

scratch

The scratch file system is designed for performance rather than reliability, as a fast workspace for temporary storage. All CSCS systems provide a scratch personal folder for users that can be accessed through the environment variable \$SCRATCH.

Alps and Piz Daint provide a Lustre scratch file system mounted on /capstor/scratch/cscs and /scratch/snx3000 respectively, while other clusters share the GPFS scratch file system under /scratch/shared.

Soft Quota

No strict quotas are enforced on scratch, but a soft quota on the number of inodes (i.e. files and folders) is enforced on Piz Daint (/scratch/snx3000) while the scratch file system on Alps (/capstor/scratch/cscs) has a soft quota on both disk occupancy and inodes (files and folders), with a grace period to allow data transfer: note that when the grace time expires, the soft quotas will become hard limits if you are over quota, therefore you won't be able to write any longer on your personal scratch folder.

In order to prevent a degradation of the performance, **users with over 1 million files and folders will be warned at submit time and will not be able to submit new jobs** on Piz Daint. Alps (Eiger) users **need to check their disk space and inodes usage** with the command quota, that is available on the front end Ela and on Eiger User Access Nodes (UANs) as well. Currently the soft quotas are 150 TB disk space and 1 million inodes on Alps scratch file system, with a grace time of two weeks.

Cleaning Policy

Please note that a cleaning policy is in place on scratch: **all files older than 30 days will be deleted by a script that runs daily, so please ensure that you do not target this filesystem as a long term storage**. Furthermore, kindly note that in order to avoid performance and stability issues on the scratch filesystem, **if the occupancy grows above the critical limit of 60% we will be forced to ask you immediate action to remove unnecessary data**: if the occupancy continues to grow and we reach 80%, we will then need to free up disk space **manually removing files and folders without further notice**

As a matter of fact, when the occupancy goes above 80% the Lustre filesystem shows a performance degradation that affects all users. The same applies with large numbers of small files, since the Lustre filesystem is not behaving ideally when dealing with high volumes of small files.

Keep also in mind that data on scratch are not backed up, therefore users are advised to move valuable data to the /project filesystem or alternative storage facilities as soon as batch jobs are completed. Please, do not use the touch command to prevent the cleaning policy from removing files, because this behaviour would deprive the community of a shared resource.