

# Jupyterhub at NHR@FAU

Johannes Veh

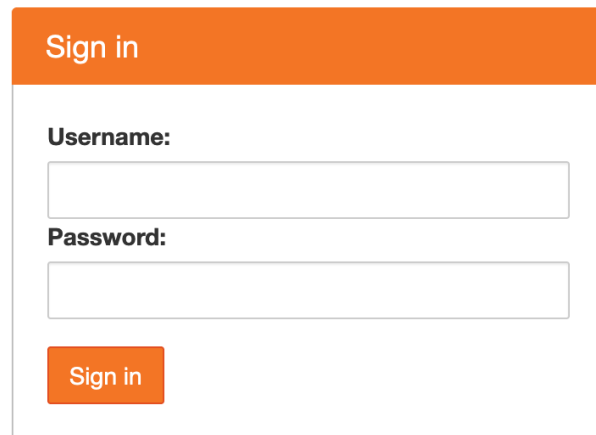
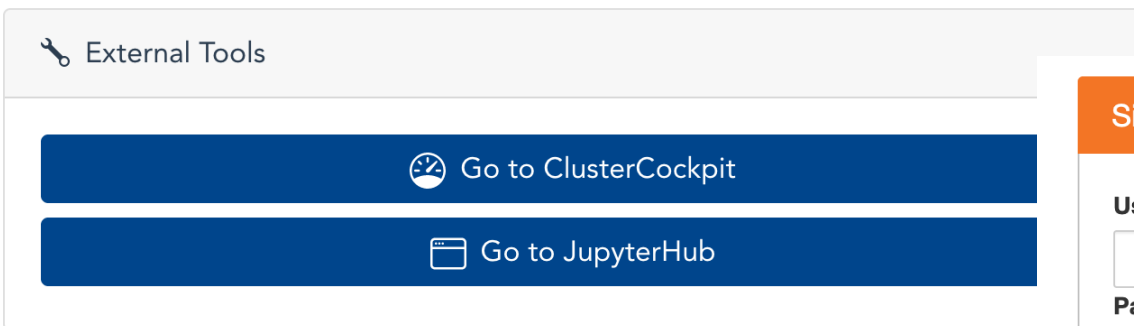
Erlangen National High Performance Computing Center (NHR@FAU)

NHR@FAU HPC Café, 2023-07-11



# Access

- You can find the Hub at: <https://hub.hpc.fau.de/jupyter/hub>
  - legacy HPC account (with a password set through [idm.fau.de](https://idm.fau.de))
    - VPN / university network only
  - new HPC account use the Jupyterhub button provided at [portal.hpc.fau.de](https://portal.hpc.fau.de)



A screenshot of a "Sign in" form. The form has an orange header with the text "Sign in". Below the header are two input fields. The first field is labeled "Username:" and the second field is labeled "Password:". Below the input fields is an orange button with the text "Sign in".

# Starting a Server

wait quite long (some or even many hours) for your job to start. There is also no way to request specific node types.

Local on jupyterhub (systemd) - 2 cores, 4 GB, unlimited  
Local on jupyterhub (systemd) with datascience-notebook\_r-4.2.3 Singularity container - 2 cores, 4 GB, unlimited  
1x GTX1080Ti, 4 hours  
2x GTX1080Ti, 3 hours  
4x GTX1080Ti, 2 hours  
1x GTX1080Ti, 4 hours, Singularity container gpu-jupyter-v1.5\_cuda-11.6\_ubuntu-20.04  
✓ TinyFAT - 1 Core (2 SMT threads) + 8 GB, 4 hours  
TinyFAT - 2 Core (4 SMT threads) + 16 GB, 4 hours  
TinyFAT - 4 Core (8 SMT threads) + 32 GB, 4 hours  
TinyFAT - 1 Core (2 SMT threads) + 8 GB, 4 hours, Singularity container Jupyter/Datascience-Notebook with R-4.2.3  
TinyFAT - 1 Core (2 SMT threads) + 8 GB, 4 hours, Singularity container xfce

The Slurm job profiles usually also do not need any of the advanced options.  
Thus, use with care!

Modules to preload

magic pre-set  ONLY if you got one e.g. for a lecture

Specific Slurm settings are usually not required. Wrong input will result in job-start failures!

--reservation=  ONLY if you got one

--account=  usually NOT required

--chdir=  \$HOME by default

Start

## Current Utilization

updated every 2 minutes

f 2023-07-07 @ 10:40

GPU resources (mainly used for Jupyterhub) =

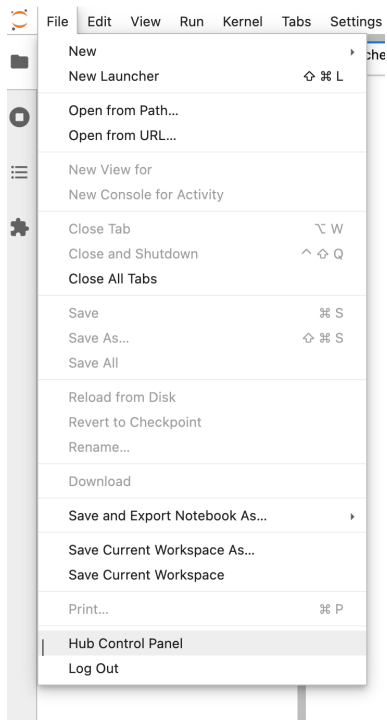
LIST	STATE
node-6	idle
node-7	idle
node-8	idle
node-9	idle

CPU-only resources (shared with TinyFAT) ===

ITION	AVAIL	NODES	STATE
*	up	3	alloc
*	up	33	idle

# Kill your Server

- Closing your browser tab will not end Slurm jobs in the background
  - Therefor kill your server



- Please be patient when pressing "Stop my server" as it will take some time to shutdown your Jupyter instance.
- Stopping is successfully started once the blue "My server" button no longer accepts clicks. Stopping is finished once the red "Stop my server" button is no longer visible.
- Servers which are waiting for the Slurm job to start, cannot be stopped; deleting the pending Slurm job on the cluster frontend manually may wait. After the timeout you can press "Stop my server" (again).
- New instances only can be started once the old one is fully stopped.

Resources are only freed if you press "Stop my server". Otherwise resources remain blocked even if you log out or close the browser tab!

# Adding a Kernel

- We use `nb_conda_kernels`
  - Load conda module on cluster

```
conda install -n <my-py> ipykernel          # python
conda install -n <my-R> r-irkernel         # R
```

Many more kernel possible:

<https://github.com/jupyter/jupyter/wiki/Jupyter-kernels>

(not every kernel was tested with our setup)

# Further Documentation

---

- [NHR-FAQ](#)
- [Cluster Overview](#)
- [Jupyter and Python](#)

Thank you



Bundesministerium  
für Bildung  
und Forschung

