

# File systems

- [Overview of available file systems](#)
- [/work storage creation requests](#)
- [Security of user data](#)
  - [Scratch automatic cleanup](#)
- [Files belonging to a former user](#)
- [How to recover snapshots](#)

## Overview of available file systems

File system + mount point  Environment Variable for Access	Purpose	Type and Size	Backup and snapshots	Intended lifetime	Cleanup strategy	Quota	Available for
<code>/home</code>  <code>\$HOME</code>	Store source files, input data, small files  Globally accessible from login and compute nodes	GPFS, 100TB	backup to disks and snapshots	Account lifetime	No cleanup	100GB  per user	All users
<code>/work</code>  <code>\$WORK</code>	Collaboration space for a group Common software, result files and data sets  Globally accessible from login and compute nodes	GPFS, 100TB	Backup upon request, at cost price.  Snapshots	Group lifetime	No cleanup	50GB*  per group	All EPFL units ( <a href="#">upon request</a> ). Not available for bachelor /master students.
<code>/scratch</code>  <code>\$SCRATCH</code>	Temporary huge result files  Accessible from the frontend and compute nodes within one cluster.	Fidis: GPFS, 375TB  Helvetios: GPFS, 186TB  Izar: GPFS, 253TB	No backup, no snapshots	2 weeks	Automatic deletion of files and empty directories older than two weeks may happen without notice	Quotas may be in place to prevent runaway tasks filling the filesystem. A typical limit would be 1/3 or 1/2 of the total space.	All users
<code>/tmp</code> <code>/\${SLURM_JOB_ID}</code>  <code>\$TMPDIR</code>	Temporary, local file space for jobs on compute nodes.  Not available on login nodes.	local (node), between 64 and 512 GB	No backup, no snapshots	Job execution	Content is deleted after job end	no	All users

\* Space on work is charged for and, as such, the quota of the group depends on the amount of space purchased. There is no backup by default but a laboratory may request such a service. The price for backup will be the cost price to SCITAS.

## /work storage creation requests

The price for work is per TB/year, usually sold for 3 years, and can be found in our website: <https://www.epfl.ch/research/facilities/scitas/getting-started/prospective-users-howto/>

Each group can have 50GB for free.

- free 50GB, please contact us at [1234@epfl.ch](mailto:1234@epfl.ch)
- paid storage: <http://go.epfl.ch/scitas-storage>

## Security of user data

The contents of the *home* file-system are backed-up on a daily basis with a six month retention period. The backed up data are held at a separate physical location to the original data.

**The contents of the *work* file-system are not backed up by default**

**The *scratch* file-systems are not backed-up under any circumstances.**

**The *scratch* file-systems are only for short-lived files and, in the case of insufficient free space, files older than two weeks may be deleted without notice in order to ensure the usability of the cluster.**

# Scratch automatic cleanup

When a scratch filesystem reaches a certain level of use (normally 90%) an automatic cleanup procedure is activated. Starting with the oldest files present deletion takes place until the occupancy has been reduced to 70%. Only files less than 2 weeks old are sure not to be deleted by the cleanup mechanism.

## Files belonging to a former user

When a user no longer has a valid account on the clusters any files belonging to him on *home* are removed from the servers. They will remain on backup for 6 months after the user has left EPFL. The head of the laboratory is responsible for ensuring that these data are correctly managed. He or she can ask for a retrieval from backup.

The *work* file-system is divided by laboratory and, as directories of users who leave EPFL cannot be deleted by their colleagues, an intervention is required. Any "professional data" of users who leave EPFL should be passed on to "a person designated by the head of the unit" (translated from [LEX 6.1.4](#), Article 25]).

Please contact us to change the ownership of any files/folders to a person designated by the responsible of the group that can assess the value of the data and handle it accordingly (delete or archive).

**Once a user is no longer accredited, files belonging to them in *scratch* can be deleted without notice.**

## How to recover snapshots

A **snapshot** is the state of a system at a particular point in time. On our clusters, the *home* and *work* filesystems are snapshotted daily and snapshots are kept for one week. This is particularly useful in case a user removes a file by mistake.

Daily snapshots of the home and work filesystems can be found in `/home/.snapshots`, `/work/.snapshots`, respectively.