

Storage on Curnagl

Where is data stored

The recommended place to store all important data is on the DCSR NAS which fulfils the UNIL requirement to have multiple copies. For more information please see the [user guide](#)

This storage is accessible from within the UNIL network using the SMB/CIFS protocol. It is also accessible on the cluster login node at /nas (see [this guide](#))

The UNIL HPC clusters also have dedicated storage that is shared amongst the compute nodes but this is not, in general, accessible outside of the clusters except via file transfer protocols (scp).

This space is intended for active use by projects and is not a long term store.

Cluster filesystems

The cluster storage is based on the IBM Spectrum Scale (GFPS) parallel filesystem. There are two disk based filesystems (users and work) and one SSD based one (scratch). Whilst there is no backup the storage is reliable and resilient to disk failure.

The role of each filesystem as well as details of the data retention policy is given below.

How much space am I using?

For the users and work filesystems the quotacheck command allows you to see the used and allocated space:

```
[ulambda@login ~]$ quotacheck
```

```
### Work Quotas ###
```

```
Project: pi_ulambda_100111-pr-g
```

Block Limits					File Limits					
Filesystem type	blocks	quota	limit	in_doubt	grace	files	quota	limit	in_doubt	grace

Remarks

work FILESET 304.6G 1.999T 2T 0 none | 1107904 9990000 10000000 0 none
DCSR-DSS.dcsr.unil.ch

Project: gruyere_100666-pr-g

Block Limits							File Limits				
Filesystem type	blocks	quota	limit	in_doubt	grace		files	quota	limit	in_doubt	grace
Remarks											
work	FILESET	0	99G	100G	0	none		1	990000	1000000	0 none
DCSR-DSS.dcsr.unil.ch											

User Quota

Block Limits							File Limits				
Filesystem type	blocks	quota	limit	in_doubt	grace		files	quota	limit	in_doubt	grace
Remarks											
users	USR	8.706G	50G	51G	160M	none		66477	102400	103424	160 none
DCSR-DSS.dcsr.unil.ch											

Users

/users/<username>

This is your home directory and can be used for storing small amounts of data. The per user quota is 50 GB and 100,000 files.

There are daily snapshots kept for seven days in case of accidental file deletion. See [here](#) for more details.

Work

/work/<path to my project>

This space is allocated per project and the quota can be increased on request by the PI as long as free space remains.

This space is not backed up but there is no over-allocation of resources so we will never ask you to remove files.

Scratch

/scratch/<username>

The scratch space is for intermediate files and the results of computations. There is no quota and the space is not charged for. You should think of it as temporary storage for a few weeks while running calculations.

In case of limited space files will be automatically deleted to free up space. The current policy is that if the usage reaches 90% files, starting with the oldest first, will be removed until the occupancy is reduced to 70%. ***No files newer than two weeks old will be removed.***

\$TMPDIR

For certain types of calculation it can be useful to use the NVMe drive on the compute node. This has a capacity of ~400 GB and can be accessed inside a batch job by using the \$TMPDIR variable.

At the end of the job this space is automatically purged.

Révision #10

Créé 21 octobre 2021 05:14:11 par Ewan Roche

Mis à jour 18 août 2022 13:55:12 par Emmanuel Jeanvoine