

Taiga User Guide

- [Lustre Documentation](#)
 - [Capacity to inode Ratio](#)
 - [Block Size](#)
 - [Default Stripe Size](#)
 - [Progressive File Layout \(PFL\)](#)
- [Taiga Access Methods](#)
 - [Lustre Native Mount](#)
 - [Globus](#)
 - [NFS](#)
 - [Taiga Allocation Access Guidelines](#)
 - [Direct Investment](#)
 - [HPC System Project Directory Access](#)
 - [Radiant Allocation](#)
- [Data Recovery](#)
 - [Snapshots \(Coming Soon\)](#)
 - [Backups](#)
- [Data Services Virtual Machines](#)

Lustre Documentation

Capacity to inode Ratio

Ratio: **1TB Quota to 1,500,000 iNodes**

An inode is a record that describes a file, directory, or link. This information is stored in a dedicated flash pool on *Taiga* and is of a finite capacity. To ensure that the inode pool does not run out of space before the capacity pools, this quota ratio is implemented. So for example if your project has a 10TB quota on *Taiga*, it will also have a quota of 15.0 million inodes. If you have any questions about this ratio please contact the storage team by opening a support ticket.

Block Size

File System Block Size: **2MB**

For a balance of throughput performance and file space efficiency, a block size of 2MB has been chosen for the *Taiga* file system. Larger file sizes help larger streaming data movement go faster, in general doing large I/O to file systems is encouraged if possible.

Default Stripe Size

Stripe Size: **1**

Number of Flash OSTs in *Taiga*: **24**

Number of HDD OSTs in *Taiga*: **28**

Lustre is capable of striping data over multiple OSTs to increase performance and help balance data across the disks. The default stripe for *Taiga* is set to 1 but this value is overridden as the file being written gets larger and larger; this behavior is determined by the PFL configured for *Taiga* which is described in the next section. If a user wants to see how many OSTs a file is striped across they can run the below command; the example shows a file that is striped across 2 OSTs.

Check Stripe Size

```
user@client# lfs getstripe -y /taiga/nsf/delta/abc123/testfile
lmm_stripe_count: 2
lmm_stripe_size: 4194304
lmm_pattern: raid0
lmm_layout_gen: 0
lmm_stripe_offset: 3
lmm_objects:
  - l_ost_idx: 0
    l_fid: 0x100000000:0x2:0x0
  - l_ost_idx: 1
    l_fid: 0x100010000:0x2:0x0
```

Progressive File Layout (PFL)

Taiga deploys a Progressive File Layout (PFL) that performs a couple key functions.

First, it allows us to keep the initial 64KB of every file on NVME flash; this increases performance for small file I/O by keeping it on faster media and keeps that noisy traffic off the spinning media that prefer larger I/O patterns. This helps improve the throughput for workloads doing large I/O by letting them have clearer access to the HDDs that make up the bulk of *Taiga*'s capacity.

Second, it allows us to dynamically set the stripe size of files so that the bigger a file grows the more stripes it gets. This helps improve the performance of the system, and helps keep the OST usage rates more balanced which leads to overall better system responsiveness. The stripe count of a file can be overridden by using either "lfs setstripe" or by using "lfs migrate" to change an existing file's stripe count; however these actions are very much discouraged. Users should use the system defaults except in rare cases.

PFL Implementation Details:

NVME Capture Size: **64KB**

Stripe for Files 0 bytes to 256MB: **1**

Stripe for Files 256MB to 4GB: **4**

Stripe for Files 4GB and Above: **All HDD** (Currently 28)

Taiga Access Methods

Lustre Native Mount

Taiga is available via a native Lustre mount on the below systems:

- Delta
- HAL
- Radiant
- HOLL-I
- NCSA Industry Systems

Native sub-directory mounts can also be requested for one-off machines via an SVC ticket, with the **STO: Taiga** component. SET will be providing a streamlined guide for Lustre client installation and configuration here (link coming soon). Requests for one-off mounts of *Taiga* will only be allowed for machines that have gone through the NCSA security hardening process.

Globus

Taiga is accessible via Globus at the endpoint name "NCSA Taiga" and the endpoint is open to the public internet for transfers from anywhere with another Globus endpoint. Authentication to the endpoint is handled by NCSA's CILogon service and requires two-factor via Duo. For shared collections, or other questions submit a ticket to help+globus@ncsa.illinois.edu. Information about Globus can be found at their site <https://www.globus.org>

NFS

Native mounts of *Taiga* using the Lustre client are greatly preferred for superior performance and increased stability, however sub-directories of *Taiga* can be mounted via NFS in cases where that is necessary. The NFS service is accessed via the HA *taiga-nfs.ncsa.illinois.edu* endpoint. The NFS endpoint is currently comprised of 4 servers that are directly connected to the 100GbE public-facing storage network and via redundant links to *Taiga*'s HDR Infiniband core fabric. If you need an NFS export please file a ticket with the **STO: Taiga** component flag and the storage team will assist in getting the exported provisioned.

Taiga Allocation Access Guidelines

Access to data on *Taiga* is ultimately governed by normal POSIX permissions and more specifically so, group membership. Some common scenario's are below:

Direct Investment

If you/your team has a direct investment on the *Taiga* system separate from a compute allocation, access to your area is governed by an LDAP group that was created during the allocation onboarding; this group name often (but not always) looks like *taiga_prj_XXXXX*. Your group's PI (and whomever they've designated) are able to add you to the LDAP group that controls access to your allocation via [NCSA's Identity Portal](#). Once you are a member of this group, you should have access to the data from any system that mounts *Taiga* across the center and the path will be the same.

HPC System Project Directory Access

Many NCSA compute environments leverage Taiga for their project space. For example, Delta's `/projects` mount maps to the full path `/taiga/nsf/delta/`. Accessing data that is within compute environment project areas requires membership to the compute/system allocation for the system in question (be it Delta, HOLL-I, etc.). The compute allocation PI should be able to get you added to that group which will gain you access to this data, an example of this LDAP group would be `delta_XXXX`. Once in that group you should be able to access this data from all NCSA compute environments you have access to (via the full path) and on the system the allocation originates from via `/projects/`.

Radiant Allocation

Radiant allocations on Taiga live underneath the full path, `/taiga/ncsa/radiant/` and are by default only available via NFS mount within Radiant instances associated with that allocation. If access of this data outside of Radiant is needed, please speak with the storage team by opening a ticket with the **STO: Taiga** component flag and someone from the storage team will assist you with your needs.

Data Recovery

Snapshots (Coming Soon)

Snapshots are run on *Taiga* once per day, and are retained for a **14** day period. The creation of new and removal of old snapshots is run automatically without intervention, it is not possible to recover deleted data from more than 14 days prior. Users wishing to recover data from snapshots should open a ticket with the storage team.

In general snapshots are designed to protect against and are useful for:

- Recovery in case of accidental delete
- Restore in case of file corruption due to application error
- Restore in case of encryption due to ransomware

Snapshots are **not** designed to protect against:

- Catastrophic file system hardware/software failure

Snapshots are also not designed to act as version control for software; all code changes should be kept in a git repository or similar version control tool.

Backups

Data on *Taiga* is not backed up; and is only a single copy. It is recommended to back up critical data to an allocation on the *Granite* tape archive system or on another system on which you have an allocation.

Data Services Virtual Machines

Some groups may want/need to have virtual machines that can directly access their data on Taiga for things like: data serving, data portals, light data analysis/indexing, etc. For these use cases groups should procure a Radiant allocation of the appropriate size to operate these services. Radiant is a separate service but is a scaleable way to give groups access to VM infrastructure that can access their data.