

## Review of using the Tufts Research Cluster

- Cluster consists of many different nodes (think of each node as a separate PC). There are groups of similar nodes and specialty nodes. Specialty nodes are login and xfer nodes (to login and to transfer files).
- Use SSH to login to cluster. Use SCP and RSYNC to transfer files (full transfer vs incremental changes).
- DO NOT run things on the login node. This is a basic hardware node not designed for computation. Same with xfer node.
- Use the login node to request resources on other nodes, then login to these other nodes to run code.
- Typical command sequence. First request resources  
salloc -N1 -c4 -t 100 -mem 10G -p batch  
(one node, 4 cores, 100 minutes, use max 10G of total RAM)  
Then, figure out which node you got allocation on:  
squeue -u username  
Then, ssh to that node hostname and navigate to the directory containing your code.

## Tufts Research Cluster Layout

