

Storage types

The DIFA-OPH cluster offers three different storage areas with different features and usage policies:

/home/

This is a storage area available from every node and this is the space that you access when you connect to the cluster frontend. This is meant to store non-reproducible data (e.g. source codes) and is regularly backed up. It should be **used for source files, compiling code**, or jobs that do not need a lot of space (see /home/work below). The /home is the area where your home folders are stored, as well as other shared areas such as /home/work/ and /home/web/ that are meant to be:

- **/home/work**: Used as a work area for jobs that do not need very big datasets or need to have lots of files in a single directory (not recommended since this might degrade performances for all users!). **Per-sector quota** of 1TB (soft) / 2TB (hard).
- **/home/web**:
 - This space is web-accessible at <https://cluster.difa.unibo.it> (and also from old <https://apps.difa.unibo.it/files/people/Str957-cluster> via an automatic redirect).
 - web access is read-only and it is not possible to create dynamic pages.
 - **Per-sector quota** of 1TB (soft) / 2TB (hard) except astro with 4.4TB/7TB.
 - Requires either index.html or .htaccess file with Options: +Indexes to be browseable

Technical characteristics

- NFS-mount via Ethernet (1Gbps which is not very fast but quite responsive);
- Quota limit to 50 GB (100GB as hard limit); check how much you're using with `quota -u`.

/scratch/

This is the fast Input/Output area to be used for direct read/write operations from the compute nodes. There is no quota in this area, but an automatic cleaning procedure is enforced on all files older than 40 days to avoid the disk space being exhausted, as this would make running jobs crash when trying to write outputs on disk. Therefore, once your jobs are finished you are recommended to archive the relevant data to /archive/ (see below) to avoid any data loss;



folders inside sectors areas **must** use the account as name or you won't get important mails ⇒ possible data loss.

Technical characteristics:

- Parallel filesystem, for quick (SSD-backed) access to the data you are working on;

- No quota, but **files older than 40 days will be automatically deleted** without further notice;
- No cross-server redundancy, just local RAID. Hence, if (when) a server (or two disks in the same RAID) fails, all data becomes unavailable.

/archive/

This is the main archive area to be used for large files, big datasets or archives; it is designed to be a distributed storage area for long-term data preservation. Data in this area should be stored in the form of compressed folders because the presence of a large number of small files will compromise its functionality. Every sector or project has a dedicated area with an associated quota on /archive/, and when the quota is exceeded no further writing is possible on the sector or project area.



folders inside sectors areas **must** use the account as name or you won't get important mails ⇒ possible data loss.

Technical characteristics:

- **To not be used to store a large number of small files**, this will compromise the functionality of the storage space eventually blocking all the reading/writing operations of the entire cluster.
- Max size for a single file is 8TB. When archiving big datasets please split them into sub-folders, compress and store them separately (preferably each compressed folder should be less than 1TB)
- Read-only access from compute nodes, read/write only from frontends and filetransfer nodes
- Quota imposed (both on file size and number of files) per sector, with extra space allocated for specific projects (in /archive/projects) or bought by individuals (in /archive/extra).
- Currently ACLs (setfacl) are not supported (cephfs exported via NFS-Ganesha does not allow to set/get ACLs)

Monitoring system of space usage

To allow users/sectors/projects to handle their /archive/ storage and avoid their sector going over quota, a monitoring and alerting system is now in place. Some of you may have already received alerting emails from the OPH cluster last night, around midnight. These emails are meant to inform you about your usage of the /archive/ system and alert you in case your sector/project is already over quota (i.e. using more than 100% of the allowed space or number of files, so that writing is already blocked) or close to the quota limit (i.e. using more than 90% of either the available space or number of files).

In particular:

- every sector/project reference person will receive an email on the first day of each month in any case (i.e. even if their sector is below 90% of disk usage) with the overview of their sector/project disk usage
- every sector/project reference person will receive an email when their sector/project is using

more than 90% of either disk space or number of files allowed for their quota; these emails will be sent with daily frequency until the sector/project disk usage is reduced below the 90% threshold

- individual users will receive an email only if their sector/project is using more than 90% of either disk space or number of files allowed for their quota; these emails will be sent with daily frequency to the users that are using (individually) more than 50% of the sector/project quota until the disk usage is reduced below the 90% threshold.

\$TMPDIR local node space (advanced)

Every node does have some available space on local storage in \$TMPDIR. This can be used to store temporary files that do not need to be shared between nodes. Being a local memory, latency is very low and with a limited capacity, usually around 200GB. It is automatically cleaned when the job terminates.

Technical characteristics:

- local space: not shared between multiple nodes, not even for a single multi-node job
- quite fast
- automatically cleaned when job ends

From:

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

Permanent link:

<https://apps.difa.unibo.it/wiki/oph:cluster:storage>

Last update: **2025/02/06 16:07**

