

Storage

Requesting Research Storage Solutions

Tufts TTS provides a networked storage solution to faculty researchers who request it based on their identified needs. This solution is often used in support of grant based projects that produce or require large data sets, and these data sets need both to be accessed by multiple users and to be backed up. Such networked storage solutions usually supports a small lab where multiple persons may interact with research data from computers at different locations. Another possibility is to request that the storage be made available to one or more accounts on the research cluster. Note, please allow up to 2 weeks for requests to be processed.

Estimation of storage needs

Please take some time in estimating your storage needs near term and longer term projections. For example, consider the following hypothetical experiment as follows:

- number of subjects including potential drop-outs (N)
- number of samples per subjects (n1)
- duration of acquisition in seconds (T)
- data rate acquisition in Hz (F)
- size of a sample in megabytes (s)

Size of dataset: $N*n1*T*F*s$

Example: If I recruit 20 subjects for three 5-min trials acquiring black and white images of 1MB each at 60Hz will give me a data set of $20*3*5*60*60*1 \sim 1000GB$ or 1TB.

For additional information, please contact Lionel Zupan, Associate Director for Research Computing, at x74933 or via email tts-research@tufts.edu.

Please work with your IT support person (**FSP**) to request or update NFS(cluster) or CIFS(Windows) **storage** request. Note, Windows storage is not mounted on the cluster and cluster NFS storage is **not** re-exported to desktops.

Storage User [guide](#)

Cluster Specific Storage

As of Oct. 10 2013, all cluster accounts are created with a fixed **5 GB** home directory disk quota.

In addition, a directory is automatically created in filesystem **/scratch** on the head node and each compute node. The directory is named with your Tufts UTLN; such as **/scratch/utln/**. There is no quota and no backup for files located there. Typically ~100 + gig is available. In addition this file system is subject to automated cleaning based on a moving window of 28 days. You may use this for temporary storage supporting your research computing needs. Note: this storage is not available to your desktop as a mounted file system, but you can access it via your ssh login or through a file transfer program such as **WinScp**.

For additional temporary storage beyond what is offered in /scratch, Tufts TTS provides a 2TB file system called **/cluster/shared/**. This file system is subject to automated cleaning based on a moving window of 28 days. There is no quota and no backup for files you located here. Every created account has access to this storage. The naming convention is the same as **/scratch/** storage and thus would be **/cluster/shared/utln/**. Please work out of your named directory. Note, **/cluster/shared/** file system is mounted on all nodes, but /tmp and /scratch/ are local to a compute node and the login node.

Another specialty storage solution is called **/scratch2**. The cluster has a limited number of nodes with local hard disk partitions dedicated to this filesystem with ~200gig of scratch space. Access to this is by request and for situations where NFS based storage is not ideal. See note below for usage example.

Another available file system is **/tmp**. This file system is used by applications and the operating system for temporary storage. **Do not use this for temporary storage**. If you fill it up, the node may stop working and affect all processes on that node. Note, **/tmp** on the headnode is a separate filesystem, however on compute nodes it is **not!!!!** Also note that it is under a 10 day cleaning rule.

Finally, filesystem **/cluster/tufts** is a NFS network storage file system available by request. This storage is only mounted on the cluster. Use this link to request or update an existing **storage** request. See below for details.

Cluster Storage Summary Table

FileSystem	Quota	Backups	Max. Capacity	Cleaning Window	Availability
/tmp	no	no	10 gig	10 days	node specific
/cluster/shared/utln	no	no	~2 TB	28 days	everywhere

/scratch/utln	no	no	~915 gig	28 days	login head node
/scratch/utln	no	no	~20 gig	28 days	compute nodes
/cluster/home	yes	yes	500meg/account	none	everywhere
/scratch2/utln	no	no	~ 200gig.	28 days	compute nodes, by request only
/cluster/tufts/	yes	NFS research storage, determined by user	amount requested by user	no	all nodes

How can I verify that my requested storage is mounted on the cluster?

Use the df command.

Since cluster node logins are unnecessary, how do I see my data on a particular node /scratch area?

From the headnode you may list files on a compute node(eg. node14) as follows:

```
> ls -l /cluster/scratch/node14/utln/
```

Where utln is your cluster username.

How do I reference my data that is located on the login node /scratch area from my process on a compute node?

The path you need to reference is: /cluster/scratch/tunic6/utln/

Where utln is your cluster username.

Sometimes I notice some latency in my program reading or writing data to my mounted storage, what could be going on?

One possibility is the NFS stale file handle issue. It is often a result of a race condition resulting in change in process ownership of file details. If there are other processes using the same data and on or about the same time frame, this is often the setting for this issue. Sometimes all it takes is just the right amount of load and access contention. One option is to move the job's data to a local disk thus avoiding network I/O traffic which is always slower than local disk. Another possible option is to have many fewer jobs in the RUN state. Contact cluster-support@tufts.edu if you have further questions.

Storage and Data Transfers

Getting your data to cluster mounted storage is normally straight forward. You may use any file transfer program such as WinScp that supports sftp or scp protocols. However if you have a large amount of data(100s of gigabytes or larger) located outside of the Tufts network domain the task may be non trivial. This can be a result of network traffic, hardware bottlenecks, and other issues beyond your control, that may result in a dropped connection; all resulting in taking a lot longer time than you thought. For transferring large amounts of data to the cluster, make sure you have the enough storage to receive it and consider using the file transfer method **rsync**. **rsync** is available on the cluster as a command line tool. Documentation for **rsync** is available as a cluster man page or on the web. A simple example: (note, that . is not a speak on your screen, it is required to indicate the target path location)

How can I transfer via rsync some files of a known pattern, for example?

```
> rsync -av ftp.ncbi.nih.gov::blast/db/nr.*.tar.gz YOUR_DEST_DIR
```

Note: embedded passwords are sometimes needed.

```
> rsync -avz yourusername@remotehost:/path-to-your-data .
or submit as a batch job should it take some time:
> bsub rsync -avz yourusername@remotehost:/path-to-your-data .
```

Suppose I want to copy data via unix command scp from my bash script that is running on some compute node to the /scratch/utln storage area of the login node. How do I reference it?

Note, your utln username is needed and since the destination is a local disk that is not mounted **everywhere**, the use of **scp** is needed as well.

```
> scp filename tunic6.uit.tufts.edu:/scratch/utln
```

Should I use scp instead of cp or mv to move data around my filesystems?

On the cluster many filesystems that you may have access to are NFS mounted on all nodes. From your shell, you can use **cp** and **mv** in the normal manner. However, for all but one exception, do not use **scp** to copy your data between filesystems. It is extra overhead that is not needed.

I am using the Tufts Box.com cloud storage option, can I transfer files that way to the cluster?

Yes. Firefox is available on the cluster. You must have X11 graphics support to receive the Firefox displays. Note that your cluster home dir. has a storage quota limit of 500meg, while your Box.com account is likely much much more than that. To transfer large files you should use the /cluster/shared/your_utln directory or one of the other options above.

How does one make use of nodes with /scratch2 local disk storage?

Access to this storage is by request. Please make this request via cluster-support@tufts.edu.

If you submit a job with the following options, LSF will place a job on nodes with /scratch2 partitions. For example, to request at least 40gig of storage for a job to run in the long_public queue try:

```
-bash_3.2$ bsub -q long_public -R "scratch2 > 40000" ./your_jobname
```

Other queues are possible as well. Note, the storage argument is in megabytes.

Cluster Mounted Storage "Ownership" and Succession

Often a request for mounted storage is to support an ongoing project. While the storage ownership is most often a faculty member, it is the project members that are often using and coordinating use of the storage. When project members status change due to graduation, etc., some account adjustments may be needed to secure access and affect appropriate file permission changes. TTS/RGTS requires the approval of the owner to make these changes. Please send a brief email to cluster-support@tufts.edu requesting what changes are needed.

If you own a research storage share that is mounted on the cluster (not your home dir.) and you leave Tufts, the following steps are taken prior to storage removal:

1. a final full backup will take place
2. taking the volume offline for two weeks
3. then deleting the storage

Research Database (HPCdb) Node

Currently there is no database service attached to the cluster for research use. However, cluster users may inquire about access to a [MySQL](#) database for supporting their research computing needs. Under some circumstances we may be able to assist in finding a host that cluster nodes may query remotely. This option is not designed for heavy I/O requests, large volume client traffic, or large storage.

Requests are treated like software requests. Please reference the Software Request Policy statement in this document.