

TOC for all the pages in the "HPC" namespace

HPC start

- [HPC User Documentation](#)
- [HPC clusters at UNIGE](#)
 - [Use the HPC resources](#)
- [RoadMap](#)
 - [Other resources](#)
- [Support - get help](#)

Baobab and Yggdrasil - our HPC clusters and their infrastructure

- [How our clusters work](#)
- [The clusters : Baobab, Yggdrasil and Bamboo](#)
- [How do our clusters work ?](#)
 - [Overview](#)
 - [Cost model](#)
 - [Price per hour](#)
 - [Free CPU Hour Allocation](#)
 - [Progressive Pricing for HPC Compute Hours](#)
 - [Purchasing or Renting Private Compute Nodes](#)
 - [Key Rules and Details](#)
 - [Cost of Renting a Compute Node](#)
 - [Usage Limits](#)
 - [CPU and GPU server example pricing](#)
 - [AMD CPU](#)
 - [GPU H100 with AMD](#)
 - [GPU RTX4090 with AMD](#)
- [Use Baobab for teaching](#)
- [How do I use your clusters ?](#)
- [For advanced users](#)
 - [Infrastructure schema](#)
 - [Compute nodes](#)
 - [CPUs models available](#)
 - [GPUs models available](#)
 - [Bamboo](#)
 - [CPU MODELS — bamboo](#)
 - [GPUs on Bamboo](#)
 - [Baobab](#)

- CPU MODELS — baobab
- CPU MODELS — baobab
 - GPUs on Baobab
- Yggdrasil
 - CPU MODELS — yggdrasil
 - GPUs on Yggdrasil
- Monitoring performance

How to use the HPC clusters

- [hpc:how_to_use_the_hpc_clusters#](#)

Access : SSH, X2GO

- Access the clusters
 - Account
 - Standard Account
 - External Account
 - Outsider Account
 - Inactivity Notice and Account Deletion Policy
 - Cluster connection
 - login nodes
 - Connect using SSH
 - ssh PublicKey
 - multiple ssh key
 - From Linux and Mac OS
 - From Windows
 - Access to the compute nodes
 - GUI access / Desktop with X2Go
 - File transfer
 - From Linux
 - From Windows
 - SSH tunnel and socks proxy
 - Alternative to using ProxyJump

Best practices and smart use of the HPC resources

- Introduction
- First steps
- Rules and etiquette
- Think green
- Stop wasting resources!

- Which resources ?
- Single thread vs multi thread vs distributed jobs
- Bad CPU usage
- Bad memory usage
- Bad time estimation
- Conclusion
- How to write a good Slurm sbatch script
- Transfer data from one cluster to another
 - Rsync
- SWITCHfilesender from the cluster
 - Configuring the CLI tools
 - Transferring files
- HPC cluster user departure procedure

Slurm and job management

- Slurm and job management
- What is Slurm ?
- Resources
 - Partitions
 - What is a partition ?
 - Which partition for my job ?
 - Partitions lists
 - Clusters partitions
 - Private partitions
 - Wall time
 - Memory
 - GPU
 - CPU
 - CPU types
 - Single thread vs multi thread vs distributed jobs
- Submitting jobs
 - Batch mode (sbatch)
 - Monothreaded jobs
 - Multithreaded jobs
 - Distributed jobs
 - GPGPU jobs
 - Interactive jobs
 - Job array
 - Advanced usage
 - Job dependency
 - Master/Slave
 - Checkpoint
- Reservation

- Job monitoring
 - Email notification of job events
 - Memory and CPU usage
 - Energy usage
 - CPUs
 - GPUs
 - Other tools
 - spart
 - pestat
 - seff
 - HDF5 profiling plugin
- Cancel jobs
- Job priorities
 - How is the priority of a job determined ?
 - Priority vs. waiting time
 - Backfill mechanism

Storage

- Storage
- Cluster storage
 - Home directory
 - Quota
 - Scratch Directory
 - Quota
 - Data Retention Policy
 - Fast directory
 - Quota
- Local storage
 - Scratch directory (local on each node)
 - Temporary private space
 - Temporary shared space
- Sharing files with other users
- Best practices
 - I/O performance
 - Check disk usage on the clusters
 - Check disk usage on home and scratch
 - Check disk usage on NASAC
 - File transfer (between nodes)
- Backup
- Archive
- Access external storage
 - NASAC
 - Troubleshooting

- Where does gio mounts my data?
 - List the user DBUS process
 - Sometimes mount is not available but you can browse/copy/interact with gio commands
- CVMFS
- EOS
- Robinhood
 - Policies
 - Report

Applications and libraries

- Applications on the clusters
 - Module - lmod
 - How to use 'module'
 - What do I do when an application is not available via 'module' ?
 - Detailed example of using 'module'
 - Loading 'R'
 - Choosing the compiler toolchain
 - FOSS toolchain
 - Intel toolchain
 - Intel compiler licenses
 - fosscuda toolchain
- Examples for selected applications
 - OpenMPI
 - Specify MCA parameters through "srun"
 - Conda
 - How to Create a Conda Environment in a Container
 - Benefits
 - Limitations
 - Step 1 - Define the Conda Environment
 - Step 2 - Build the Container
 - Step 3 - Use the Container
 - Conda environment management
 - Package management
 - ADF
 - Gaussian
 - Git
 - NVIDIA HPC SDK Installation in Home Directory on a Heterogeneous Cluster
 - Prerequisites
 - Installation Steps
 - Verify Installation
 - Gurobi
 - Jupyter notebook and Jupyter lab

- [Mathematica](#)
- [Matlab](#)
 - [Parallel with Matlab](#)
 - [Pass sbatch arguments to Matlab](#)
 - [Compile your Matlab code](#)
 - [Matlab PATH](#)
 - [Matlab java.opts](#)
 - [CHROMIUM mailbox/texture errors](#)
 - [Wavelab](#)
- [OpenCL](#)
- [Distant Paraview](#)
- [Python](#)
 - [Custom Python lib](#)
 - [Pip install from source](#)
- [R project and RStudio](#)
 - [RStudio](#)
 - [R packages](#)
- [Variant Effect Predictor \(VEP\)](#)
 - [Install species](#)
- [Apptainer \(was Singularity\)](#)
 - [Intro](#)
 - [Pull an existing image](#)
 - [Convert a Docker image](#)
 - [Run a container](#)
 - [Modify the image \(not persistent\)](#)
 - [Modify the image \(persistent\)](#)
 - [References](#)
- [Stata](#)
- [TensorFlow](#)
- [Compile and install a software in your /home](#)

Glossary

- [Glossary](#)

FAQ

- [FAQ: Frequently Asked Question](#)
 - [General](#)
 - [Cost](#)
 - [Account](#)
 - [Connection to Cluster](#)
 - [X2GO-Desktop](#)

- [Storage](#)
- [Applications](#)
- [Slurm: job scheduler](#)
- [Issues](#)
- [Switch edu-ID Login Issues](#)
- [HPC community forum](#)

From:

<https://doc.erresearch.unige.ch/> - **eResearch Doc**

Permanent link:

<https://doc.erresearch.unige.ch/hpc/toc?rev=1602166240>

Last update: **2025/06/11 12:27**

