

- General
 - I have a problem, what should I do?
 - Which cluster should I use ?
 - Must I include citations and acknowledgments in my publication?
 - Why is the cluster running slowly ?
- Cost
 - I'm using Open XDMoD to check how many hours my group used this year, but the amount is quite different from what appears on the invoice or in "ug_slurm_usage_per_user.py".
 - My group only used our private partition. I heard from a colleague that it is free of charge, yet I still received an invoice.
 - Your documentation states the cost per hour is 0.0157 CHF, but my invoice shows 0.02 CHF.
 - I'm a PI and want to see usage details per user.
 - I have no idea why I received your email about 'HPC billing'.
 - I'm not interested in receiving further information about HPC at UNIGE, can you please remove me from the hpc-community mailing list?
 - I'm a PI, how do I know which users are associated with me on Baobab?
 - I'm a faculty/group manager, how may I have a list of every PI of a given dept?
 - If multiple PIs jointly purchase computing resources on the Baobab cluster, who receives the invoice?
 - I'm a PI, I tried to use OpenXDmoD to see the past usage of my group without success
 - With OpenXDmoD how can I check usage on more than one partition?
 - I want to login to OpenXDmoD, what are the login details?
 - I'm a user and I've noticed that I'm connected to two PIs, how is this possible?
 - I'm organizing a course and we need some HPC resources for the students. Do we have to pay for it?
- Account
 - Who can be registered as PI (Primary Investigator) at Baobab?
 - When does my account expire ?
 - I'm leaving UNIGE, can I continue to use Baobab HPC service?
- Connection to Cluster
 - When I type my password, no characters are printed. Why?
 - When I tried to connect to the cluster, there is no response.
 - SSH "Could not resolve hostname XXXX: Name or service not known"
 - When I try to connect to Clusters using "ssh" or "sftp", I see the message: Connection refused
 - I have forgotten my password, can the HPC team reset it?
 - How to check my SshPublicKey ?
 - Is it possible to access my account from more than one SSH key?
 - I tried to connect without success.
 - I'm having trouble connecting with software like PuTTY, FileZilla, or X2Go. What should I do?
 - I'm an "outsider" and I can't connect to open on demand
- X2GO-Desktop
 - Why I can't connect with x2go ?
- Storage
 - I have a question about the storage !?
 - How can I request for a shared directory
 - How can I access to a shared directory?
 - How can I copy data from one cluster to another one?

- Applications
 - What applications are installed on Clusters ?
 - The software I need is not available on Clusters: what should I do ?
 - Can I use any Microsoft Windows software ?
 - Can I use a proprietary licensed software ?
 - I need a different Linux distributions/version, am I stuck ?
 - Illegal instruction
 - How can I use another Python version ?
 - Can I load two versions of the same software? How can I use two different software versions with different GCC dependencies?
- Slurm: job scheduler
 - What is Slurm ?
 - I am already familiar with "torque/pbs/sge/lsf/...", what are the equivalent concepts in slurm ?
 - Can I run some small test runs in the login node ?
 - What partition should I choose ?
 - Can I launch a job longer than 4 days ?
 - How are the priorities computed ?
 - Why My jobs stay a long time in the pending queue ?
 - Can I run interactive tasks ?
 - I want to run several time the same job with different parameters
 - Is the nodes from my partition in use?
 - Why I'm not able to use all the cores of a compute node ?
 - How can I access to a private slurm partition?
- Issues
 - I have a keyboard issue using a Mac.
 - When I ssh, I get the message : "cannot change locale (UTF-8): No such file or directory"
 - When I try to connect to the cluster from a Mac using "ssh -Y" and I receive an error like:
- Switch edu-ID Login Issues
 - I get an error message from Switch edu-ID while trying to access:
- HPC community forum
 - I don't find a way to receive email summary of new post

FAQ: Frequently Asked Question

General



I have a problem, what should I do?

Please follow these steps:



- Review this FAQ to see if your issue is addressed.
- Check the current issues on the cluster here:
<https://hpc-community.unige.ch/t/2024-current-issues-on-hpc-cluster/> (A new post is created each year for reference).
- Post in the [HPC-community](#) under the category **HPC issue > HPC support** using the



Template.



Which cluster should I use ?



You can use the three clusters, but see [this link](#) to help you choose the right cluster.



Must I include citations and acknowledgments in my publication?

Yes, according the [terms of use](#) you **must** include at least:



"The computations were performed at University of Geneva using Baobab HPC service."



Why is the cluster running slowly ?

There could be several reasons for the cluster to slow down. It's important to figure out where the slowness is happening:



- **Login Node:** If the login node feels slow, it might be because someone is running heavy processes on it, which isn't recommended. The login node is meant for tasks like file editing, job submission, and monitoring—not running jobs. If another user is hogging the CPU resources, it could affect your experience, but this won't impact the performance of jobs on the compute nodes.
- **Compute Nodes:** Slowness on the compute nodes might be due to high CPU usage, storage issues, or other factors, which could cause your jobs to run more slowly.
- **Storage (Home, Scratch, Other):** If there's a problem with storage (like home directories or scratch space), it can slow down the entire cluster and affect your job performance.

What You Can Do: Make sure you're not contributing to the slowdown. Use the `htop` command on the login node to check CPU usage. If you see that all the CPUs are in use, take a screenshot and send it to us at hpc@unige.ch so we can look into it.

Cost



I have no idea why I received your email about 'HPC billing'.

The message is about the fact that the high performance computing service known as Baobab will become a paid service after a free quota has been used. We sent the announcement to two mailing lists:



- baobab-announce: which includes all users of the Baobab service.
- hpc-community: very low-traffic mailing list containing all PIs and people interested in the HPC community. It may happen that you belong to the two mailings.



I'm not interested in receiving further information about HPC at UNIGE, can you please remove me from the hpc-community mailing list?

If you are a UNIGE member or have a [switcheduid](#) account, you can unsubscribe from the "hpc-community" list on [sympa web interface](#).

An alternate method is to send an email to sympa@listes.unige.ch with the following mail body "UNSUBSCRIBE hpc-community". This mail must be sent using the email you wish to unsubscribe from.



If you are not a UNIGE member or if none of the previous steps worked, please send a request to hpc@unige.ch, subject: "please unsubscribe me from the hpc-community mailing list".

Please note that you can't unsubscribe from the "baobab-announce" list if you still have an account on the Baobab.



I'm a PI, how do I know which users are associated with me on Baobab?

If you have access to one of the clusters, you can use the sshare command:

```
(baobab)-[root@admin1 ~]$ sshare -a -A <your_isis_username>
```

Account	User	RawShares	NormShares	RawUsage
EffectvUsage	FairShare			

isis_pi		41	0.014594	73169235
0.031775				
isis_pi	user1	1	0.000768	130935
0.000239				
isis_pi	user2	1	0.000768	5069653
0.000300				
isis_pi	user3	1	0.000768	0
0.000000				
isis_pi	user4	1	0.000768	0



```
0.000000  1.000000
isis_pi      user5      1      0.000768      1707102
0.000285  0.773432
[...]
```



You can also use [OpenXDmoD](#) to check user usage. Note that the list may be incomplete: for example, if a registered user has never used the cluster in the time period you specify, they won't appear at all.



I'm a faculty/group manager, how may I have a list of every PI of a given dept?

You can use `sacctmgr` for that purpose

```
sacctmgr show assoc where parent=<your_department_name>
cluster=baobab format=account
```

If you don't know the name of your department as registered in our cluster, you can list them by faculty:



```
sacctmgr show assoc where parent=sciences cluster=baobab
format=account
Account
-----
astro
biad
biani
bicel
[...]
```



I'm a PI, I tried to use OpenXDmoD to see the past usage of my group without success



We have a [tutorial](#) which explain how to do that.



How can I check usage on more than one partition?



Unfortunately, it seems that you need to do this operation for each partition separately.

I want to login to OpenXDmoD, what are the login details?

-  User authentication isn't available at the moment. You can access all metrics without authentication. In the future, you'll be able to connect using your [switcheduid](#) credentials, with the benefit of being able to create custom dashboards.

I'm a user and I've noticed that I'm connected to two PIs, how is this possible?

-  The PI must be seen as a project. You can be part of two projects, and when you submit a job to the cluster, you can specify which project to charge to using the `--account` flag.

I'm organising a course and we need some HPC resources for the students. Do we have to pay for it?

-  The Baobab service is free for courses as long as the usage is low and for a defined period of time. Check [How our clusters work](#).

Account

Who can be registered as PI (Primary Investigator) at Baobab?

- 
- Anyone with an ISIs account (even an external one) with a fairly long validity period: teachers rather than assistants.
 - Someone who knows the users they are going to invite. All their users will be under their responsibility.
 - Someone who knows the service(s) for which they will be inviting outsiders. HPC in this case
 - Person responsible for teaching or research (therefore responsible for the data generated and able to know what to do with it when a user leaves).

When does my account expire ?

-  * If you have a non student account (Phd, postdoc, researcher), your account will expire at the same time your contract expire at UNIGE. Right now, there is a grace period after the end

of your contract of around 6 months.

- If you have an outsider account, you need to check the expiration date you received when you filled the invitation.
- If you have an unige student account, you can check the expiration date with the chage command:

(baobab)-[yourusername@login2 ~]\$ chage -l yourusername

Last password change	: Apr 01, 2022
Password expires	: never
Password inactive	: never
Account expires	: never
Minimum number of days between password change	: 0
Maximum number of days between password change	: 99999
Number of days of warning before password expires	: 7



I'm leaving UNIGE, can I continue to use Baobab HPC service?



Yes it is possible as long as you collaborate tightly with your former research group. Your PI must [invite](#) you as [outsider](#). For technical reason, your account needs to be expired prior doing the request for the invitation. We'll then reactivate your account. You'll keep your data.

Connection to Cluster



When I type my password, no characters are printed. Why?



Unlike Windows systems, Linux and Unix systems **do not display** any characters (not even *) when you enter your password in a terminal. The field remains blank, and the cursor will not move.

Simply type your password and press Enter. Your connection should be successful.

Please be cautious not to mistype your password multiple times, as you may be temporarily blocked (see below).



When I tried to connect to the cluster, there is no response.



We employ fail2ban on the clusters to prevent brute-force attacks.

If you enter the wrong password three times consecutively, you will be banned for 15 minutes (fail2ban will blacklist your IP address). After 15 minutes, you can attempt to connect again.

If you are still unable to connect after 15 minutes, please contact us with the following information:



- Your username
- Your IP address (you can find it using [this web service](#)).
- The cluster you are attempting to connect to.



SSH "Could not resolve hostname XXXX: Name or service not known"

It means the specified hostname cannot be found, either due to a typo or because the DNS can't resolve it.



- check the login node [hostname](#)

PS: Keep in mind that baobab2 has been decommissioned for 2 years.



When I try to connect to Clusters using "ssh" or "sftp", I see the message: Connection refused

Connection refused



This may occur because you attempted to connect multiple times with incorrect credentials (e.g., wrong username or password), causing your IP address to be blacklisted. Your IP address will be automatically unblocked after 15 minutes.

Please note that your Baobab/Yggdrasil password is the same as your ISIS password, which we do not manage. If you forgot your password or need to verify it, please use the following service: mdp.unige.ch.



I have forgotten my password, can the HPC team reset it?



No, your Baobab/Yggdrasil/Bamboo password is your ISIS password, and we do not manage it.

If you **forgot your password** or need to verify it, please use the following service:

- mdp.unige.ch.



How to check my SshPublicKey ?



- If you are a **collaborator/student/external** user Check on [my-account](#)
- If you are an **Outsider** user Check on [applicant](#)

For more informations please refer to [ssh PublicKey](#) page.



Is it possible to access my account from more than one SSH key?



Yes, please check [Access the clusters](#)



I tried to connect without success.

There are three possible reasons why you may not be able to connect:



- **The cluster is under maintenance.** Maintenance occurs periodically. Please check your email (including junk/spam folders) or visit the [HPC-community](#) for announcements.
- **Your network is blocking access to our clusters or the SSH protocol.** We use public IP addresses for the login nodes. If you cannot connect, please contact your local network administrator to determine if there are any restrictions on accessing [login1.baobab.hpc.unige.ch](#), [login1.yggdrasil.hpc.unige.ch](#), or [login1.bamboo.hpc.unige.ch](#), or if port 22 is blocked. you can receive this message: `ssh: connect to username@login1.baobab.hpc.unige.ch port 22: Connection timed out`
- **The login node is down.** While unlikely, if this occurs, please wait a little or contact us if the issue persists beyond 15 minutes.



I'm having trouble connecting with software like PuTTY, FileZilla, or X2Go. What should I do?



The most often reason is your is not up to date. Check update and try again. The most common reason is that your software is outdated. Please check for updates and try again. If the issue persists, refer to the FAQ section on [connection_to_cluster](#) for more troubleshooting steps.



I'm an "outsider" and I can't connect to open on demand

This should be fixed in the future, but in the meantime, we have a workaround:



- connect to <https://openondemand.baobab.hpc.unige.ch> using the same email you used when register as outsider. You'll get the following error: Error – failed to map user ()
- Then go to [the session page](#) and send us a screenshot. We'll activate manually your account.

X2GO-Desktop



Why I can't connect with x2go ?

We have already identified a number of common problems:

- Check the general FAQ: [connection_to_cluster](#)
- [Check your quota](#); reaching the limit will prevent you from writing to your directory, which means X2Go won't be able to initialize the necessary configurations.
- If you're using Anaconda/conda, try commenting out the conda block in your .bashrc file.



```
# >>> conda initialize >>>
# !! Contents within this block are managed by 'conda init' !!
__conda_setup="$(('/path/to/your/anaconda3/bin/conda' 'shell.bash'
'hook' 2> /dev/null)"
if [ $? -eq 0 ]; then
    eval "$__conda_setup"
else
    if [ -f "/path/to/your/anaconda3/etc/profile.d/conda.sh" ]; then
        . "/path/to/your/anaconda3/etc/profile.d/conda.sh"
    else
        export PATH="/path/to/your/anaconda3/bin:$PATH"
    fi
fi
unset __conda_setup
# <<< conda initialize <<<
```

- Make a backup(**steps by steps**) of the following files or directories and try to login again:
 1. **~/.bashrc**
 2. **~/.Xauthority**
 3. **~/.x2go**
 4. **~/.local/session**
 5. **~/.config/xfce**



Storage

I have a question about the storage !?



- Where should I store my files?
- What should I do if I delete something by mistake?
- Is there a backup?
- How can I restore a deleted file?
- How much storage space is available?
- My job creates lots of temporary small files, and everything is slow. What should I do?



For detailed information on all storage-related topics, please refer to our [Storage page](#). This page provides comprehensive guidance on file storage, recovery, and managing storage space efficiently.

If you need to store a large amount of data, consider using the “Academic NAS” service, which you can find here: [Academic NAS](#).



How can I access to a shared directory?

To access a **shared directory**, you need to be added to the appropriate group.



Please send an email to hpc@unige.ch including relevant information (Username, Group, private_partition etc...) with the responsible person for the share or partition in CC. The responsible person **must** approve the modification.



How can I copy data from one cluster to another one?



If you have a lot of data, the best way is to use rsync between both clusters, so you won't have to copy the data to your laptop first. [Transfer data from one cluster to another](#)

Applications



What applications are installed on Clusters ?



You can find information about available applications [here](#)



The software I need is not available on Clusters: what should I do ?



Please check [this documentation](#).



Can I use any Microsoft Windows software ?



Baobab is a GNU/Linux only machine, like the majority of academic clusters. If you have a windows software that could run on a Windows cluster, contact us at hpc@unige.ch, perhaps we could find some solutions.



Can I use a proprietary licensed software ?



Yes we can install it, but you should pay the required license. Send us a request at hpc@unige.ch.



I need a different Linux distributions/version, am I stuck ?



No, please check the [Apptainer](#) documentation.



Illegal instruction



If you run a program and it crashes with an error “Illegal instruction” the reason is probably because you have compiled your program on Baobab login node and your program is running on an older compute node on which the CPU lacks some specialized functionality that were used during the compilation.

You have two possibilities:

1. Recompile your program with less optimization, or compile on an older node. See [Advanced users](#)



2. Only run your program on newer servers. See [Specify the CPU type you want](#) and [Compute nodes](#).



How can I use another Python version ?

You need to distinguish between the system-installed Python package and the Python versions provided by module or easybuild. Since we support a variety of software needs for our users, we use module to manage different software versions, including multiple Python versions. To switch between them, you can use the module command to load the specific Python version you need.

Python:

Description:

Python is a programming language that lets you work more quickly and integrate your systems more effectively.



Versions:

Python/2.7.11
[...]
Python/3.11.5

For detailed information about a specific "Python" package (including how to load the modules) use the module's full name.

Note that names that have a trailing (E) are extensions provided by other modules.

For example:

```
$ module spider Python/3.11.5
```



Can I load two versions of the same software? How can I use two different software versions with different GCC dependencies?



No, you cannot load two versions of the same software simultaneously. Additionally, if two software packages depend on different GCC versions, you will not be able to load them at the same time.

- ✓ In this case you need to check if there is another version available compatible with the toolchain (GCC, foss etc...) you want to use. If not, please refer to [The software I need is not available on Clusters: what should I do ?](#).

Slurm: job scheduler

? What is Slurm ?

Slurm is a job scheduling system used to manage and allocate resources in a computing cluster. It helps you submit, monitor, and control jobs (tasks) on the cluster. Please take a moment to review this very important section: [Slurm and job management](#)

- ✓ **As a reminder:** It is **forbidden** to run heavy compute jobs on the login nodes, you **must** use a compute node instead.

? I am already familiar with "torque/pbs/sge/lsf/...", what are the equivalent concepts in slurm ?

Have a look at this scheduler "rosetta stone", available here:

- ✓ <http://slurm.schedmd.com/rosetta.pdf>

? Can I run some small test runs in the login node ?

- ✓ **No never.** You **must** use SLURM to run any test. The debug partition is dedicated to small tests.

What partition should I choose ?

- ? See our documentation about [Slurm Partitions](#).

Can I launch a job longer than 4 days ?

No. Unfortunately you can't. If we raised this limit, you will have to wait longer before having your pending jobs started. We think that the 4 days limit is a good trade-off.

However there could be two work-around if you experience an issue with this limit:

-  1. Some software feature **checkpointing**. During runtime, the program will periodically save its current state on the disks. In that case, this snapshot may be used to resume the computation by another job. Check if your program allows checkpointing. If you cannot find the information, try contacting the developer or ask us at hpc@unige.ch.
- 2. You could add private notes to Baobab. In that case the limit will be raised to 7 days or even higher. If you are interested, contact us.

How are the priorities computed ?

See [here](#)

-  To get the priority calculation details of the jobs in the pending queue, you can use the command: `sprio -w`. You can also have a look at the weights, by typing `sprio -l`.

Why My jobs stay a long time in the pending queue ?

See

- 
 - [Which partition for my job](#)
 - [Job priorities](#)
 - [Stop wasting resources](#)

Can I run interactive tasks ?

Yes, you can. But it is really awkward because you cannot be sure when your job will start.

-  See [Interactive jobs](#)
- You may be interesting about [OpenOnDemand](#) which provide a graphical to start Interactive session (JupyterLab, MatLab, VScode, R etc...)



I want to run several time the same job with different parameters



In that case you can use the **job arrays** feature of SLURM. Please, have a look at the documentation [Job array](#)



Why I'm not able to use all the cores of a compute node ?

Indeed, we are reserving two cores per node for system tasks such as data transfer, and os stuff.

```
(yggdrasil)-[root@admin1 ~]$ scontrol show node cpu001
NodeName=cpu001 Arch=x86_64 CoresPerSocket=18
CPUAlloc=0 CPUEfctv=34 CPUTot=36 CPULoad=0.01
AvailableFeatures=GOLD-6240,XEON_GOLD_6240,V9
ActiveFeatures=GOLD-6240,XEON_GOLD_6240,V9
Gres=(null)
NodeAddr=cpu001 NodeHostName=cpu001 Version=23.02.1
OS=Linux 4.18.0-477.10.1.el8_8.x86_64 #1 SMP Tue May 16 11:38:37
UTC 2023
RealMemory=187000 AllocMem=0 FreeMem=185338 Sockets=2 Boards=1
CoreSpecCount=2 CPUSpecList=17,35 <===== this means
we have two specialization cores <<<<
State=IDLE ThreadsPerCore=1 TmpDisk=150000 Weight=10 Owner=N/A
MCS_label=N/A
Partitions=debug-cpu
BootTime=2023-08-10T12:08:11 SlurmdStartTime=2023-08-10T12:09:00
LastBusyTime=2023-08-11T10:06:42 ResumeAfterTime=None
CfgTRES=cpu=34,mem=187000M,billing=34
AllocTRES=
CapWatts=n/a
CurrentWatts=0 AveWatts=0
ExtSensorsJoules=n/s ExtSensorsWatts=0 ExtSensorsTemp=n/s
```

If you really need to use all the cores of a compute node, you can override this parameter: `-core-spec=0`. This will implicitly lead to an exclusive allocation of the node.

ref: https://slurm.schedmd.com/core_spec.html



How can I access to a private slurm partition?



To use a **private Slurm partition**, you need to be added to the appropriate group.



Please send an email to hpc@unige.ch including relevant information (Username, Group, private_partition etc...) with the responsible person for the share or partition in CC. The responsible person **must** approve the modification.

Mac Issues



I have a keyboard issue using a Mac.



Please refer to this [keymap-issues-with-nx-from-mac-os-x](#) for a potential solution.

When I ssh, I get the message : "cannot change locale (UTF-8): No such file or directory"



```
-bash: warning: setlocale: LC_CTYPE: cannot change locale (UTF-8): No such file or directory
```



You can resolve this issue by following Step #1 [here](#).

Please ensure that you close all open terminals on your Mac and relaunch them.

When I try to connect to the cluster from a Mac using "ssh -Y" and I receive an error like:



```
Can't connect to X11
```



This issue likely arises because Xorg is no longer provided natively on macOS. You need to install XQuartz.

Refer to this solution: [macOS High Sierra and X11 Forwarding](#).

Switch edu-ID Login Issues



I get an error message from Switch edu-ID while trying to access:



1. <https://hpc-community.unige.ch/>
2. <https://openondemand.baobab.hpc.unige.ch/>
3. <https://openondemand.bamboo.hpc.unige.ch/>

Please follow these links for support:



1. [SWITCH edu-ID Account - Welcome](#)
2. [How to Create or Verify Your SWITCH edu-ID Account and Link It to UNIGE](#)

Ensure that you are using the email address linked to your Switch edu-ID account.


Please also note that your ISIS (UNIGE) password and your Switch edu-ID password are not the same. Verify that you are using the correct password when logging in.

HPC community forum



I don't find a way to receive email summary of new post

You can activate the email summary following those steps:



The screenshot displays the user profile for Yann Sagon. The top navigation bar includes links for Message, Admin, and Expand. Below this, a secondary navigation bar lists various settings: Summary, Activity, Notifications, Messages, Invites, Badges, and Preferences (highlighted with a red circle and number 1). Under the Preferences tab, there are sub-tabs for Account, Security, Profile, Email (highlighted with a red circle and number 2), and Notifications. The Email settings section includes options for when to receive emails (always, only when away, never) and whether to include previous replies at the bottom of emails. The Activity Summary section (highlighted with a red circle and number 3) includes a checkbox for 'When I log in here, send me an email summary of popular topics and replies' and a dropdown menu for the frequency (daily). A green checkmark icon is visible on the left side of the page.

From:

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

Permanent link:

<https://doc.ereseach.unige.ch/hpc/faq?rev=1749026116>

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

