

HPC User Documentation

Welcome on UNIGE's High Performance Computing (HPC) Baobab service documentation. You will find here all the information you need to use the unige's clusters Baobab, Yggdrasil, Bamboo and get started.

In the documentation, when we talk about Baobab we may refer as the HPC service or one of our cluster.. Baobab.

Before accessing our HPC clusters, you need to :

- Have an account - if you don't have one yet, request an account [here](#).
- Read and agree to the [Terms of Use](#)

HPC clusters at UNIGE

[HPC for dummies, an introduction to HPC](#)

Use the HPC resources

- [How to use Linux](#) - for Linux beginners
- [How our clusters work](#) - description of the HPC environment
- [Access the clusters](#) - how to connect with SSH or X2Go, transfer files
- [Applications and libraries](#) - how to use module to load the available applications on our clusters
- [Storage](#) - where to store your data on the different storage available on the clusters
- [Data Life cycle](#) - How to manage your data
- [Slurm and job management](#) - how to request computing resources with Slurm
- [Resource utilization and accounting](#) - cpu hours used, utilization reporting
- [Best practices](#) - best practices and smart use of the HPC resources
- [HPC team](#) - HPC team and local HPC contact
- [FAQ](#) - Frequently Asked Questions



[How to use OpenOnDemand](#)

ganttt title HPC Roadmap dateFormat 2025-01-01 section Maintenance section Robinhood Deploy on Boabab :B1, 2025-02-20, 30d deleted old files - 1 year :B2, after B1 , 31d deleted old files - 6 month :B3, after B2 , 31d deleted old files - 3 month :B4, after B3 , 31d Deploy on Yggdrasil :Y1, 2025-05-20, 30d deleted old files - 1 year :Y2, after Y1 , 31d deleted old files - 6 month :Y3, after Y2 , 31d deleted old files - 3 month :Y4, after Y3 , 31d

Other resources

- [VDI Infrastructure](#) - for GPU computation that cannot be done on the clusters

Support - get help

- **FAQ** : First of all, check the [FAQ for general questions and troubleshooting](#).
- Still couldn't find what you need ?
 - use the **search engine** (use @hpc to only search HPC related pages)
 - check the detailed [Table of content](#).
 - The [Glossary](#) if everything sounds like gibberish.
- **Forum/Community** (mandatory for non-private requests): The [HPC community forum](#) is the required platform for asking questions if you didn't find a solution here. The forum is quite active and is the best way to ask for help, share, or discuss with the community.
 - Please use the **Support>issue category** and follow the template for any problems you encounter.
- **HPC-Lunch** : First Thursday of each months, you can join our zoom session to discuss with us about HPC. More [here](#)
- **Email** : For any other requests (confidential questions, meeting request, etc.) you can contact the HPC team by email (see below)

[Contact the HPC team by email](#)

README CAREFULLY

WELL WELL WELL, if you read this message is that you have run out all available documentation ressources, isn't it ?

That is to say:

- You have googled your issue and found nothing
- You have read or search in the official documentation (slurm/linux centos|redhat/etc...)and found nothing
- You have post your issue on [hpc-community forum section hpc-support>hpc-issues](#) , and you get no answer.
- You have read or search in the Unige HPC documenation, and guess what ? nothing neither
- You have a specific request 

Here he is! You are The one and only winner of the Gemini Croquette Contest, and earn the righ to contact the HPC-team. what luck!

You can contact the HPC team by email :

When you reach out for help, we always do our best to get back to you as soon as possible, but please

bear in mind that :

- **Hundreds of other people** use our clusters and almost each of you use some special combination of softwares (compilers, libraries, using specific version, etc.). **We don't have experts for each software.**
- The best way to get help quickly is to read this documentation or search/post on [HPC community forum](#).
- When you ask for help, please **give us enough details** so we know exactly what you did, such as :
 - What did you try, what didn't work, what is the expected result, what error message you have, etc.
 - If this information is missing, most of the time, **we will just reply and tell you we need this informations... !**
- **Respect the following mail template**, use the mail balise in the subject to be more specific like
 - Subject: **[baobab][yggdrasil][slurm] error launching jobs with srun**

```
<Hello formula >
```

```
Username: $USER
```

```
Cluster: $CLUSTER
```

```
Subject: <slurm|software|system|connexion| (an explanatory keyword) >
```

```
jobid: <separate by space if you have more than one>
```

```
<
```

```
<Here the mail body explaining ans describing everething with the  
MAXIMUM details. the MAXIMUM!>
```

- Explain us your experience like we were newbies.
we don't spy your work so we don't know what is going wrong :/
- Explain the purpose of your mind, sometimes what we want is not what we think/do/are doing.
- Give us the command you executed, your sbatch, and all relevant informations
- Give us the path to the relevant files in Baobab (sbatch, error logs, etc.) and/or add them as attachments.

```
>
```

```
<Respect formula (it's always appreciated)>
```

```
<Bye Bye formula>
```

```
--
```

```
$USER
```

Here our contact hpc@unige.ch

For more general information about the HPC at UNIGE and eResearch : unige.ch/eresearch/hpc
(Digital infrastructures and services for research)

From:

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

Permanent link:

<https://doc.eresearch.unige.ch/hpc/start?rev=1739543196>

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

