

# Architecture

## Tufts Linux Research Cluster

Tufts Technology Services **TTS** provides a wide array of services in support of Tufts research community. High Performance Computing(**HPC**) hardware from Cisco and IBM is used to create the cluster. The hardware complement includes Cisco blades, IBM M3 and M4 iDataplexes, nVidia GPUs and a 10Gb/s interconnect Cisco network.

**IBM m4 nodes:** Intel(R) Xeon(R) CPU E5-2670 0 @ 2.60GHz  
**Cisco nodes:** Intel(R) Xeon(R) CPU E5-2660 v2 @ 2.20GHz  
**IBM m3 nodes:** Intel(R) Xeon(R) CPU X5675 @ 3.07GHz

By late Dec. 2014 there is approximately 163 compute nodes. Total **slurm** managed cpu/core count is ~4000+ and a peak performance of roughly 60+ Teraflops. In this HPC environment, TTS also provides researchers with access to commercial engineering software, popular open-source research software applications and tools for bioinformatics and statistics. Secure networked storage for research data (400+ TB CIFS desktop on NetApp appliances and 511 TB GPFS cluster storage ) is available.

Each cluster node has 12, 16 or 20 cores using three different Intel CPUs. Compute node memory ranges from 24 to 384 gigabytes of memory.

GPU computing is supported by 12 nVidia GPU models:

**K20, M2070, M2050**

The Linux operating system(RedHat 6.7) on each node is configured identically across every machine. In addition there is a login node and a file transfer node supporting the compute nodes. Client/user workstations access the cluster via the Tufts Network using **ssh** based connection client software. Remote ssh access for researchers is also supported. The user login node has an additional network interface that connects to the compute nodes using private IP addressing via 10Gig network hardware. This scheme allows the compute nodes to be a "virtual" resource managed by **slurm** job queuing software. This approach also allows the cluster to scale to a large number of nodes thus providing the structure for future growth. The login node of the cluster is reserved for the use of compilers, running shell tools, and launching and submitting programs to compute nodes. The login node is not intended for running research programs, or for general computing purposes, and all jobs are to be submitted to compute nodes using **slurm**.

A separate file transfer node, [xfer.cluster.tufts.edu](http://xfer.cluster.tufts.edu), is also provided to accommodate large data transfers.

See the Conceptual diagram and layout of cluster nodes .