

# Helmholtz GPU Hackathon 2023 User Guide

Andreas Herten <a.herten@fz-juelich.de>

## Abstract

Some hints for participants and mentors of the GPU Hackathon 2023 at Jülich Supercomputing Centre.  
*Hack Away!*

## Contents

<b>1</b>	<b>Helmholtz GPU Hackathon 2024</b>	<b>1</b>
1.1	Sign-Up . . . . .	1
1.2	HPC Systems . . . . .	1
1.3	Access . . . . .	2
1.4	Environment . . . . .	2
1.5	Containers . . . . .	2
1.6	Batch System . . . . .	2
1.7	Etc . . . . .	3
1.7.1	Previous Documentation . . . . .	3
1.7.2	PDFs . . . . .	3

## 1 Helmholtz GPU Hackathon 2024

This repository holds the documentation for the Helmholtz GPU Hackathon 2024 at CASUS Görlitz.

For additional info, please write #cluster-support on Slack.

### 1.1 Sign-Up

Please use JuDoor to sign up for our training project, training2406: <https://judoor.fz-juelich.de/projects/join/training2406>

Make sure to accept the usage agreement for JURECA-DC and JUWELS Booster.

Please upload your SSH key to the system via JuDoor. The key needs to be restricted to accept accesses only from a specific source, as specified through the `from` clause. Please have a look at the associated documentation ([SSH Access](#) and [Key Upload](#)).

### 1.2 HPC Systems

We are using primarily JURECA-DC for the Hackathon, a system with 768 NVIDIA A100 GPUs.

For the system documentation, see the following websites:

- [JURECA-DC](#)
- [JUWELS Booster](#)

## 1.3 Access

After successfully uploading your key through JuDoor, you should be able to access JURECA-DC via

```
ssh user1@jureca.fz-juelich.de
```

The hostname for JUWELS Booster is `juwels-booster.fz-juelich.de`.

An alternative way of access the systems is through *Jupyter JSC*, JSC's Jupyter-based web portal available at <https://jupyter-jsc.fz-juelich.de>. Sessions should generally be launched on the login nodes. A great alternative to X is available through the portal called Xpra. It's great to run the Nsight tools!

## 1.4 Environment

On the systems, different directories are accessible to you. To set environment variables according to a project, call the following snippet after logging in:

```
jutil env activate -p training2406 -A training2406
```

This will, for example, make the directory `$PROJECT` available to use, which you can use to store data. Your `$HOME` will not be a good place for data storage, as it is severely limited! Use `$PROJECT` (or `$SCRATCH`, see documentation on [Available File Systems](#)).

Different software can be loaded to the environment via environment modules, via the `module` command. To see available compilers (the first level of a toolchain), type `module avail`.

The most relevant modules are

- Compiler: GCC (with additional CUDA), NVHPC
- MPI: ParaStationMPI, OpenMPI (make sure to have loaded MPI-settings/CUDA as well)

## 1.5 Containers

JSC supports containers thorough Apptainer (previously: Singularity) on the HPC systems. The details are covered in a [dedicated article in the systems documetnation](#). Access is subject to accepting a dedicated license agreement (because of special treatment regarding support) on JuDoor.

Once access is granted (check your groups), Docker containers can be imported and executed similarly to the following example:

```
$ apptainer pull tf.sif docker://nvcv.io/nvidia/tensorflow:20.12-tf1-py3
$ srun -n 1 --pty apptainer exec --nv tf.sif python3 myscript.py
```

## 1.6 Batch System

The JSC systems use a special flavor of Slurm as the workload manager (PSSlurm). Most of the vanilla Slurm commands are available with some Jülich-specific additions. An overview of Slurm is available in the according documentation which also gives example job scripts and interactive commands: <https://apps.fz-juelich.de/jsc/hps/jureca/batchsystem.html>

Please account your jobs to the `training2406` project, either by setting the according environment variable with the above `jutil` command (as above), or by manually adding `-A training2406` to your batch jobs.

Different partitions are available (see [documentation for limits](#)):

- `dc-gpu`: All GPU-equipped nodes
- `dc-gpu-devel`: Some nodes available for development

For the days of the Hackathon, reservations will be in place to accelerate scheduling of jobs.

- Day 1: `--reservation gpuhack24`
- Day 2: `--reservation gpuhack24-2024-04-23`

- Day 3: `--reservation gpuhack24-2024-04-24`
- Day 4: `--reservation gpuhack24-2024-04-25`
- Day 5: `--reservation gpuhack24-2024-04-26`

X-forwarding sometimes is a bit of a challenge, please consider using *Xpra* in your Browser through Jupyter JSC!

## **1.7 Etc**

### **1.7.1 Previous Documentation**

More (although slightly outdated) documentation is available from the 2021 Hackathon [in the according JSC Gitlab Hackathon docu branch](#).

### **1.7.2 PDFs**

See the directory `./pdf/` for PDF version of the documentation.