

INTRODUCTION TO HPC CLUSTER USE

HOW TO TRANSFORM SNĚHURKA INTO CHIMERA

JAROSLAV HRON

FEBRUARY 6, 2026

GOALS FOR TODAY

By the end of this session, you should be able to:

- ▶ Log into a Chimera cluster
- ▶ Navigate storage (home, scratch, project) and move data in/out
- ▶ Run software with environment modules
- ▶ Submit and manage batch jobs (Slurm examples)

- ▶ Typical components as represented on Sněžurka:
- ▶ **Login nodes** `r3d3.karlin.mff.cuni.cz` (interactive: edit code, prepare jobs, file/data management)
- ▶ **Compute nodes** `r41, r42, ...` (where your jobs run)
- ▶ **Shared storage** (`/home` , `/work` , `/scratch`)
- ▶ **Scheduler** Slurm (allocates resources and queues work to partitions `express` , `long`)

SOME DEFINITIONS....

Sněhurka (Snow White and the Seven Dwarfs)

Oxford Learner's Dictionaries: a traditional children's story. Snow White is a beautiful princess. Her stepmother is jealous of her beauty and orders a man to kill her. He feels pity for her and leaves her alive in a forest, where she lives happily with seven dwarfs.

Chimera (Χίμαιρα)

Oxford Learner's Dictionaries: a **chimera** is

- ▶ a fire-breathing female monster with a lion's head, a goat's body, and a snake's tail
- ▶ something that is hoped for but is impossible to achieve.

Merriam-Webster Dictionary: a **chimera** is

- ▶ a mythical fire-breathing monster made of parts of different animals
- ▶ something made up of very different parts (or an unrealizable illusion).

BASIC COMPARISON

Sněhurka

- ▶ about 1024 cores (intel, amd)
- ▶ internal net 100 GB/s InfiniBand
- ▶ Ubuntu 22.04 LTS

login node: r3d3.karlin.mff.cuni.cz

compute nodes: r30, r31, ..., r50

home path: /usr/users/<login>

work path: /usr/work/<login>

local temp path: node local /tmp

shared temp path: none

Chimera

- ▶ about 3064 cores (intel, amd)
- ▶ internal net 200 GB/s InfiniBand
- ▶ Rocky Linux 9 OS (EL9 Linux)

login node: hpc.troja.mff.cuni.cz

compute nodes: various names + (r42)

home path: /home/<login>

work path: /work/<login> (on request)

local temp path: memory mapped /tmp

shared temp path: /scratch

CHIMERA: INFORMATION RESOURCES

Our info page monitoring the transition:

- o <https://cluster.karlin.mff.cuni.cz/relocation/>

Main MFF HPC info page:

- 1 <https://www.mff.cuni.cz/en/hpc-cluster/general-information>
- ▶ <https://www.mff.cuni.cz/en/hpc-cluster/jupyter-hub>

Faculty gitlab service:

- 2 <https://gitlab.mff.cuni.cz/mff/hpc/clusters#mff-hpc-metacenter>
- 3 <https://gitlab.mff.cuni.cz/mff/hpc/clusters/-/issues>

Existing tutorials from other departments:

ucjf https://ipnp.cz/?page_id=8244

CHIMERA: INITIAL STEPS TO CREATE YOUR ACCOUNT

1. Initialize your gitlab account at `https://gitlab.mff.cuni.cz`

(Note: Karlin local gitlab is separate service `https://gitlab.karlin.mff.cuni.cz`)

2. Create your account on Mattermost .

This is useful to get notifications from admins about planned outages and report possible issues with Chimera and/or JupyterHub.

3. Initialize your Chimera account by logging to JupyterHub service at `https://hpc.troja.mff.cuni.cz:8000`

See `https://www.mff.cuni.cz/en/hpc-cluster/jupyter-hub` for detailed guide.

Credentials

All credentials are your SIS login name and password.

LOGGING IN (SSH)

- ▶ From macOS/Linux/WSL (or Windows with an SSH client): use terminal with `ssh` client
- ▶ Windows - MobaXterm (<https://mobaxterm.mobatek.net/>), Bitvise ssh (<https://bitvise.com/>)
- ▶ Windows VisualCode - caution!

CHIMERA: SSH ACCESS

Once you have your account activated:

```
# 1) Connect to a headnode
ssh <CASname>@hpc.troja.mff.cuni.cz

# 2) Then access the cluster environment
[<CASname>@hpc-head ~]$ pwd
/home/<CASname>
[<CASname>@hpc-head ~]$ ~hron/bin/freenodes
```

Credentials

All credentials are your SIS login name and password.

CHIMERA: SLURM PARTITIONS (QUEUES)

Chimera uses Slurm (<https://slurm.schedmd.com/documentation.html>).
The cluster is divided into several partitions (\approx queues)

- ▶ "free-for-all": `ffa` (1h-12h), `ffa-short` (30min-2h), `ffa-preempt`,
`ffa-checkpoint`, `gpu-ffa`
- ▶ matematicka sekce: `express` (6h-12h), associated only with our nodes
- ▶ Education: `edu` (2h-4h)

Always request the **smallest** resources/time that fit your job to start sooner.

CHIMERA: SLURM QUICK COMMANDS

Partitions and node state

sinfo

Your jobs

squeue -u \$USER

Interactive allocation (example)

salloc -p ffa -n 4 --mem=8G --time 02:00:00

Batch submission (example)

sbatch -p express job.**sbatch**

CHIMERA: SOFTWARE STACKS OPTIONS

Chimera provides:

- ▶ **Environment modules**
- ▶ **Apptainer/Singularity** containers (reproducible environments)
- ▶ **Charliecloud** based on a user-defined software stack (UDSS)

CHIMERA: STORAGE BASICS

Common storage areas:

- ▶ `/home/<name>` (backed up, smaller quota): suitable for configs, scripts, source code.
- ▶ `/work/<name>` (optional, medium/large): Bigger datasets, computed results.
- ▶ `/scratch/` (fast, temporary): job I/O, checkpoints, intermediates.
- ▶ `/archive/` (slow, permanent): Archive for bigger data not need immediately

```
df -h
```

```
# lfs quota -h /home/<name>  
# du -sh /home/<name>/folder
```

MOVING DATA TO/FROM THE CLUSTER

Typical options:

- ▶ **scp** (simple copies)
- ▶ **rsync** (fast incremental sync)

```
# Copy a file from local machine to the cluster
```

```
scp local.dat <name>@hpc.troja.mff.cuni.cz:/path/to/
```

```
# Sync a directory (recommended for repeated transfers)
```

```
rsync -avzP ./results/ <name>@hpc.troja.mff.cuni.cz:/path/to/results/
```

SOFTWARE ENVIRONMENTS: MODULES

Many clusters use **environment modules** to manage compilers and applications.

- ▶ You typically start with a clean environment and load what you need.
- ▶ Modules ensure consistent versions across users.

```
module avail  
module spider gcc  
module load gcc/12.2.0  
module list  
module purge
```