order to let the user request resources
- A table is displayed parsing information about available resources from the scheduler
- Session can be restored when the browser is closed, thanks to a job dispatched in oversubscription on the backend nodes by the slurm controller in the cloud environment



Interactive session parameters setting

Slurm Reservation:  Slurm Account:

Number of cores:  Memory (MB):

GPU configuration:  Time (hours):

ICE4HPC Backend Environment:

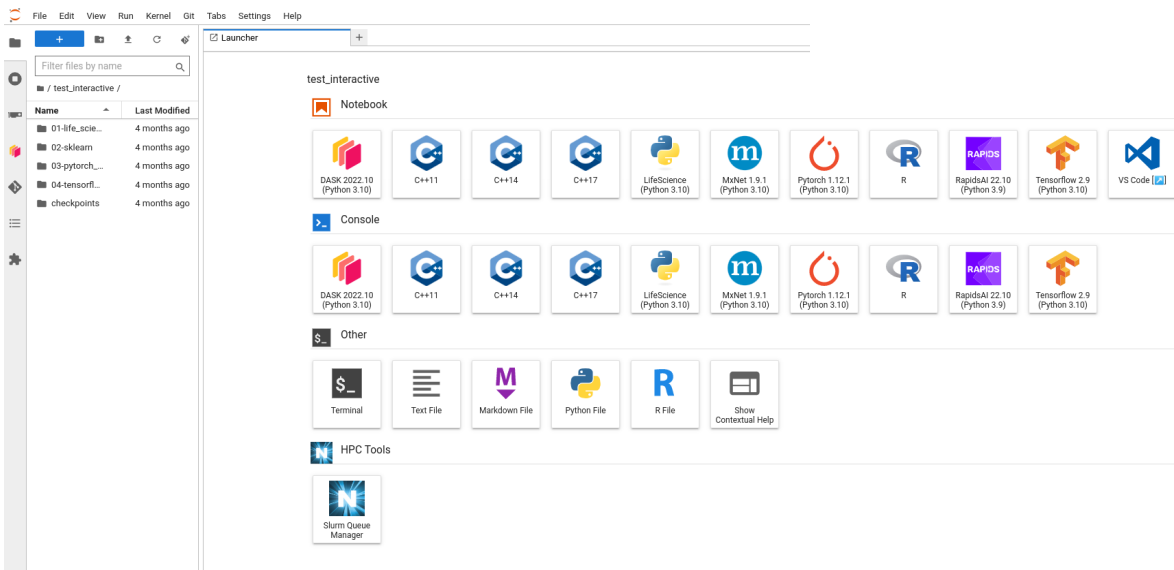
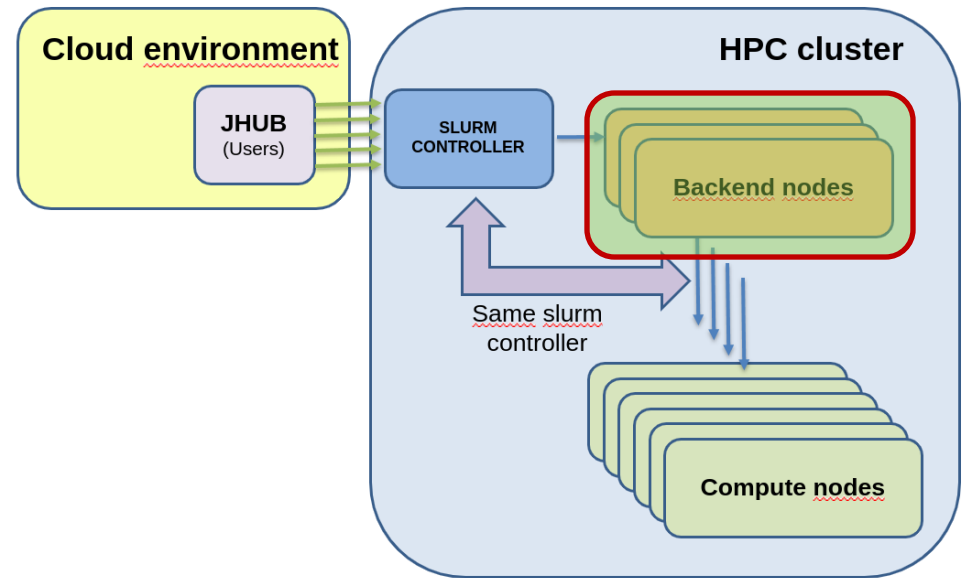
User interface:

AVAILABLE COMPUTATIONAL RESOURCES

	Nodes	Total Number of cores	Total Free Memory (MB)
No available GPUs	0	/	/
1 available GPU	0	/	/
2 available GPUs	5	228	1.904.996

# Allocation sequence

- Oversubscription is currently set to 5 jobs per core
- User session (i.e. backend job) expires currently in only 8 hours for the pre-production phase (to be increased to 36 hours in production)

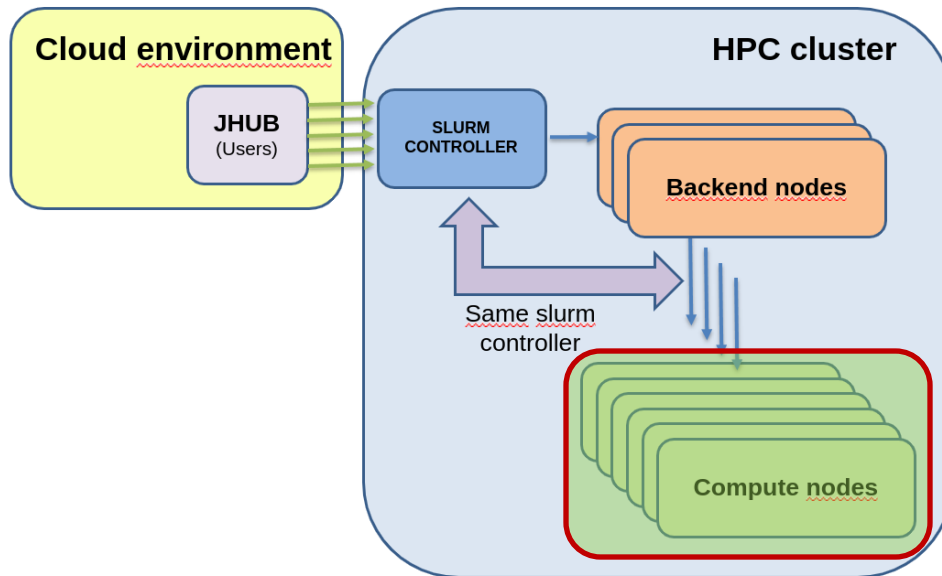


# Additional features

- Batch jobs can be submitted to the cluster via a dedicated interface
- Slurm queues and notifications can be also monitored
- Proxy for additional web interfaces from the launcher (currently just VSCode; R-Studio and VNC incoming)

```

1 #include <stdio.h>
2
3 int bsearch (int *a, int n, int x) {
4     int i = 0, j = n - 1;
5     while (i <= j) {
6         int k = i + (j - i) / 2;
7         if (a[k] == x) {
8             return k;
9         }
10        else if (a[k] < x) {
11            i = k + 1;
12        }
13        else {
14            j = k - 1;
15        }
16    }
17    return -1;
18 }
19
20 int bsearch_r (int *a, int x, int i, int j) {
21     if (j < i) {
22         return -1;
23     }
24     int k = i + ((j - i) / 2);
25     if (a[k] == x) {
26         return k;
27     }
28     else if (a[k] < x) {
29         return bsearch_r(a, x, k + 1, j);
30     }
31     else {
32         return bsearch_r(a, x, i, k - 1);
33     }
34 }
35
36 int main () {
37     int a[] = {-31, 0, 1, 2, 2, 4, 65, 83, 99, 782};
38     int n = sizeof a / sizeof a[0];
  
```



JOBID	PARTITION	NAME	USER	ST	TIME	NODES	NOELIST(REASON)
8602647	g100_usr_	jupyterh	fpitar1	R	0:18	1	r505r017

The interface shows a form for submitting a batch job. It includes a 'Submit a File' button and a 'Submit Text' button. Below these, there is a text input field for the path to the slurm script, with the value 'test\_interactive/03-pytorch\_GPU/test.png' entered. A 'Submit Job' button is located at the bottom of the form.

# Current status

- Security assessment concluded, waiting for response
- Expected to be publicly available before the end of the month (May 2023) at the following link:

<https://jupyter.g100.cineca.it/>

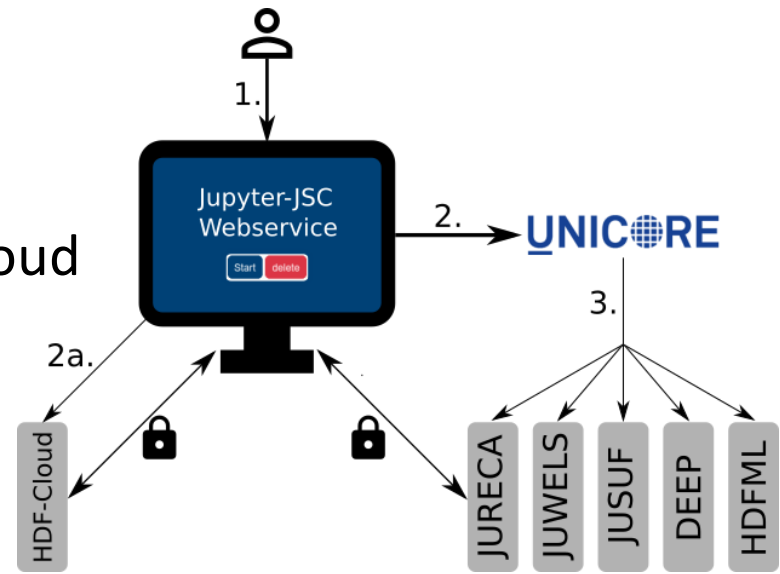
- Membership in an active computing project at CINECA will be required
- Further extensions in the future (Julia, VNC, R-Studio, Octave...)



# **IAC IMPLEMENTATIONS AT FENIX SITES**

# Interactive computing at JSC

- Based on Jupyterlab
- Frontend runs on the cloud (HDF-cloud at JSC, deployed via Kubernetes)
- Main components:
  - Jupyterhub
  - Spawner:
    - via Unicare for multiple clusters
    - Via UserLab Manager for cloud-systems (starts a kubernetes pod)
  - Ad-hoc tunneling to compute nodes
- Access from the following link:



<https://jupyter-jsc.fz-juelich.de/>

Membership in an active computing project at JSC is required

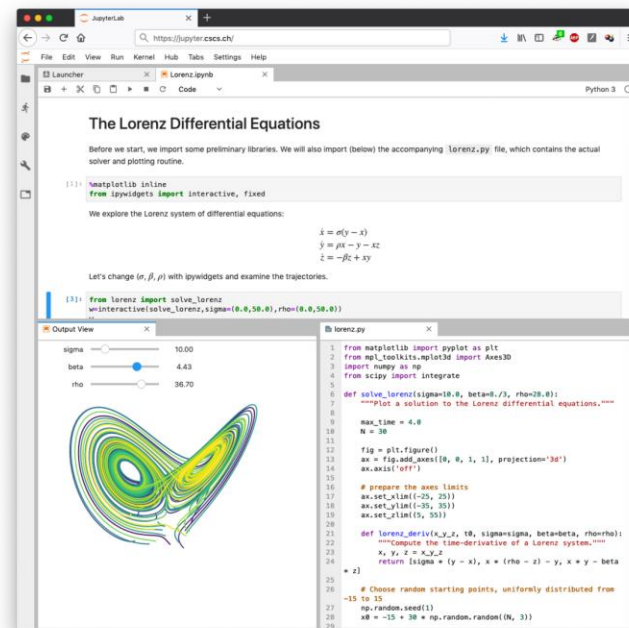
# Interactive computing at CSCS

- Similar to Cineca implementation (Jupyterlab + Slurm spawner)
- Cluster used is Piz Daint
- Job is redirected to different queues depending on the amount of requested resources
- Cli tools to ease the customization of environments
- Access from the following link:

<https://jupyter.cscs.ch/>

(Documentation here: <https://user.cscs.ch/tools/interactive/jupyterlab/>)

Membership in an active computing project at CSCS is required



# Interactive computing at BSC

- Interactive sessions are driven via ssh
- Cluster employed is HSM
- Jupyterhub can be initialized via ssh on a compute node
  
- More information here:

<https://www.bsc.es/supportkc/docs/HSMCompute/intro/>

Membership in an active computing project at BSC is required.  
You need also to communicate a public ip to be whitelisted



# Interactive computing at CEA

- Based on NiceDCV to access a Remote Desktop (Gnome environment)
- Running on Irene cluster
- Access in several steps
  - Access a TGCC machine via ssh
  - Initialize a NiceDCV session via specific command line script
  - An url is generated; the user can connect to the url from his/her browser (new authentication is required)
- More information here:

[https://www-hpc.cea.fr/tgcc-public/en/html/toc/fulldoc/Interactive\\_access.html](https://www-hpc.cea.fr/tgcc-public/en/html/toc/fulldoc/Interactive_access.html)

Membership in an active computing project at CEA is required



**Thanks for listening!**

Aknowledgements: Anne Nahm, Valentina Armuzza, Björn Hagemeyer (JSC), Luciana Marques, Javier Bartolome (BSC), Alex Upton (CSCS), Thomas Leibovici (CEA), Mirko Cestari, Andrea Pieretti, Alessandro Federico (CINECA), Mario Rosati (E4)