

Duke Office of Information Technology

Storage Usage

In January 2007, the DSCR storage system was upgraded to a [Network Appliance FAS-3020](#) filer. This new unit greatly expanded our existing storage capacity as well as offer a much needed boost in performance.

With the new file server, the storage allocation approach was also reorganized. Currently the server implements a storage quota for all groups. The intent of this quota is not to limit in any way a researcher's access to the disk space that they need, but rather to protect all users and prevent the "file system full" issues that we have seen in the past.

Additionally, the NetApps provides a feature called "snapshots" which allow us to retain several *on-line back-up copies of your data* in addition to regular tape back-ups. If you accidentally delete a file, you can immediately retrieve a copy from yesterday. Each group's root directory, e.g. /home/csem/, has a hidden subdirectory called .snapshot (note the period at the beginning).

- ? Each research group will now have all of their data/directories contained within their own separate file system.
- ? The amount of disk storage available to the group is limited by the size of this file system.
- ? If you fill up your filesystem space, two things may happen
 - o New jobs may fail due to not being able to write to their output files.
 - o Existing jobs that are running and writing data to existing file may have the data corrupted or truncated.
- ? The initial size of the filesystem is based upon that group's maximum usage over the last 12 months, padded upwards. New groups will be allocated an initial partition of 100GB, which will be grown as their needs increase.
- ? The amount of available filesystem that you have may be seen by using the command: " df -h ~/ " (without the quotes). Example:

```
-bash-3.00$ df -h ~/ Filesystem Size Used Avail Use% Mounted on filer1:/vol/csem
100G 5.1G 95G 6% /home/csem
```

- ? The filesystem size limits are for the entire group. There are no quotas for individual users. We ask that groups police their own storage usage locally.
- ? The size of your filesystem may be increased on an as-needed basis. If you need more disk space to complete your work, the PI for the group should send an email to hpc-support@duke.edu and make the request.
- ? While we make every effort to adequately provision storage for our users, space is a limited commodity. We would ask groups that expect to have large long-term storage needs (0.5TB or greater) to contact us about procuring your own additional disk space. The new Network Application files gives us the ability to much better

manage our storage resources and we are now in a position to accept contributions of storage as part of the cluster participation model. Email us for more details.

Duke Office of Information Technology - www.oit.duke.edu - (919) 684-2200 - help@oit.duke.edu