



HPC JupyterHub

Short Introduction to the HPC JupyterHub at the RWTH

Alvaro Frank | 05.02.2024

[hpcjupyterhub](https://www.hpcjupyterhub.com)



RWTHAACHEN
UNIVERSITY

NHR4
CES

Topics

Covered:

- What is [hpcjupyterhub](#)
- What hardware it offers
- How to connect
- How to use the JupyterHub
- How to access and share files
- How to obtain help

Not covered:

- How to code (Python, C++, Notebooks)
- How to create an account
- How to get compute time
- How to fully use Slurm
- How to create user profiles (future feature)

Jupyter

- **JupyterHub**

- Web page login
- Configure software and hardware (Configure Job)
- Starts JupyterLabs (Submit, Wait, Start Job)

- **JupyterLab**

- Run Jupyter notebooks (Job runs)
- Run consoles
 - Run linux commands, code, etc
- Browse files

- **Jupyter Notebooks**

- Source code, plots, etc.
- Uses Jupyter **Kernels** (python, R, C++, etc).

Hardware

- **Hardware:**
 - **Claix 2018**
 - ~500x(24c,128GB)
 - ~9x(...,2x Pascal GPUs)
 - ~1200x(48c,192GB)
 - ~54x(...,2x Volta GPUs)
 - **Claix 2023 (soon)**
- **Software:**
 - **System Provided Profiles**
 - pip install --user
 - **User defined Profiles (soon):**
 - pip + conda + modules
- **Partitions:**
 - **CLAIX-2018**
 - c18g - GPUs
 - c18m - CPUs
 - Special Partitions
 - **CLAIX-2023 (soon)**
 - c23g? - GPUs
 - c23m? - CPUs
 - c23i? - GPUs

Overview


Interactive access

- > VPN, Office, Network
 - > Browser
 - > Login **JupyterHub**
 - > Choose Profile
 - > SLURM Scheduler: Start **JupyterLabs**
 - > Code **Jupyter Notebooks**
 - > Run Console
 - > Run Kernels (python, R, c++, etc).


Limitations:

- No remote VisualStudio
- No Token remote access

Infrastructure



HPC JupyterHub: starts JupyterLab(s)



Please note that 2FA/MFA is mandatory now to login to the HPC system, including the HPC JupyterHub.

JupyterHub

HPC JupyterHub

HPC Username:

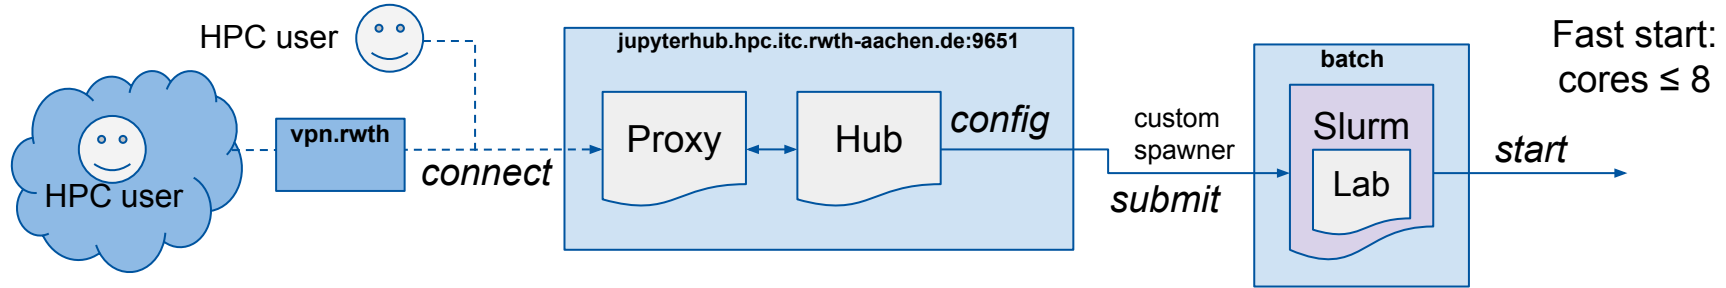
HPC Password:

2FA Code:

[Sign in](#)

Need help? Read our [Documentation](#).
Login problems? Enable your [Account](#).
Startup issues? See [Maintenance Log](#).
Further help? Contact the [ServiceDesk](#).
Please review our [Privacy Policy](#)

Infrastructure



Slurm Batch Scripts:

[ITC - SchedMD](#)

Slurm Commands:

[ITC - SchedMD](#)

HPC JupyterHub: starts JupyterLab(s)

JupyterHub

JupyterLab Server: Main

JupyterLab Profile: ?

Enable User PKGs? (No support!) ?

Keras + TensorFlow in CUDA within conda for GPUs, native on HPC compute nodes.
Kernels: Python 3.9 + Keras + TensorFlow

Billing Project: ?

Partitions: ?

Available core-hours: 5.48k ?

Advanced Settings: ?

Simple:

Max duration: ?

Number of Cores: ?

Node Memory: ?

Number of GPUs: ?

?

Starting a Server:

Server: Main

Your server is starting up.

You will be redirected automatically to your compute node with your JupyterLab when it's ready.

You can cancel/stop your submission before it starts [Cancel](#)

SLURM Job id 32356105 PENDING . Expected start: 2023-01-24 @ 16:11:52

Event log - click to hide -

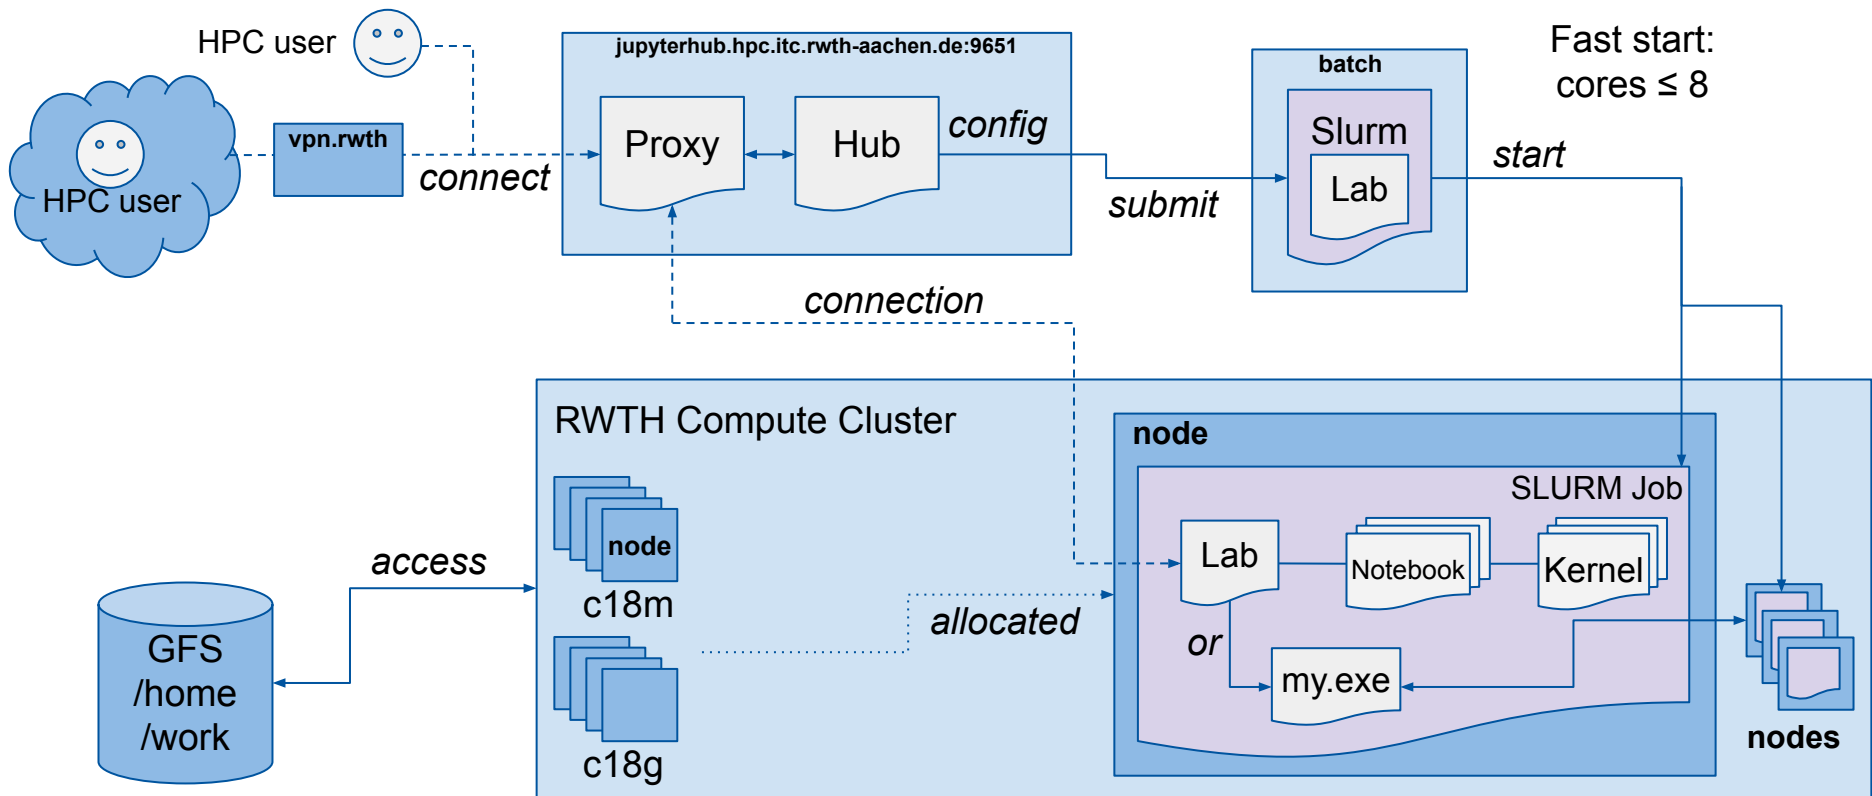
Server requested

SLURM Job id 32356105 PENDING . Expected start: N/A

SLURM Job id 32356105 PENDING . Expected start: 2023-01-24 @ 16:11:52

JupyterHub

Infrastructure



JupyterLab and Notebooks

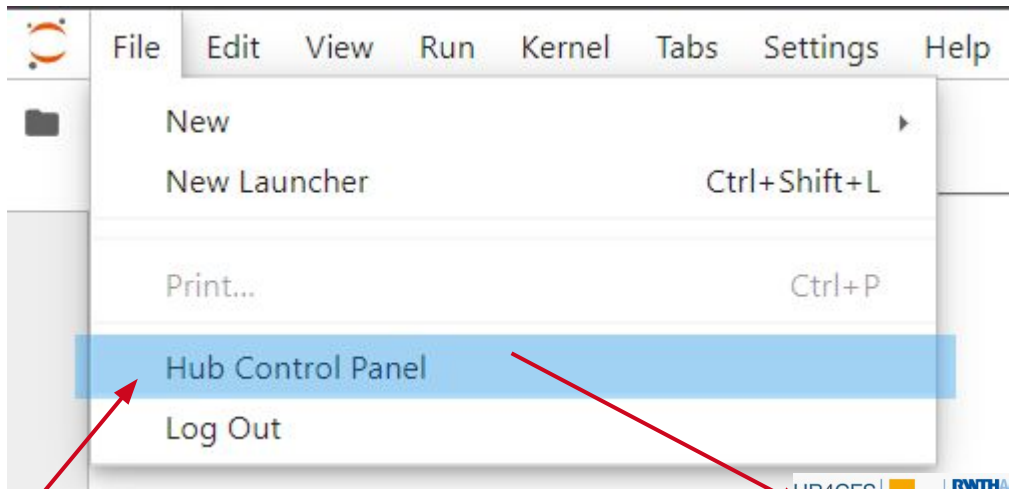
The screenshot displays the JupyterLab interface. On the left, the file browser shows a directory structure with a file named 'ExampleNotebook.ipynb' selected. A red arrow points from this file to the notebook editor on the right. The notebook editor shows a code cell with the following Python code:

```
[2]: i=9
[4]: i+=1
[5]: print(i)
10
[ ]:
```

The status bar at the bottom indicates the kernel is 'Python 3 (ipykern...)' and the mode is 'Com...'. The text 'JupyterLab' is visible on the left side of the interface, and 'Jupyter Notebook' is visible on the right side of the notebook editor.

DEMO

Access HUB



[Stop My Server](#) [My Server](#)

Named Servers

In addition to your default server, you may have additional server(s) with names. This allows you to have more than one server running at the same time.

Server name	URL	Last activity	Actions
<input type="text" value="Name your server"/>	Add New Server		
SecondServer	 SecondServer	a few seconds ago	stop
ThirdServer		Never	start delete

Important!

- Closing the browser window/tab will eventually close the running lab:
 - Labs do not run in the background!
- You can run multiple labs at the same time
- Data is shared through GIT/GFS
- Use !<cmd> to run shell commands
- Python is inherently single-threaded
 - joblib, multiprocessing, threading and asyncio, external MPI
- Shell dir changes don't apply to next cmd!

**Thank you for your attention.
Any questions?**