

How to use the HPC clusters

This page gives you directions to start using the clusters **Baobab** and **Yggdrasil**, from your first baby steps to a more advanced use.

Introduction

An HPC cluster is an advanced, complex and always-evolving piece of technology. It's easy to forget details and make mistakes when using one, so don't hesitate to check this section every now and then, yes, even if you are the local HPC guru in your team! There's always something new to learn !

The HPC utopia, citizens and basic users

As a user, you are sharing a fabulous technological beast with hundreds of other researchers. Any research is important (and yes we realise *your* research is more important!), nevertheless we expect all of you to act not only as basic users, but as good HPC citizen.

Indeed, it is possible to use the cluster by just asking your favorite colleagues to give your pointers and adapt the script they wrote 5 years ago. BUT it's unlikely that will do the job for your project. Worse, this might even have severe consequences on your jobs and impact other users. So don't be like that! You can easily become a good HPC citizen by reading the [Best practices and smart use of the HPC resources](#) section. We also encourage you to have at the very least some basic knowledge of Linux, command line (without GUI) and scripting.



add links to Linux and Bash tutorials.

As a good HPC citizen at the UNIGE, those are your rights and obligations :

- **Use the clusters only for research purposes**
 - This means you **cannot** mine Bitcoins here, nor can you store your holiday pictures.
- Use `module` to load the libraries or applications you need.
 - Pro tip: you can even force the version to have consistent results and survive an OS migration, for example.
- You can compile binaries in your `$HOME` directory and use them on any node of the cluster (since your `$HOME` is accessible from any node).
- You can request the HPC team to install a new software or version of a library in order to load it through `module`.
 - Just use the [HPC Wish-list](#).
- You **cannot** install new software with `yum` (and don't bother with `apt-get`, we are not running Ubuntu).
- **You are not sudo** and will never be
 - N.B. this is only available for the sacred HPC engineers (Linux-type beard and hoody mandatory)
- And you can vote when we choose the name for a new cluster every 10 years!

First steps

You are now ready to start using the HPC clusters, and hopefully to be a good HPC citizen. But how do you start ?

For your first steps we recommend the following :

- (have basic knowledge of Linux and Bash scripting)
- Connect to the login node of the cluster you are planning to use :
 - [Access : SSH, X2GO](#)
- Check the page [Best practices and smart use of the HPC resources](#) to make the most of the cluster.
 - This page contains important information! You can hog resources and/or wait much longer than you could if you don't request the right amount of time for your job, if you ask too much (or not enough) resources, if you are not using the right partition, etc.
- Understand how to load your libraries/application with module
 - [Applications and libraries](#)
- Learn how to write a Slurm sbatch script
 - [Slurm and job management](#)
- Check again the [Best practices and smart use of the HPC resources](#) (you thought you could just skip it earlier, didn't you? It's like you don't even want to be a good HPC citizen...).



From:

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

Permanent link:

https://doc.eresearch.unige.ch/hpc/how_to_use_the_hpc_clusters?rev=1592427382

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

