

Tufts HPC Cluster New User Guide

Welcome to the Tufts High Performance Compute(HPC) research cluster.

Command Line

If you are using a Mac, a Linux, or PuTTY/Cygwin ssh/SecureCRT/.etc on a Windows machine:

Login:

Please do not run any software on the login nodes.

Shell environment (default: bash):

```
ssh utln@login.cluster.tufts.edu
```

or

```
ssh utln@login.cluster.tufts.edu
```

With GUI (graphical user interface):

```
ssh -YC utln@login.cluster.tufts.edu
```

or

```
ssh -X utln@login.cluster.tufts.edu
```

Now you are on a **login node** of the cluster and in your Home Directory.

Storage:

Your home directory (max 5GB) should be /cluster/home/utln

To see your home directory quota:

```
$quota
```

Your research storage (from 50GB and up) should be /cluster/tufts/yourlabname/

To see your research storage quota:

```
$df -H
```

Note: research storage is shared within a research group.

Data Transfer:

Please use Xfer for transferring files between the HPC cluster and your local machine.

Scp [from] [to]

Do this on your local machine.

Transfer for your Home Directory:

File Transfer:

```
$scp utln@xfer.cluster.tufts.edu:/cluster/home/utln/[FileName] [ToWhere]
```

```
$scp [FromWhereOnYourComputer]/[FileName]
```

```
utln@xfer.cluster.tufts.edu:/cluster/home/utln/[DirectoryName]
```

Directory Transfer (with content):

```
$scp -r utln@xfer.cluster.tufts.edu:/cluster/home/utln/[DirectoryName] [ToWhere]
```

```
$scp -r [FromWhereOnYourComputer] utln@xfer.cluster.tufts.edu:/cluster/home/utln/[DirectoryName]
```

Transfer for your Research Storage:

File Transfer:

```
$scp utln@xfer.cluster.tufts.edu:/cluster/tufts/yourlabdir/yourdir/[FileName] [ToWhere]
```

```
$scp [FromWhereOnYourComputer]/[FileName]
```

```
utln@xfer.cluster.tufts.edu:/cluster/tufts/yourlabdir/yourdir/[DirectoryName]
```

Directory Transfer (with content):

```
$scp -r utln@xfer.cluster.tufts.edu:/cluster/tufts/yourlabdir/yourdir/DirectoryName] [ToWhere]
```

```
$scp -r [FromWhereOnYourComputer]
```

```
utln@xfer.cluster.tufts.edu:/cluster/tufts/yourlabdir/yourdir/DirectoryName
```

Run Jobs:

Should always run your jobs on compute nodes.

You can start a simple interactive session with typing from your login node command line:

```
$ srun -p interactive --pty bash
```

This will get you to a compute node in the same bash environment as your login node.

or be more specific to specific # of cores or memory

or with X11 (if you need GUI for display)

```
$ srun -p interactive --pty --x11=first bash
```

You will see your location has changed from [utln@loginXXX] to [utln@<nodename>].

In the future, you can run your test cases here with no trouble before you submit larger jobs.

When you are done with the interactive session, just say "\$exit" and you will be back to the login node again.

There are more options for the interactive sessions, you can find more information on our wiki <https://wikis.uit.tufts.edu/confluence/display/TuftsUITResearchComputing/Slurm>, the "QuickStartSlurm" is very helpful if you are not ready to go through all the documentations.

Note: interactive sessions do have limits, the time limit for every interactive session is 4 hours.

Once your time is up, you will be logged out from the current compute node you are at, but you are still on the login node of the cluster.

If you want to submit a batch script, which will allow the jobs run in the background. Here is a sample script that you can modify it to fit your needs:

```
#!/bin/sh
#SBATCH -J Your_Job_Name
#SBATCH -p batch #run on batch queue
#SBATCH --time=0-1:00:00 #day-hour:minute:second
#SBATCH -n 12 #request 12 cores
#SBATCH --mem=2000 #request 2000MG memory
#SBATCH --output=PName.%N.%j.out #your output file
#SBATCH --error=PName.%j.err #your error file
[commands_you_would_like_to_exe_on_the_compute_nodes]
```

For more information about Slurm and batch scripts, please check out:

<https://wikis.uit.tufts.edu/confluence/display/TuftsUITResearchComputing/Slurm>

Modules/Applications/Software:

To check what's already installed for you on the cluster:

```
$module avail
```

To check if a particular software is installed on the cluster for you:

```
$module avail softwarename
```

To load some software/application/module onto your account:

```
$module load softwarename
```

To unload some software/application/module from your account:

```
$module unload softwarename
```

To unload all software/application/module from your account:

```
$module purge
```

To check what's installed on your account:

```
$module list
```

OnDemand (Tufts HPC cluster web interface)

Please go to <https://ondemand.cluster.tufts.edu/>

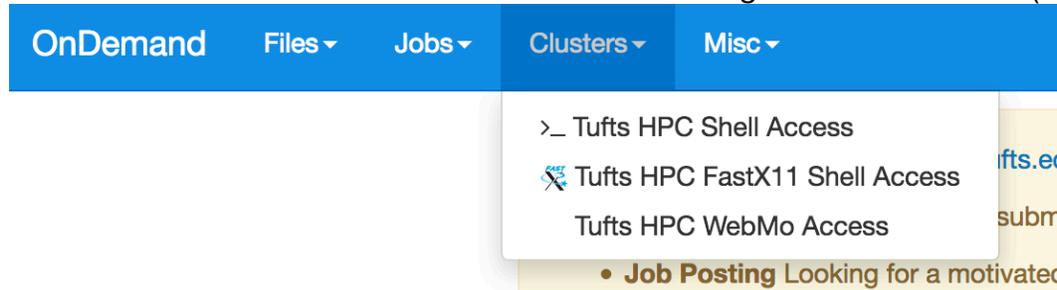
Use your Tufts UTLN and password to login.



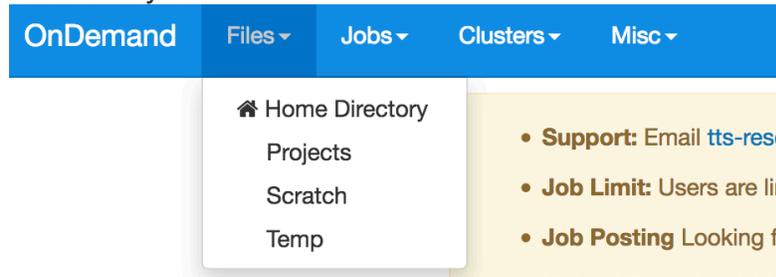
In Clusters, you can start a shell access to the HPC cluster.

“Tufts HPC Shell Access” = `ssh utln@login.cluster.tufts.edu`

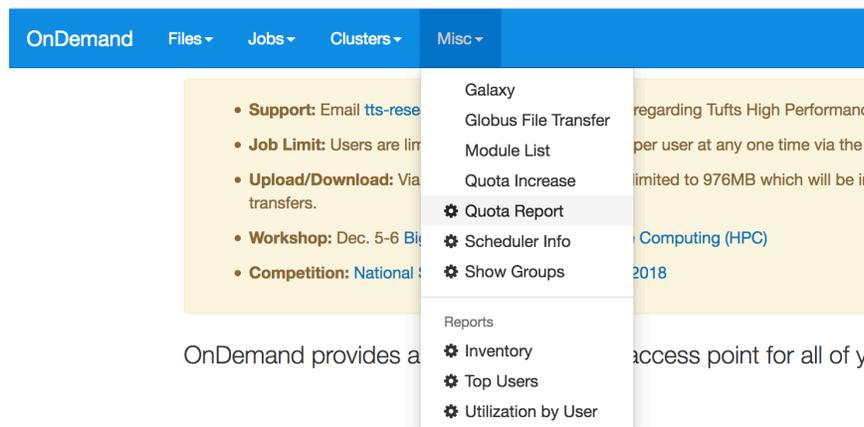
“Tufts HPC FastX11 Shell Access” = `ssh -X utln@login.cluster.tufts.edu` (with GUI)



Under Files, you can transfer your files between local machine and the Cluster, either you home directory or your research storage(Projects). Note, if you don't see your research storage under Projects, don't panic, please let us know, we will make sure it's visible to you.



To See your quota, you can go to “Misc”/ “Quota Report”



Message of the Day