

Startup Guide

This is a quick guide to getting started with the new cluster. Any further information or insights can be found in the extended help sections.

Cluster Access

To access the cluster, two things are required: that you have an Unibo account (of the type *name.surname00@unibo.it*) and that this account has been authorised for access by the OPH responsible person of your research sector. See [accessing the cluster](#) for the contact details of the OPH responsible persons.

Once authorised, one could access the cluster via ssh protocol. The required password is the one of your university e-mail. From the Linux terminal, type the following instruction:

```
ssh name.surname00@137.204.165.41
```

(you can use any other ophfeX address instead of 137.204.165.41)

Once logged in, you will land on the [Frontend](#), the workspace shared by all users that is employed to submit jobs and access datasets stored in memory.



Do not use the Frontend to execute long and demanding jobs, as this can lead to the shutdown of the entire system!

Setup the Environment

The first time you access the cluster, you must set up your working environment. Here are a few tips to help you manage your account and data correctly. In particular, attention must be paid to the correct use of data storage areas, see [Storage types](#). The OPH cluster currently has two main memory spaces with two different functions:

Working storage area: /home

At the first connection to the cluster, the system automatically generates a folder for you in the /home partition. This is a limited memory area and should not be used to store massive amounts of data. Typically, software codes and intensive-use documents or scripts that do not take up much space is saved under this area.

You can manage the data in this folder directly from the Frontend and to your liking, as long as you limit the storage space used.

Data storage area: /scratch

This is the main archive area, that must be used for large datasets and archives. This storage area is accessible from the Frontend but you need to create your personal folder manually. To facilitate access to the data in /scratch, it is recommended to use a symbolic link to a personal data folder. To do this:

Create your own folder in /scratch within the general folder of your research sector (see under research sector names [here](#)). For example, if you work in the *astro* sector, type:

```
mkdir -p /scratch/astro/name.surname00
```

Then create the symbolic link in the folder accessible from the Frontend:

```
ln -s /scratch/astro/name.surname00 run
```

in this way, when you access the Frontend, you will immediately find a folder named *run* from which you can access the data saved in /scratch (note that *run* is not a folder but a symbolic link). In any case, if you need to work with some big files (e.g. a large dataset) use symlinks from /home to the files in /scratch to get access.



The /scratch area cannot handle folders with a large number of files. Data folders in this area must be compacted into archives (e.g., .tgz or .zip).

Please pay attention to this policy: the stable number of files saved in /scratch shall not exceed a few thousand per user. Otherwise, the system becomes incredibly slow and unstable. Periodic checks on the number of files of each user are carried out automatically.

Note to student supervisors : Once you have created the data write folder for your students, you can request read and write rights to access the files through `setfacl -m u:name.surname0:rw /home/pathToFolder`, where `name.surname0` is your account name and `/home/pathToFolder` is the universal path to the folder you want to access.

Run a Job

The job executed (in parallel or serial) on the cluster is managed by [Slurm WorkLoad Manager](#). The submission of a job is done via a bash-type script, consisting of: the header with metadata for users, the execution settings (e.g. number of processors, memory, execution time), the modules and the executable to be run. See the section [Run a Job](#) for more details.



An example job script with comments can be downloaded [here](#) and adapted to personal needs: [runParallel.sh](#)

To run the job, the script has to be submitted by:

```
sbatch runParallel.sh
```

The output of the job execution is redirected into two files: `infoRun000` which contains the output and the job number in the name, `errRun000` which contains the error messages.

Job Monitoring and Management

Once a job has been submitted, it is possible to monitor the priority and progress status using the following:

- `sllurmtop`, displays the status of the cluster in a 'semigraphic' fashion
- `squeue`, displays queue status
- `scancel <job-ID>`, cancels the execution of a job with a given identification number (ID)

Additional information and functions can be found in the official [documentation of Slurm](#).

Problems and Troubleshooting

The cluster management requires quite a lot of time and energy at this stage. The management team kindly asks not to contact the technical administrators except for urgent matters or serious problems (which do not allow work to continue). Reports of malfunctions may be sent without guarantee of an immediate response.

- For information on accounting and cluster access problems, please contact the [reference person for your research area](#).
- For problems accessing memory and executing jobs on the cluster, contact the system administrators at `difa.csi@unibo.it`.

The technical administrators do not offer assistance for problems related to the use of Slurm (see [online documentation](#)) or related to your personal code or software.

Thank you for your cooperation and understanding.

From:
<https://apps.difa.unibo.it/wiki/> - Wiki DIFA

Permanent link:
<https://apps.difa.unibo.it/wiki/oph:cluster:startupguide>

Last update: **2025/02/10 11:29**

