Research Grant Information

Institutional Computing Resources

Grant project support may leverage the research computing capacity of Tufts University to provide a dedicated computing cluster and access to Tufts research storage area network (SAN).

The Tufts University Linux Research Cluster is comprised of 103 IBM Linux systems (1032 compute cores) interconnected via 10Gig network. Each cluster node has eight or twelve 2.8Ghz+ Intel cores and 16,24,32,48 or 96 gigabytes of memory. As of Sept. 2011 one compute node was provisioned with two Nvidia Tesla M2050 GPU processors and by Jan. 2012 an additional GPU node will be installed. The Linux operating system on each node is RedHat 5 configured identically across every machine. In addition there is a login node and a management node supporting the compute node array. Client/user workstations access the cluster via the Tufts Network or remotely with ssh. The user/login node has an additional network interface that connects to the compute nodes using private non-routable IP addressing via the 10 Gigabyte hardware. This scheme allows the compute nodes to be a "virtualized" resource managed by the job queueing software LSF, and abstracted away behind the user node.

In addition to the research cluster, this project will also leverage the Tufts research networked storage. Currently, Tufts University offers a total of 192 TB of storage capacity on a Network Appliance(NetApp) SAN to help researchers safely store their data. This storage appliance is backed up on a daily basis with up to 1 year of back-ups kept off site at any moment. For this project, storage will be provisioned for free to up to 500GB. Above 500GB, storage will be provisioned at a fixed recovery rate, and financial details will be finalized at time of award. Provisioned research storage space will also be available for mounting on the Tufts research cluster to leverage the computing power of the latter for data analysis. Details regarding the Tufts research cluster and associated research storage can be found at http://go.tufts.edu/cluster .

Network Security in relation to Research Cluster, Storage Services

Tufts University maintains a distributed information technology environment, with central as well as local aspects of overall planning and control. Tufts' information security program is structured in a similar manner. Operationally, Tufts central IT organization (TTS) and each local IT group maintain standards of quality and professionalism regarding operational processes and procedures that enable effective operational security. For TTS managed systems, the emphasis is on centralized resources such as administration and finance, telecommunications, research computing and networking, systems and operations as well as directory, email, LDAP, calendaring, storage and Windows domain services. TTS also provides data center services and backups for all of these systems. Additionally, a large number of management systems (for patching), anti-virus and firewall services are centrally provided and/or managed by TTS. Within TTS, processes and procedures exist for managed infrastructure changes, as change control is required for all critical central systems. Tufts University provides anti-virus software for computers owned by the University, and makes anti-virus software available at no charge for users who employ personally owned computers in the course of their duties at the University.

Tufts Research Storage services is based on a Network Appliance(NetApp) storage infrastructure located in the Tufts Administration Building(TAB) machine room. Provisioned storage is NFS (Network File System) mounted on the Research Computing Cluster for project access. NFS exports are not exported outside of TTS managed systems. Tufts Research Computing Cluster is also co-located within TAB's machine room. Network based storage connected to the cluster is via a private(non public) network connection.

Access to the Tufts IP network itself is controlled via MAC address authentication which is performed via the Tufts login credentials and tracked in the TUNIS Cardinal system; this system uses an 8 character password scheme. A switched versus broadcast hub network architecture is in place limiting traffic to just the specific ports in use to transport data from source to destination. Access to Tufts LAN network resources is controlled via Active Directory where applicable or LDAP, which requires the user to authenticate each time a system joins the domain. All of these controls are identically implemented on the wired as well as wireless Tufts networks.

Both Research Storage and linux based Cluster Compute server operating systems are kept current via sound patch management procedures. For example, PC's owned and managed by Tufts are automatically patched via the Windows Server Update Service. All other computing platforms are required to be on a similar automated patching schedule. From an operational standpoint, most central and local systems are maintained and managed using encrypted communications channels. For UNIX/linux servers, SSH is utilized; on Windows, Microsoft Terminal Services is utilized. User access to cluster services is via SSH and LDAP. No direct user login access to central Research Storage services is possible.

Additional user related Cluster information can be found here: http://go.tufts.edu/cluster

All devices and users are subject to the Tufts Acceptable Use policy found on **TTS** website:

How to reference the computing and storage resources for grant purposes.

Please reference this resource as: Tufts High-performance Computing Research Cluster