

Supercomputer NFS Scratch Allocation

The Joule Supercomputer offers two NFS Scratch filesystems, `/nfs/scr/1` and `/nfs/scr/2`, each providing an allocation of 3TB of disk space and 5Million file and directory entries, i.e. inodes. A total of 6TB of disk space and 10Million files and directories per user.

- NFS Scratch is scratch space, it is NOT BACKED UP, a failure of only one component of the scratch filesystems could result in data loss.
- NFS Scratch is subject to a 60 DAY PURGE POLICY. Data older than 60 DAYS WILL BE DELETED automatically from the NFS Scratch filesystems.
- Data results you wish to keep should be post-processed and/or moved to the long-term filesystem, `/nfs/long`, which is a redundant, backed-up filesystem.
- The long-term filesystem utilizes a redundant 10GbE link to all login systems to make migration of data as fast as possible. Long-term is also housed in a separate, physically segregated, data center with backup power generation to make storage of your data as safe as possible.
- If your scratch data consists of millions of files and directories, we ask that you compress this data using tar when moving it to long-term storage. This reduces your overall file count in the long-term filesystem and also allows for optimal performance of our backup processes. Users can tar directly from NFS Scratch to long-term on their login node, tar without compression is recommended, e.g. `tar cf /nfs/long/1/username/data.tar /nfs/scr/1/username/scratch_data`, replacing "1" and "username" with your login node number and username.

Failure to follow these policies could result in DATA LOSS or IMPACT THE STORAGE of other user data. The Joule 2.0 supercomputer is a resource shared among multiple groups and users, please be a good citizen to your fellow researchers.

Please sign below to request access to the NFS Scratch filesystem and acknowledge you understand and will adhere to these policies.

User: _____

Supervisor: _____

Signature: _____

Signature: _____