

# New Cluster

## May 2014 Announcement

Every 3 years, the Tufts high-performance computing cluster environment is partially refreshed to offer increased capacity. We did so in 2008 and 2011, and we have just completed our competitive bid for our 2014 refresh purchase. As you may know, we currently have 4 generations of IBM hardware and we are proud to announce that we will soon welcome a new generation of Cisco UCS hardware.

Here is how we are planning to retire our oldest IBM hardware, install our newest Cisco hardware, and integrate both IBM and Cisco environments while allowing for a smooth and non-disruptive transition.

First, we will retire the oldest IBM hardware on June 15, 2014. This represents a removal of ~500 cores or about 25% of the current cluster capacity. We have adjusted LSF queues to help accommodate the decreased compute capacity while maintaining a similar quality of service from June 15 thru August 2014. We will refer below to this IBM cluster with retired nodes as the "current cluster".

The new Cisco hardware will be delivered in the second half of June, then installed and configured over the Summer. By mid-September, this second Cisco cluster with the new 1000 cores will be made available in production to our users to form the foundation of our new Tufts high-performance computing environment. We will refer to this Cisco cluster below as the "new cluster".

The current cluster and the new cluster will be available concurrently during the 2014 fall semester. This will allow for a smooth user transition with ample time for testing. Finally, on January 1, 2015 the current cluster hardware will transition to the new cluster production environment for a total of 2680 cores (w/ 1680 IBM and 1000 Cisco cores) as compared to 2110 today. Due to a significantly different and newer architecture, the 1000 Cisco cores are a lot more powerful than twice the computing power of the 6-year-old 500 cores being retired. In practical terms this refresh represents the largest increase in compute power for the Tufts HPC community.

As part of this hardware refresh, we will also be replacing our LSF scheduler with slurm on our new cluster. We made this decision as did many other high-performance computing centers across the US since slurm is considerably cheaper while offering similar features as LSF and while being quickly improved by a large user community to adapt to newer HPC paradigms such as accessing cloud resources.

To allow for a smooth transition to the new slurm scheduler, LSF will remain on the current cluster until December 31, 2014 at which time it will also be retired and replaced by slurm. We are creating this 3-month+ transition period from mid-September to December with both LSF and slurm available to allow users to easily switch from the former to the latter. Workshops, tip sheets and training will be organized during the transition period to make sure no Tufts cluster user is left behind.

## Update Announcement 2011

TTS is pleased to announce the completion of the Tufts research cluster upgrade project. The new High-Performance Computing (HPC) research environment is in production bringing more than 1000-cores of computing power to the Tufts community. Over the past three years, TTS has observed an increasing demand for our HPC research cluster. In anticipation of the ongoing need for additional resources, we began a project to increase resources more than three fold.

In addition to an increase in capacity, we have made the following changes:

- The cluster network interconnect has changed from InfiniBand to 10 Gigabit Ethernet to allow our high-bandwidth communication needs to grow in conjunction with our Data Center.
- Several computing nodes will have 96GB of RAM per node (or 8GB per core) for memory-intensive jobs.
- An experimental GPU research environment that can deliver up to 1 Teraflops of computing power in double-precision floating point mode has been added for exploration of alternate parallel computing methods.

The upgrade has been both extensive and complex. Following the acquisition, which started in November 2010 and was completed in May 2011, three additional phases during the month of August were needed to transition from the old to the new HPC environment:

**Phase 1**, completed on August 13, brought the newly acquired IBM hardware into a production-ready configuration. Once the new cluster was fully configured and running, all cluster users were switched from our legacy cluster to this new environment in a transparent manner. At the end of Phase 1, the new Tufts HPC research cluster environment totaled 46 new iDataPlex servers for 552 cores (@ 3.06GHz).

**Phase 2**, completed on August 17, consisted of powering off the legacy hardware to change the cluster network interconnect from Infiniband to 10GigE. As the new interconnect hardware was fully configured and tested on the legacy hardware and the newly-configured legacy nodes were integrated in the new cluster environment, the number of cores available to Tufts researchers increased accordingly. During this process, priority was given to faculty-contributed nodes.

As of August 29, 2011, the new cluster environment is at 77% of total capacity with 776 cores in production (out of 1004) and we anticipate full completion with a full 1004-core production capacity by early September.

Please read the information below carefully since some of the changes to the new cluster environment may affect you.

## User Interface:

There are no changes to your login home directory, user account names or passwords. If you don't remember your password, please visit Tufts

tools at <http://tuftstools.tufts.edu/> . Home directories have retained their previous backup schedules and the default login shell environment remains bash.

Please note that you may receive a warning that the SSH key has changed when logging for the first time to the new research cluster at **cluster6.uit.tufts.edu**. This is the result of the new installation and is not a security issue. Your login credentials and login hostname have not changed, and simple fixes are detailed below.

## Storage:

Data on temporary file systems (/tmp, /scratch or /scratch2) were not migrated to the new cluster. Users were encouraged to transfer any important data prior to August 12. Temporary storage on file system /cluster/shared/your\_user\_name are available on the new cluster as before. Similarly, any directory on file system /cluster/tufts/ are available on the new cluster.

## Compute Nodes:

There has been a slight change to the compute node naming convention to help reflect hardware differences:

New iDataPlex nodes are now named: nodeNN (node01, node02, etc)

Legacy public BladeCenter nodes are: nodebNN (nodeb01, nodeb02, etc)

Legacy private BladeCenter blades contributed by Tufts faculty are renamed to: contribbNN (contribb01, contribb02, etc)

Going forward, new iDataPlex nodes contributed by Tufts faculty will be named: contribNN (contrib01, contrib02, etc)

## LSF Queues:

All LSF queues keep the same name on the new cluster.

## Parallel related:

Access to the new GPU resource will be activated in September after the new cluster is stable and in full production.

## SSH related login key issues

We are aware that changes to the ssh keys have caused some connectivity problems with the new cluster head node. A simple solution is to remove old cluster keys from your ssh client. This process depends on how you connect. For those that use ssh directly, (Mac OS X or Windows with CygWin) one may remove/edit out the old key in the following manner:

Launch the terminal application:

Mac OS X launch Terminal.app

Windows with CygWin launch CygWin and then startx

From the shell prompt, run the following command:

```
rm -f ~/.ssh/known_hosts
```

If you are using a Windows ssh or sftp client you may need to use the embedded key management options/tool. Unfortunately, there are too many ssh clients for us to provide exact steps for each. Your client documentation is the best source of information for making this change.

Once you have connected to the new head node, you may need to clear old ssh keys to ensure connectivity with all the new cluster nodes. The easiest way to do this is to run the following command when you have connected:

```
cleanSSH.sh
```

## SSH debug information

Should the above removal of ssh keys not work upon login, try connecting with ssh in verbose mode.

```
ssh -v -Y -C cluster6.uit.tufts.edu
```

or

```
ssh -v -v -Y -C cluster6.uit.tufts.edu
```

or

```
ssh -v -v -v -Y -C cluster6.uit.tufts.edu
```

where each additional -v adds more detail. Usually one -v is sufficient. Login, then logout and cut-n-paste the total debug output into a mail message and send to cluster-support@tufts.edu.

## **Support**

Please check your mailbox for a follow-up announcements. If you have any questions or concerns with Tufts new HPC research cluster environment, please direct them via email to [cluster-support@tufts.edu](mailto:cluster-support@tufts.edu).