

Filesystem User Guide

Filesystems and data organization in the Cirrus HPC and HTC clusters

- [Disk Quota Policy](#)
- [Directories and filesystems](#)
- [How to transfer files between Cirrus and a local machine](#)
- [How to access files belonging to someone else or another project](#)
- [Lustre](#)
 - [Lustre Basics](#)
 - [Lustre best practices](#)

Disk Quota Policy

To prevent storage misuse and problems resulting from filesystems becoming full, Cirrus has implemented the following data usage policies and limits.

File system	Hard Limit Exceeded
/home	Once quota or number of files exceeded, new files will not be allowed. No automatic warning implemented. This may lead to login problems.
/data/unixgrp	Once quota or number of files exceeded, new files will not be allowed and all jobs will fail. No automatic checks implement and users are allowed to submit jobs and run then even no quota available.

- There are two conventions for defining digital storage capacities, namely, base 2 (1KB = 1024 bytes) and base 10 (1KB = 1000 bytes), and both are in common use. In Cirrus we follow Base 2, the following definitions are used when referring to digital storage capacities (filesystem quota limits, usage, etc.):

Conversion Table	--	--
1 KiloByte (KB)	2^{10}	1024 bytes
1 MegaByte (MB)	2^{20}	1048576 bytes
1 GigaByte (GB)	2^{30}	1073741824 bytes
1 TeraByte (TB)	2^{40}	1099511627776 bytes
1 PetaByte (PB)	2^{50}	1125899906842624 byte

- To know more about quota usage and filesystem please see the [directories and filesystems section](#)

“ **NOTE** Users should check in advance if they have enough quota for their jobs since this may lead to job failures and loss of CPU hours pledged to them.

Directories and filesystems

The Cirrus filesystems are based on the Lustre shared filesystem which is mounted in the compute nodes and in the submission/login nodes. Lustre is mostly suited to the storage of large/huge files.

Summary

CNCA-Lisbon cluster

Name[1]	Purpose	Availability	Quota [2]	TimeLimit [3]	Backup
/home/unixgrp/user	User default home	always	20GB/user	none	no
/data/unixgrp	Group data large data files [4].	always	100GB/group	temporary (will be deleted after 6 months)[4]	no
/exper-sw/unixgrp	Install software for groups	on request	20GB/group	permanent	no

CNCA-Vila Real cluster

Name[1]	Purpose	Availability	Quota [2]	TimeLimit [3]	Backup
/home/unixgrp/user	User default home	always	200GB/user	none	no
/scratch/unixgrp/user	User working directory	always	10TB/group	none	no

[1] Each user belongs to a Unix group: **unixgrp** - determined by the project they are attached to

[2] Increases to these quotas will be considered on a case-by-case basis.

[3] Time limit defines time after which a file is erased on the file system since its most recent access time, as defined by the file access timestamp.

[4] This folder has read+write permission to all users belong same group



NOTE: To know your unix group id type the command: **id**

Filesystems

/home

- Intended to be used for source code, executables and immutable data (input files etc), NOT large data sets.
- Globally accessible from all nodes within a system.
- Quotas apply. To know your actual quota limit:

```
lfs quota -uh username /home/unixgrp/username
```

```
Disk quotas for usr username (uid XXXXXX):
```

Filesystem	used	quota	limit	grace	files	quota	limit	grace
/home/unixgrp/username	11.6G	20G	20G	-	1395	10400	10400	-

NOTE: if your used \geq quota users are not allowed to write

/data

- Each project has a directory with pathname `/data/unixgrp` on each compute node. Users connected to the project have `rwx` permissions in that directory and so may create their own files in those areas.
- Globally accessible from all nodes within a system.
- Quotas apply. To know your actual quota limit:

```
lfs quota -uh username /data/unixgrp
```

```
Disk quotas for usr username (uid XXXXXX):
```

Filesystem	used	quota	limit	grace	files	quota	limit	grace
/data/unixgrp/	219.6G	300G	300G	-	81395	1102400	102400	-

NOTE: There are also limits on the number of files (files) that can be owned by a group (project) on `/data`.

/exper-sw

- Intended to be used for software, shared by group
- Globally accessible from all nodes within a system.
- Permissions to use based on request

How to transfer files between Cirrus and a local machine

“ **NOTE** Please note that the Cirrus Lustre storage is not designed for small files. Attempting to store or retrieve files less than a few megabytes will result in extremely poor performance for all users. If you wish to store lots of small files to massdata, please use a utility such as tar to combine them into a single, larger file.

rsync

To transfer files from your local machine to Cirrus, we recommend to use rsync. With appropriate options, rsync is resumable allowing the transfer to continue in case of drop mid-transfer. The recommended command line to use on your local machine is:

```
rsync -avPS src dst
```

where **src** his the path of the files on your local machine and **dst** will be the path to your destination on the appropriate host.

For example:

```
//1. Transfer files from local machine to cirrus.ncg.ingrid.pt:
```

```
rsync -avPS /home/myfiles user@cirrus.ncg.ingrid.pt:/data/unixgrp/my_folder/my_files  
//where user is your CNCA username, and unixgrp is the unix group identifying your project.
```

```
//2. Transfer files from Cirrus to your local machine:
```

```
rsync -avPS user@cirrus.ncg.ingrid.pt:/data/unixgrp/my_folder/my_files /home/mylocalfolder/  
//where user is your CNCA username, unixgrp is the unix group identifying your projec and /home/mylocalfolder/
```

it's a local machine folder.

sshfs

The use of *sshfs* is a secure convenient way to share external volumes as long the user is able to open a *ssh* session.

Suppose we want to share a directory */remote/dir* from some remote server named *server.remote.pt* on the local user interface *cirrus* at subdirectory *mydata*. If the remote username is *remuser*, execute on CIRRUS user interface:

```
$ mkdir mydata
$ sshfs remuser@server.remote.pt:/remote/dir mydata
$ df
Filesystem                1K-blocks    Used Available Use% Mounted on
server.remote.pt:/remote/dir 841572128 83471036 715328408  11% /users2/<group>/<user>/mydata
$ ls -l mydata
$ cp mydata/some_file local/workdir
```

when finish the user can unmount the remote volume with the command

```
$ fusermount -u mydata
```

The user will have on local mount directory the same permission to access the volume as on the remote server.

“ The mount point will be available only on the local node. For example, if the user mount the volume on the user interface this directory content will not be available on the workernodes. This is a convenient way to access and copy files between sites but it is not suited to be used within a batch job, particularly on a MPI batch job with multiple nodes.

How to access files belonging to someone else or another project

Users are allowed to change the default permissions file permissions in all folders they own. The default access rights to Cirrus filesystem is the following:

/home

- All files under /home are only accessible by the user.

/data

- All files under /data/unixgrp/ are rwx for all users belonging to that group.

Example on how to change permissions for user abc123 belonging to group xyz:

```
//1 listening the files on a given folder
```

```
ls -l
```

```
drwxr--r-- 1 abc123 xyz    4096 Nov 25 11:03 mydir
-rwxr-xr-- 1 abc123 xyz  4126231 Nov 25 15:42 myfile
```

- On this example the file owner (abc123) has read, write and execute (rwx) for myfile while the group (xyz) has read and execute (r-x) rights and the remaining of users has read access (r--).
- Now giving permissions for user def456 to read my files and folders:

```
//2 Changing file permissions
```

```
setfacl -Rm u:def456:rwx mydir
```

```
setfacl -Rm d:u:def456:rwx mydir (this options only applies to new files)
```

- Now giving permissions for a group abc:

//2 Changing file permissions

```
setfacl -Rm g:abc:rwx mydir
```

```
setfacl -Rm d:g:abc:rwx mydir (this options only applies to new files)
```

- for further details on file permission and attributes in linux filesystems [click here](#)

“ **NOTE** We don't recommend using `chmod o+r`, `chmod o+w`, or `chmod o+x` to give non-group members access to your project's files. Instead, you should use access control lists to limit the access privileges of specific users

Lustre

Lustre basic description of the Cirrus filesystem

Lustre

Lustre Basics

The latest Lustre documentation is available at [Lustre](#)

Lustre best practices

Distributed filesystems such as Lustre are ideal for HPC and HTC environments. In these environments the typical workload consists of large files that have to be accessed from many compute nodes with very high bandwidth and/or low latency. Therefore these filesystems are very different from the filesystems used on desktop computers or isolated servers. Although they excel at handling large files, they also have strong limitations when handling small files and access patterns more commonly found in enterprise and desktop environments. Operations that can be extremely fast on a workstation local disk can be painfully slow and expensive on a Lustre filesystem, affecting both the users performing those operations and eventually all other users. These best practices and recommendations are aimed to enable a smooth use of Lustre by minimizing or avoiding unnecessary or very expensive filesystem operations.

Avoid accessing attributes of files and directories

Accessing metadata information such as file attributes (e.g. type, ownership, protection, size, dates, etc) in Lustre is resource intensive and can degrade the filesystem performance, especially when performed frequently or over large directories. Minimize the use of system calls that access or modify these attributes such as `stat()`, `statx()`, `open()`, `openat()`, etc.

The same applies to commands like `ls -l` or `ls --color` that make use of above mentioned calls. Instead use a simple `ls` or `ls -l filename`.

Avoid using commands that access metadata massively

Avoid using commands such as `ls -R`, `find`, `locate`, `du`, `df` and similar. These commands walk the filesystem recursively and/or perform heavy metadata operations. They are very intensive on accessing filesystem metadata and can degrade badly the overall file system performance. If walking the filesystem recursively is absolutely required, then use the Lustre provided `lfs find` instead of `find` and similar tools.

Use the Lustre lfs command

To minimize the number of Lustre RPC calls, whenever possible use the `lfs` commands instead of the system provided commands:

- `lfs df` => instead of `df`
- `lfs find` => instead of `find`

Avoid using wild cards

Expanding the wild cards is resource intensive. Executing commands with wildcards on a very large numbers of files may take a very long time and badly impact the filesystem performance. Instead of using wild cards, create a list of the target files and apply the command to each of these files.

Read Only Access

Whenever possible open the files as read-only using `O_RDONLY`, furthermore if you don't need to update the file access time then open the files as `O_RDONLY | O_NOATIME`. If access time information is needed while performing parallel I/O then, let the master process open the files as `O_RDONLY` and all other ranks open the same files as `O_RDONLY|O_NOATIME`.

Avoid having a large number of files in a single directory

When a file is accessed, Lustre places a lock on the parent directory. When many files in the same directory are to be opened this creates contention. Writing thousands of files to a single directory produces massive load on Lustre metadata servers, often resulting on taking filesystems offline. Accessing a single directory containing thousands of files can cause heavy resource contention degrading the filesystem performance.

The alternative is to organize the data into multiple sub-directories and split the files across them. A common approach is to use the square root of the number of files, for instance for 90000 files the square root would be 300, therefore 300 directories should be created containing 300 files each.

Avoid small files

Accessing small files on the Lustre filesystem is very inefficient. The recommended file size is above 1GB. Reorganize the data in large files or use file formats such as **HDF5** or **NetCDF**. Alternatively if the total size of the files is small such as a few gigabytes then copy the small files to `/tmp` or to a scratch directory local to each compute node at the beginning of the job (do not forget to transfer and/or delete the files in the end). This approach can be combined with the use of archival tools such as `tar` to store the small files in one or more large tarballs that can be kept on Lustre more efficiently. When data is read-only, another alternative is to create a disk image and mount it read-only through loopback in each cluster node as described in (ref 4). Container tools such as `singularity` can also enable the use of loopback mounted disk images.

Avoid small buffer sizes

When reading or writing to files, Lustre performs much better with large buffer sizes ($\geq 1\text{MB}$). Aggregating small read and write operations into larger ones is highly recommended. MPI-IO Collective Buffering enables aggregate I/O.

Avoid small repetitive file operations

Avoid performing repetitive small I/O operations such as frequently opening files in append mode, write small amounts of data, and close the file. Instead open the file once perform all I/O operations and close.

Avoid multiple processes opening the same files at the same time

Multiple processes opening the same files at the same time can create contention and file open errors. Instead perform the open from a single process (master), or open the file read-only to avoid locking, or implement the open with a try and error approach with a sleep in case of error.

Avoid accessing the same file region from many processes

If multiple processes access the same file region at the same time, the Lustre distributed lock manager enforces coherency so that all clients see consistent results. Having many processes trying to access the same file region simultaneously can cause performance degradation.

In this case it might be preferable to either: replicate the file, split the file, perform the I/O operations from a single process rank or make sure simultaneous access will not occur. In any case it is recommended to keep the amount of file-open and file-lock operations in parallel as small as possible to reduce contention.

Appending to the same file from many processes

This is similar to the previous recommendation. If multiple processes try to append to the same file this will trigger locking and this can cause heavy contention. Ideally only one process should append each file.

File operations through the master

When accessing small shared files in a parallel job it is often more efficient to perform all the required operations through the master process and if needed broadcast the data to the other ranks, instead of accessing the same files from every rank. Similarly, if multiple ranks of a parallel job require information regarding a given file, the most efficient approach is to have the master process performing the required calls (e.g. `stat()`, `fstat()`, etc) and then broadcast the information to the other ranks. For an example see “Broadcast Stat” in (ref 3).

File Striping

In Lustre large files can be split into segments that in turn can be automatically spread across multiple storage devices. File striping is useful for parallel I/O over large files. For this to work the mount point in question must be backed by multiple storage devices (OSTs). The command `lfs df` can be used to verify if a given mount point is backed by multiple OSTs. To obtain file striping information for a given file use:

- `lfs getstripe filename`

The file striping can be set using the command `lfs setstripe`. If the command is applied to a directory it will define the default stripe settings for files created in that directory. A subdirectory inherits all stripe settings from its parent directory. If the command is applied to a file it will stripe that file across OSTs according to the specified settings.

- `lfs setstripe -s 64m -c 4 filename` => split filename in 64MB segments and spread it across 4 OSTs

If a large file is to be shared in parallel by several ranks (processes) having each rank working on its own portion of the file, then it might be useful to stripe the file in a number of segments equal to the number of processes, or a multiple of the number of processes.

For maximum performance I/O requests should be stripe aligned, this means that the processes accessing the file should do it at offsets that match the stripe boundaries. This minimizes the chances of a process having to access more than one segment (and more than one OST) to obtain the required data.

For small files striping should be disabled, this can be achieved by setting a stripe count of 1. The same applies if a large file will be accessed by a single process.

- `lfs setstripe -s 1m -c 1 mydirectory/smallfilesdir/`

Avoid installing software on Lustre

Software is usually composed of many small files and as previously mentioned accessing many small files on Lustre can place a huge load on metadata servers. The software compilations in particular can be better performed locally by copying or untaring the software to the `/tmp/$USER/` of a login server.

Furthermore, under high load I/O access to the Lustre file systems may block. If executables are stored in Lustre and access to the filesystem fails the executables may crash. Therefore, whenever possible, it is better to copy the executables to the `/tmp` of the cluster nodes.

To address these issues the INCD compute environment includes a special read-only caching filesystem (CVMFS from CERN) that provides scalable software distribution within and across compute clusters. Contact the INCD support if deployment of software on the INCD CVMFS is desired.

File locking with flock

If applications are carefully designed the use of flock is not strictly necessary since Lustre can keep non-overlapping writes consistent, and can handle concurrent file append operations. Depending on the Lustre version the use of flocks may have a performance impact.

Backups

Perform regular backups of your data to a safe location. INCD does not perform data backups.

References

These best practices have been compiled from the INCD experience and from the following sources:

1. https://www.nas.nasa.gov/hecc/support/kb/lustre-best-practices_226.html
2. <https://hpcf.umbc.edu/general-productivity/lustre-best-practices/>
3. http://researchcomputing.github.io/meetup_fall_2014/pdfs/fall2014_meetup10_lustre.pdf
4. <http://www.prace-ri.eu/IMG/pdf/WP245.pdf>
5. [https://opus.nci.org.au/display/Help/Lustre Best Practices#LustreBestPractices-7\)LimittheNumberofProcessesPerformingParallelI/O](https://opus.nci.org.au/display/Help/Lustre+Best+Practices#LustreBestPractices-7)LimittheNumberofProcessesPerformingParallelI/O)
6. <https://wiki.gsi.de/foswiki/bin/view/Linux/LustreFs>