

Storage

Except in particular cases, the accounts will have the following storage scheme associated with them:

Storage	Size	Description
/home	5 GB	It is recommended that you do not store data or code in this partition. When the maximum capacity is reached, you will need to move the files to /data .
/data	X TB	Partition with a " soft quota " on the hired storage. For large volumes it will be necessary to request it through the support email (support@hpc.iter.es). It is accessible in <code>/home/<user>/data</code> .
/scratch	--	Storage only available on request
/local/ <jobid>	300GB	On the local HDD of the node assigned to the execution by the scheduler. The content of this partition is deleted after the execution is finished When launching an execution to the queue manager, the folder is created in the indicated path. The jobid is accessible through the <code>\$SLURM_JOBID</code> environment variable
/lustre	X TB	High performance storage based on LUSTRE with higher bandwidth performance. It should be used as data storage in use by computational tasks, once available the results have to be transferred to home or data, as the permanence of data in this partition is not assured and it may undergo periodic clean-up operations. It is accessible <code>/home/<user>/lustre</code>

Any changes to this storage scheme will be notified.

 **Your *home* directory has a limit of 5GB.**

We recommend saving your virtual environment, application data and results in *data* or *lustre* (if you need high performance storage) directory !!

Backups

User data is protected up by the backup system of the storage system. The backup scheme is as follows:

- For **/home**, 6 hourly copies, 2 daily copies and 2 weekly copies.
- For **/data**, 6 hourly copies, 2 daily copies and 2 weekly copies.

For `/local` there are no backups, because it is a storage located on the hard disks of the compute nodes, focused on improving the performance of the jobs. Once the jobs are finished, the content of this partition is completely erased. The idea is that, once the execution of the job is started, the user copies the data is going to work with to this partition of the node and, before the job is completely finished, that is, in the same execution script, moves the results obtained to his **data** space.

For `/lustre` there is also no backup system for the same reason. It is a storage dedicated to optimising user runs, but this time, using the infiniband network that interconnects all the nodes and has much higher speeds than other networks. The idea is that the user copies the data of the executions here, and when they finish, moves the results back to his **data** space. In this case, we don't delete anything, but if for some reason a problem occurs and the data is affected or the user deletes the data by mistake, there are no backups to recover the data.