

FAQ

This is a page collecting answers to requests arrived to the [HPC Helpdesk](#). Please check here before sending a specific request.

In this page:

- **General:**
 - How can I add a collaborator to my project?
 - I still didn't receive the username and the link for the 2FA (two factor authentication) configuration?
 - Backup Policy
 - Information about my project on CINECA clusters (end data, total end monthly amount of hours, the usage?)
 - I have finished my budget but my project is still active, how can I do?
 - Which filesystems do I have available? Which usage is intended?
 - How can I check how much free disk have I available?
- **Connection/login**
 - How to change my password?
 - My new password isn't accepted, with error "Could not execute the password modify extended operation for DN"
 - RCM does not connect on the new infrastructure GALILEO100
 - I receive the error message "WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!" when trying to login
 - I keep receiving the error message "WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!" even if I modify known_host file
 - Windows WSL issue DNS resolution failing
 - I have the message "perl: warning: Setting locale failed" when I login. How do I solve?
- **2FA:**
 - Windows PowerShell:
 - verify smallstep error
 - use X11
 - Macintosh:
 - undefined method 'cellar' when installing step
- **Executions/scheduler:**
 - I was copying data or executing something in the login and the process was killed. Why?
 - My job has been waiting for a long time.
 - Troubles running a dynamically linked executable
 - Error invalid account when submitting a job: Invalid account or account/partition combination specified.
 - How can I list all my jobs currently in the queue?
 - Can I modify SLURM settings of a waiting job?
 - How can I place and release a job from hold state?
 - Jobs in interactive mode: launching the program with srun but it remains stuck
 - I get the following message: "srun: Warning: can't honor --ntasks-per-node set to xx which doesn't match the requested tasks yy with the number of requested nodes yy. Ignoring --ntasks-per-node." What does it mean?
 - I'm trying to setup a conda environment but I get connection errors from the package repository. How can I solve this?
- **Performance:**
 - I have found performance problems, what should I do?

General:

- **How can I add a collaborator to my project?**

Project group leaders can also manage their user's membership on their [UserDB](#) page.

- **I still didn't receive the username and the link for the 2FA (two factor authentication) configuration?**

You have to do the complete registration on the UserDB page and to be associated with a project (PI has to add you). Once you have inserted all the necessary information and you are associated with a project a new access button will appear, just click on it and you will receive in two mails the username and the link for the 2FA configuration.

- **Backup Policy**

Just the \$HOME filesystem is guaranteed by daily backups. The backup procedure runs daily and we preserve a maximum of three different copies of the same file. Older versions are kept for 1 month. The last version of deleted files is kept for 2 months, then definitely removed from the backup archive.

- **Information about my project on CINECA clusters (end data, total end monthly amount of hours, the usage?)**

You can **list all the Accounts** attached to your username on the current cluster, together with the "budget" and the consumed resources, with the command:

```
> saldo -b
```

More information in our [documentation](#).

- I have finished my budget but my project is still active, how can I do?

Non-expired projects with exhausted budgets may be allowed to keep using the computational resources at the cost of minimal priority. Ask superc@cineca.it to motivate your request and, in case of a positive evaluation, you will be enabled to use the qos_lowprio QOS.

- Which filesystems do I have available? Which usage is intended?
 - **\$HOME**: to store programs and small light results. This is permanent, backed-uped, user-specific, and local area.
 - **\$CINECA_SCRATCH**: where you can execute your programs. This is a large disk for the storage of run time data and files. It is a temporary area
 - **\$WORK**: An area visible to all the collaborators of the project. This is a safe storage area to keep run time data for the whole life of the project and six months after the end of the project.
 - **\$DRES**: An additional area to store your results if they are heavy. This space is not automatic. You need to request for it writing to superc@cineca.it

More detailed information can be found [here](#).

- How can I check how much free disk have I available?

You can check your occupancy with a command "cindata". Option "-h" shows the help for this command and in our [documentation](#), you can find the description of the [output](#).

Connection/login

- How to change my password?

There are several ways to change the password:

- by clicking on Forgot password while authenticating on <https://sso.hpc.cineca.it> (Only if you have already configured 2FA)
- by clicking on Update button in the Password section of your Account Security page in <https://sso.hpc.cineca.it>

If you still need to configure your 2FA, please write to superc@cineca.it to receive the temporary configuration link.

- My new password isn't accepted, with error "Could not execute the password modify extended operation for DN"

The error message can be difficult to interpret, but it means that the new password you have chosen does not respect our password policies. Please check the [password policy](#) and choose your new password accordingly.

- RCM does not connect on the new infrastructure GALILEO100

Please check that you are using the latest version of RCM. You can download the application compatible with your operating system in the [download page](#). Moreover, if you still experienced the issue, please delete .rcm from your home on the Galileo100 (it was copied from the home of the old infrastructure Galileo)

- I receive the error message "WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!" when trying to login

The problem may happen because we have reinstalled the login node changing the fingerprint. We should have informed you through an HPC-news. If this is the case you can remove the old fingerprint from your known_hosts file with the command

```
ssh-keygen -f "~/ssh/known_hosts" -R "login.<cluster_name>.cineca.it"
```

- I keep receiving the error message "WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!" even if I modify known_host file

The procedure we have found so far is the following:

- 1) remove from the known_hosts file in ~/.ssh all the lines associated to the cluster you would like to login (or remove the whole known_hosts file)
- 2) login to the cluster directly to all the login nodes in the following way (for example on Leonardo)

```
ssh <username>@login01-ext.leonardo.cineca.it -o hashknownhosts=no  
ssh <username>@login02-ext.leonardo.cineca.it -o hashknownhosts=no  
ssh <username>@login05-ext.leonardo.cineca.it -o hashknownhosts=no  
ssh <username>@login07-ext.leonardo.cineca.it -o hashknownhosts=no
```

in all the steps above a new line in the known_host file will be added with the fingerprint of all the specific login nodes. Please check the [cluster specific guide](#) for the naming of all the login nodes of the cluster you would like to login.

3) edit the known_hosts file replacing 'login01-ext.<cluster_name>.cineca.it' with 'login*.<cluster_name>.cineca.it' and repeat the same for the other logins.

You may also obtain the same result with a single command.

On UNIX systems (Linux, MacOS):

```
for KEYAL in ssh-rsa ecdsa-sha2-nistp256; do for n in 1 2 5 7; do ssh-keyscan -t $KEYAL login0${n}-ext.leonardo.cineca.it | sed s"/0${n}-ext/\*/" >> ~/.ssh/known_hosts; done; done
```

On Windows Powershell:

```
foreach ($key in "ssh-rsa", "ecdsa-sha2-nistp256") {foreach ($node in "1", "2", "5", "7"){ssh-keyscan -t $key "login0$node-ext.leonardo.cineca.it" | foreach {$_ -replace ("0$node-ext", "**")} | Add-Content -Path "$HOME\.ssh\known_hosts"}}
```

As an alternative (valid for all Cineca machines except Leonardo) you may create a config file, located in \$HOME/.ssh/config for Linux/mac OS, and C:\Users\username\.ssh\config for Windows, with the following content:

```
Host login*.cineca.it
    HostKeyAlgorithms ecdsa-sha2-nistp256
```

Now the problem should not appear again.

• Windows WSL issue DNS resolution failing

If the DNS resolution fails with *Temporary failure in name resolution* or *resolution timing out*, an automatic change in /etc/resolv.conf occurred. You can change it manually by replacing the name server value with 8.8.8.8. This file is automatically generated by WSL: to stop the automatic generation of this file, add the following entry to /etc/wsl.conf: [network] generateResolvConf = false. Then, add in your .bashrc the following commands for the automatic creation of the name server value in the resolv.conf file:

```
> if [ ! -f /etc/resolv.conf ]; then
>     echo "nameserver 8.8.8.8" > /etc/resolv.conf
> fi
```

• I have the message "perl: warning: Setting locale failed" when I login. How do I solve?

This warning is typical of Mac users (but can happen with other OS too). It is actually innocuous and can be safely ignored, but if you want to get rid of it you can add these lines to the .bashrc of your workstation, or in general you can execute

```
them before trying to login to our systems:
export LANGUAGE=en_US.UTF-8
export LANG=en_US.UTF-8
export LC_ALL=en_US.UTF-8
```

if you try to login afterwards, the warnings should have disappeared.

2FA:

- Windows PowerShell:
 - **verify smallstep error**

If running the command *step* to verify the installation of smallstep you incur in the following error:

```
PS C:\Users\user > step
step : The term 'step' is not recognized as the name of a cmdlet,
function, script file, or operable program. Check
the spelling of the name, or if a path was included, verify that the path
is correct and try again.
At line:1 char:1
+ ~~~~
+ CategoryInfo          : ObjectNotFound: (step:String) [],
ParentContainsErrorRecordException
+ FullyQualifiedErrorId : CommandNotFoundException
```

check if you find the executable `step.exe` inside the folder:

```
C:\Users\user\scoop\shims
```

The installation command should have placed it there. If you don't find it, run on your Powershell the command:

```
scoop install smallstep/step
```

- **use X11**

- install Xming: <https://sourceforge.net/projects/xming/>, it will open a window in the background that you won't be able to see but you can see that it's there looking between the icons in the Windows' applications bar (bottom right)
- follow the steps reported at <https://x410.dev/cookbook/built-in-ssh-x11-forwarding-in-powershell-or-windows-command-prompt/> for PowerShell
- then you can run the command `ssh` to login on the cluster

- **Macintosh:**

- **undefined method 'cellar' when installing step**

You may encounter an error that looks like this:

```
Error: step: undefined method `cellar' for #<BottleSpecification:0x000000012e579660>
```

In this case, the problem is in your homebrew. It may refer to the directories of different processes, e.g. Intel, while you need to make it refer to AMD. You can reinstall homebrew:

```
brew tap homebrew/core
```

and then set the proper paths with simple shell commands:

```
(echo; echo 'eval "$(/opt/homebrew/bin/brew shellenv)"') >> /Users/<my_user>/.zprofile
```

```
eval "$(/opt/homebrew/bin/brew shellenv)"
```

If this is not the solution of your error, the command "`brew doctor`" should give you an hint about how to proceed in your specific case.

Executions/scheduler:

- **I was copying data or executing something in the login and the process was killed. Why?**

The login nodes in our facilities have 10 minutes CPU time limit. This means that any execution requiring more than that is automatically killed. You may avoid this restriction by using the batch script on the SLURM partitions or an interactive run. The partition and the resources **depend on the machine** you are considering, so please see the "UG3.0 System specific guide" page. Important details and suggestions on how to transfer your data can be found on the "Data management" page.

- **My job has been waiting for a long time.**

The priorities in the queue are composed of several factors and the value may change due to the presence of other jobs, of the resources required, and your priority. You can see the reason for your job in the queue with the `squeue` command. If the state is PD, the job is pending. Some reasons for the pending state that could be displayed:

- Priority= The job is waiting for free resources.
- Dependency= It is depending to the end of another job.
- QOSMaxJobsPerUserLimit = The maximum number of jobs a user can have running at a given time.

More information about the reason meanings can be found on the SLURM resource limits page. You can also consult the estimated starting run time with the SLURM command `scontrol`:

```
> scontrol show job #JOBID
```

or you can see the priority of your job with the `sprio` SLURM command:

```
> sprio -j #JOBID
```

Here you can find more info on Budget linearization.

- Troubles running a dynamically linked executable

In order to use the available libraries dynamically, you have to add the library directory to the path list of the environment variable `$LD_LIBRARY_PATH` before the execution of your job starts. To see your library path, please load the library module and then use

```
> module show <libraryname>
```

Then export or set the path:

- `export LD_LIBRARY_PATH=${LD_LIBRARY_PATH}:/path>`
- `set LD_LIBRARY_PATH ${LD_LIBRARY_PATH}:/path> (csh)`

- Error invalid account when submitting a job: Invalid account or account/partition combination specified.

- The error **Invalid account** might depend on the lack of resources associated to your project or there is an error with the account name in your batch script. Just use the `saldo` command. If the account is correct and valid, are you launching the job on the right partition? To see which partition is right for your case and account, please consult the Summary paragraph inside the System Specific Guide.

- How can I list all my jobs currently in the queue?

Just do the command

```
> squeue -u $USER
```

and you will see the list of your running or pending jobs.

- Can I modify SLURM settings of a waiting job?

Some Slurm settings of a pending job can be modified using the command `scontrol update`. For example, setting the new job name and time limit of the pending job:

```
> scontrol update JobId=2543 Name=newtest TimeLimit=00:10:00
```

- How can I place and release a job from hold state?

In order to place a job on hold type

```
> scontrol hold JOB_ID.
```

To release the job from the hold state, issue

```
> scontrol release JOB_ID.
```

- Jobs in interactive mode: launching the program with `srun` but it remains stuck

After creating an interactive job with `srun` command

```
srun -N 1 -n X ... --pty /bin/bash
```

if I launch my job with `srun`

```
srun -n X program
```

the command remains stuck and nothing happens.

This is because the resources requested with the second `srun` are already allocated for the first `srun` and the command remains stuck waiting for the resources to be freed.

There are two possible solutions:

1. create the interactive job with `salloc` instead of `srun`

```
salloc -N 1 -n X
```

please refer to [this page](#) with more details about how to use `salloc`.

2. launch the second command with the flag `--overlap`

```
srun -n X --overlap program
```

Note: if you use mpirun instead of srun to launch your program inside an interactive job allocating more than one node you will encounter the same error. If you would like/need to use mpirun on multiple nodes, you can follow the first solution above (salloc). On a single node, instead, the command will work (because in this last case mpirun does not call srun).

- I get the following message: "srun: Warning: can't honor --ntasks-per-node set to xx which doesn't match the requested tasks yy with the number of requested nodes yy. Ignoring --ntasks-per-node." What does it mean?

This is a message that can appear when using mpirun and Intelmpi parallel environment. It is a known problem that can be safely ignored, since mpirun does not read the proper Slurm variables and thinks that the environment is not set properly, thus generating the warning: in reality, the instance of srun behind it will respect the setting you requested with your Slurm directives.

While there are workarounds for this, the best solution (apart from ignoring the message) is to use srun instead of mpirun: with this command, the Slurm environment is read properly and the warning does not appear.

- I'm trying to setup a conda environment but I get connection errors from the package repository. How can I solve this?

Due to the recent modifications to Anaconda's Terms of Service, the use of the anaconda public package repository on CINECA platforms is no more permitted without the purchase of a license for commercial use. As a consequence, the service was disrupted by Anaconda. Please create your environment using alternative channels, such as conda-forge, adding the --override-channels flag and add the channel to your conda environment (.condarc) via:

```
$ conda config --add channels conda-forge
```

As an alternative you may recur to python virtual environment and pip installations.

Performance:

- I have found performance problems, what should I do?

First, you should find a method to reproduce the problem and confirm after some tests that it is reproducible. After that, you should provide all relevant information to the Support Team by e-mail together with instructions on how to reproduce your tests. Support Team will investigate the issue and contact you as soon as possible.