

Set file striping

When you are writing large files to /scratch or /group, it is recommended to stripe these across multiple partitions (Lustre OSTs). When you have **very** large files (>1TB), you **must** stripe them as they can rapidly fill the OST and cause I/O issues for you and other users.

By default, files on /scratch and /group are not striped, so when a large file is written to a single OST, it can limit your I/O performance. You will need to set the file striping (or file layout) before writing the file, or migrate the file if it has already been written.

Step-by-step guide - creating and setting a new striped directory

For creating a new directory and setting the file layout for new files. This will only affect newly written files, see migration below to apply striping to already existing files.

1. Create a directory which will contain your files:

```
> mkdir /scratch/pawsey0000/data
```

2. Designate the directory for striping and the number of OSTs to stripe it over, 4 in this example:

```
> lfs setstripe --stripe-count 4 /scratch/pawsey0000/data
```

3. Newly written files in this directory will be striped over the specified number of partitions. This can be confirmed:

```
> lfs getstripe /scratch/pawsey0000/data
/scratch/pawsey0000/data
stripe_count: 4 stripe_size: 1048576 stripe_offset: -1
/scratch/pawsey0000/data/test.dat
lmm_stripe_count: 4
[...]
```

Step-by-step guide - striping an existing file through migration

For retroactively setting the file layout on existing files. Migration will copy the file in place and spread it across the specified number of OSTs.

1. Check the current layout on the existing file(s) to confirm they do not have the desired file layout:

```
> lfs getstripe /scratch/pawsey0000/data/test.dat
/scratch/pawsey0000/data/test.dat
lmm_stripe_count: 1
lmm_stripe_size: 1048576
[...]
```

2. Migrate the existing file(s) with the desired file layout:

```
> lfs migrate --stripe-count 4 --stripe-size 4m /scratch/pawsey0000/data/test.dat
```

3. Confirm that your changes have been set

```
> lfs getstripe /scratch/pawsey0000/data/test.dat
/scratch/pawsey0000/data/test.dat
lmm_stripe_count: 4
lmm_stripe_size: 4194304
[...]
```

Related articles

- [How to transfer data from IRDS at UWA to Zeus and Magnus](#)
- [Transferring to Pawsey via Clouster](#)

- [Simple guide to slurm job arrays on Zeus](#)
- [Set file striping](#)
- [Automate the restart of jobs from a checkpoint](#)