

Tufts Cluster News

Sept. 30, 2015 Cluster Maintenance Window Announcement

Some compute node capacity was recently added as part of an end of Spring 2015 compute node purchase. Additional nodes will be added during the next several weeks to the compute node pool. The net effect of this additional capacity is to help reduce slurm job placement times, especially during busy periods of use. As the cluster environment matures it has become increasingly clear that there is a need for a standing maintenance window, similar to that found at other HPC sites. Thus far we have found ways to minimize user impact due to system changes without such a window. Our plan is to target Wednesday mornings at 6-7 AM starting 9/30/15. For the most part, maintenance activities during this window will not have any direct effects on access or job placements. Network concurrent license service is another possible service interruption during these windows. With history as a guide, these events are on the order of several minutes or less and don't happen all that often. In cases where user impact is expected, sufficient notification will be given. If you would like to express any concerns please contact us at cluster-support@tufts.edu.

May 12, 2015 Nercomp Data Challenge

Tufts, Harvard and Yale have organized an event entitled "Research Technology Services: The Data Challenge". Additional information can be found at <http://nercomp.org/index.php?section=events&evtid=443>.

Login node health monitoring policy!

Historically Tufts Technology Services relied on manual monitoring of the login node to support an acceptable quality of service for all. In addition we have asked users to be cognizant of the need to submit significant work to compute nodes via slurm (and formerly Isf). We will take an additional step in an automated monitoring solution on Monday 1/26/15 to promote the highest quality of service on the login node of the cluster. User processes exceeding a reasonable cpu utilization or memory limit will be terminated. An email will be generated and sent to the user with a brief summary and reminder to utilize the cluster compute nodes. If you have further questions regarding this new health monitoring feature, please send your questions/comments to cluster-support@tufts.edu.

Final migration of LSF cluster to Slurm based cluster completed!

Fall 2014 was a semester of transition. All IBM compute nodes have been re-imaged and added to the Cisco Slurm based cluster. The result is approximately 170+ compute nodes, 2600 cores and a peak performance of roughly 60+ Teraflops. In this new HPC environment, Tufts Technology Services provides researchers with access to scientific and engineering commercial and open-source research software applications. Cluster nodes are connected with high speed Cisco networking and access to secure storage for research data (400+ TB CIFS and NFS storage on NetApp appliances). The top ten users of compute cycles for the fall semester are:

utln	cpu minutes
hyu04	22810497
easgar01	16042764
smchug04	10054129
dsloug01	9554376
hyu04	9045697
smchug04	5487652
ksliwa	4869539
amasci01	4836376
hcui01	4027368
cburke05	2650937

Note: multiple utln entries correspond to different slurm accounts and is not to be confused with utln user accounts.

It is expected that HPC usage during the Spring 2015 semester should be larger than ever.

XSEDE - 2013 International Summer School on HPC Challengers in Computational Sciences (New York)

As part of Tufts participation in the XSEDE program I would like to make everyone aware of a very interesting opportunity. Feel free to pass this on to anyone who may be interested.

-- Shawn G. Doughty
Tufts University
Tufts Technology Services
Research and Geospatial Technology Services
Senior Research Technology Specialist
XSEDE Campus Champion
<http://it.tufts.edu/>
617-627-5462

Training available for students in U.S., Europe, and Japan at International Summer School on HPC Challenges in Computational Sciences

Graduate students and postdoctoral scholars in the United States, Europe, and Japan are invited to apply for the fourth International Summer School on HPC Challenges in Computational Sciences, to be held* **June 23-28, 2013, at New York University in New York City. The summer school is sponsored by** *the U.S. National Science Foundation's Extreme Science and Engineering Discovery Environment (XSEDE) project, the European Union Seventh Framework Program's Partnership for Advanced Computing in Europe (PRACE), and RIKEN Advanced Institute for Computational Science (RIKEN AICS). Leading American, European and Japanese computational scientists and high-performance computing technologists will offer instruction on a variety of topics, including: · Access to EU, U.S., and Japanese cyberinfrastructures · HPC challenges by discipline (e.g., bioinformatics, computer science, chemistry, and physics) · HPC programming proficiencies · Performance analysis & profiling · Algorithmic approaches & numerical libraries · Data-intensive computing · Scientific visualization

The expense-paid summer school will benefit advanced scholars from European, U.S., and Japanese institutions who use HPC to conduct research.

Further information and to apply for the 2013 summer school, visit* <https://www.xsede.org/web/summerschool13>. *Applications are due by March 18.

Contacts:

PRACE: Hermann Lederer RZG, Max Planck Society, Germany lederer@rzg.mpg.de

Simon Wong ICHEC, Ireland simon.wong@ichec.ie **RIKEN AICS:** Mitsuhsa Sato AICS, RIKEN msato@riken.jp

XSEDE:

Scott Lathrop NCSA, University of Illinois at Urbana-Champaign, United States lathrop@illinois.edu

*About PRACE: *The Partnership for Advanced Computing in Europe (PRACE) is an international non-profit association with its seat in Brussels. The PRACE Research Infrastructure provides a persistent world-class high performance computing service for scientists and researchers from academia and industry in Europe. The Implementation Phase of PRACE receives funding from the EU's Seventh Framework Programme (FP7/2007-2013) under grant agreements RI-261557, RI-283493 and RI-312763. For more information, see www.prace-ri.eu.

About RIKEN AICS*: *RIKEN is one of Japan's largest research organizations with institutes and centers in locations throughout Japan. The Advanced Institute for Computational Science (AICS) strives to create an international center of excellence dedicated to generating world-leading results through the use of its world-class supercomputer "K computer." It serves as the core of the "innovative high-performance computer infrastructure" project promoted by the Ministry of Education, Culture, Sports, Science and Technology.

*About XSEDE: *The Extreme Science and Engineering Discovery Environment (XSEDE) is the most advanced, powerful, and robust collection of integrated digital resources and services in the world. It is a single virtual system that scientists can use to interactively share computing resources, data, and expertise. The five-year project is supported by the U.S. National Science Foundation. For more information, see www.xsede.org.

Tufts Research Cluster Operating System upgrade:

Now that Tufts TTS has addressed with IBM all our concerns that resulted in the delay of our planned cluster upgrade, we are happy to inform you that we are ready to upgrade our cluster operating system to RedHat Enterprise License 6.2 (RHEL 6.2). We will do so progressively between October 11 and December 17 of 2012 to make it as safe and convenient as possible to our cluster user community. Indeed, to affect a smooth upgrade, we will not convert all nodes at once to RHEL 6.2. Instead, we will maintain both the current RHEL 5.5 environment and the new RHEL 6.2 environment to allow you to test all your installed packages, and to give you the possibility to revert to your currently stable environment while we resolve with you any issues discovered in RHEL 6.2. To start, the current RHEL 5.5 environment will have 42 fewer nodes (about a 35% reduction) and these 42 nodes will be upgraded to RHEL 6.2 and accessible through a separate RHEL 6.2 head node named cluster6.uit.tufts.edu. Logins to cluster6 will start Oct. 11. Note: your home directory, username and password will remain as it is under the new RHEL 6.2 environment.

Between Oct 11 and Dec 17, an average of six compute nodes will migrate on a weekly basis from RHEL 5.5 to RHEL 6.2. The supporting LSF queues will be duplicated to allow usage on RHEL 6.2 nodes with the number six added as a suffix to distinguish them from the original RHEL 5.5 queues. For example, the LSF queue 'normal_public' in the RHEL 5.5 environment will be renamed to 'normal_public6' in the RHEL 6.2 environment. Existing user-contributed nodes and the two GPU nodes will also migrate sometime during this transition period once we have agreed with their users of the most convenient time to do so.

In preparation for this cluster upgrade, Tufts TTS has tested all commercial packages and the most used open source software installed on the cluster. However, it is materially impossible to test every package installed for individual users. Therefore, during this transition period, we encourage all cluster users to try their software under RHEL 6.2 as soon as possible to allow us as much time as possible to help resolve any issue encountered. To help us, please report any issues encountered under the new RHEL 6.2 environment to cluster-support@tufts.edu. Please don't forget: your continued access to a RHEL 5.5 environment using the current cluster. uit.tufts.edu head node will offer you production continuity while any reported issue is addressed. Once you are satisfied with your RHEL 6.2 testing, please transition all your computational work to the RHEL 6.2 environment as soon as possible. Our current plan is to have transitioned all compute nodes to RHEL 6.2 by December 17, thus ending on that day all access to the old RHEL 5.5 environment. Ahead of that deadline, we will regularly poll the cluster user community to make sure that a full transition can take place as planned on December 17, 2012.

Tufts Research Cluster update:

In the past 5 years, we have seen steady increasing demand for our High-Performance Computing(HPC) research cluster. In parallel, we have steadily increased our HPC capacity from 80 cores in 2007 to 320 cores in 2008 and over 1000 cores in Sept 2011 with our latest major upgrade that tripled the size of the cluster we acquired in 2008. By the end of May 2012, the addition of 8 new additional compute nodes will bring the capacity of Tufts' Research Cluster to 1128 cores. This expansion of resources provides a better quality of service during periods of heavy computational use and increases the throughput for many types of computational jobs.

Tufts Visualization Awards

Congratulations to the winners of the 2012 Visualizing Research@Tufts Awards Program! Find out who won and view the [winning entries](#)

Please join us for the awards ceremony where all the entries will be showcased:

When: 3-5 pm, Thursday, April 12, 2012

Where: Alumnae Lounge, Medford Campus

Matlab's parallel licenses upgrade

Recently we have seen increased demand for access to Matlab's MPI Parallel computing capability. To help alleviate the concurrency clashes, we have upgraded our license. This upgrade will allow for two concurrent 16 core Matlab MPI jobs on the cluster instead of the current one job. On the other hand, if one constraints jobs to 8 cores, then 4 concurrent jobs are possible. Depending on the nature of what is computing, that might be a good compromise trade off for additional throughput. Furthermore, while one 32 core job is possible, it will cause others to wait, as before, and experience what appears as errors when it is really a failure of job placement due to no extra licenses. If you plan to use this capability please send us an email note to cluster-support@tufts.edu to discuss your options.

Qualtrics: New University-wide Survey Tool Now Available

TTS is pleased to announce that Tufts has purchased a university-wide license for Qualtrics, an easy to use, full-featured, web-based tool for creating and conducting online surveys. Tufts' customized version of Qualtrics is now available to all current faculty, clinical affiliates*, staff, and students via a dedicated website. To access Qualtrics, navigate to <https://tufts.qualtrics.com> and login with your Tufts Username and Tufts Password.

Qualtrics features go far beyond those available through free online survey tools.

- Over 100 different types of survey questions to choose from.
- The ability to embed multimedia into surveys.
- Completely customizable branding with your school or department logo.
- Support from Qualtrics' Research and Support Specialists.
- Robust online training and documentation.
- Direct exports to Excel and SPSS and an API that integrates with virtually any other system.
- Online collaboration in real-time as you and your colleagues author surveys.
- 48 supported languages.
- A tool to track participation and reminders.
- Accessible with your Tufts Username (UTLN) and Tufts Password---no new login information to remember.

To learn more about Qualtrics, user support for the product or security considerations, visit: <http://it.tufts.edu/>.

If you have any questions about this program or encounter any difficulty ordering software, please contact the TTS Client Support Center by email at uitsc@tufts.edu or by phone at 617-627-3376.

All use of Qualtrics must be in support of the University's administrative, teaching, and/or research work. Tufts' license does not extend to affiliate organizations unless use of the license is in support of research that originated within an academic unit at Tufts University.

Ansys IcePak software

A single concurrent seat of **IcePak** was added to address computational fluid dynamics (CFD) problems for electronics thermal management. This should be available early Feb. 2012.

Additional GPU node added

With the start of spring semester 2012, a second identical GPU node was added to the cluster to allow for additional use and use that intentionally spans nodes.

GPU resources available

As part of the recent research cluster upgrade this summer one compute node was provisioned with two Nvidia Tesla M2050 GPU processors. As promised, this note is to inform users that this new feature of the cluster is now functional.

GPU processing is an excellent means to achieve shorter run times for many algorithms. For example, for some Matlab codes the approximate performance enhancement is about an 8X speed-up over non GPU Matlab equivalents. Note, Nvidia Cuda(gpu language) and applications such as Matlab require specific coding to use gpu resources.

Cuda provided applications compile and run on the cluster. You'll find the CUDA toolkit in /opt/shared/cudatoolkit and the GPU computing SDK in /opt/shared/gpucomputingsdk. The SDK contains a number of CUDA sample C applications that can be found under/opt/shared/gpucomputingsdk/4.0.17/C. Compiled samples can be found in /opt/shared/gpucomputingsdk/4.0.17/C/bin/linux/release.

Matlab's Parallel Toolbox GPU demonstration applications also work. Additional applications such as Mathematica, Ansys, Maple and others offer various levels of support within their product.

To support GPU access new LSF GPU queues have been installed: short_gpu, normal_gpu and long_gpu. Additional GPU related information will be documented on the cluster wiki in the next few days.

HPC developments around GPU computing is growing rapidly. TTS/RTGS plans to add another identical GPU node before the end of 2011.

Expansion of Tufts High Performance Computing Cluster to Begin August 2011

Over the past three years TTS has observed an increasing demand for our High-Performance Computing research cluster. In anticipation of the ongoing need for additional resources, we are very proud to announce that TTS will be increasing Tufts' HPC capacity to over 1000 cores in early September. The new cluster capacity will be more than triple the size of the cluster we acquired three years ago.

In addition to increasing our capacity, we will make the following changes:

- The cluster network interconnect will change from InfiniBand to 10 Gigabit Ethernet.
- Several computing nodes will have 96GB of RAM for memory-intensive jobs.
- Experimental GPU research environment will be added for exploration of alternate parallel computing methods.

This upgrade is both complicated and extensive. As a result, there are two phases for the changes that will unfold during the month of August.

Phase 1, in progress, will bring the newly-acquired IBM hardware into a production ready configuration. Once the new cluster is fully configured and running, we will switch from our legacy cluster to this new environment in a transparent manner to cluster users.

Phase 2 will start on August 15 and will consist of powering off the legacy cluster to change the network interconnect. During Phase 2, the new cluster will be available through the same processes as the legacy cluster. As the hardware changes to the legacy cluster are completed, the number of cores available will increase. Priority will be given to faculty-contributed nodes during this process, and we anticipate full completion by early September.

If you have any questions or concerns, please direct them via email to cluster-support@tufts.edu.

Matlab toolbox additions

Five seats of Instrument Control, Real-Time Workshop and Real-Time Workshop Embedded Coder have been added to the Tufts network concurrent license.

Matlab Distributed Computing Toolbox

The license for this toolbox was upgraded from 8 to 16 cores for MPI jobs.

Deform software

Due to changing priorities and usage requirements, Deform software has been removed from the cluster.

GPU computing now and later

The latest versions of Matlab(R2010b) and Mathematica(ver.8) both support GPU computation. This option offers the promise of shorter duration computation times for some mathematical operations. Both products support several nVidia Cuda capable graphics cards. Note, at this time, there are no Cuda capable graphics cards on cluster hardware. However when installed on supported hardware such as your desktop or laptop you could then make use of GPU computing. **Note** During the Summer of 2011 the Cluster will be upgraded and some GPU capability will be added to provide approximately 2 Teraflops of additional computing power.

Visualization competition

To showcase Tufts research projects, enable opportunities for collaboration, and promote the use of visualization as a research tool, **TTS** is launching the first Visualizing Research @Tufts Awards program in Spring 2011.

For more information see: <http://sites.tufts.edu/vrta>

Another private node added to cluster

Since the Fall of 2009, Tufts Professors Khardon(CS), Miller(EE), Abriola(Civil) and Cohen(CS) have contributed compute nodes to support their research work. As explained on the cluster main page, arrangements have been made to allow their groups priority access. During Dec. 2010 Physics Professors Austin Napier and K. Sliwa contributed one additional node to the cluster.

Matlab upgrade

The Oct. 2010 upgrade of Matlab to ver. R2010b brings several **enhancements** over previous versions. In addition an increase to Tufts' base Matlab license from 40 to 65 network concurrent seats was addressed to mitigate license contention during periods of peak demand. Two toolbox enhancements to note are the **Distributed Computing** and **Curve Fitting** Toolboxes. The former allows for Matlab MPI parallel computing job on the cluster with 8 cores while the latter is a new contribution to ver. R2010b.

Tufts Carma subscription

New for Fall 2010 is **Carma**, the Center for the Advancement of Research Methods and Analysis, is an interdisciplinary center devoted to helping faculty, graduate students and professionals learn of current developments in various areas of research methods and statistics. This subscription provides to the Tufts community web access to Carma video lectures. A broad range of applied statistical and methodological topics is presented covering the social, management and behavioral sciences. Anyone at Tufts interested in viewing these lectures must register with a Tufts email address. Directions and further information is available on the website. For additional information please contact durwood.marshall@tufts.edu.

Bioinformatic tool additions

Recently the following software items have been added to the cluster:
TopHat, BowTie, Cufflinks, blast, MrBayes, Velvet, PolyPHred and BioPerl.

Abaqus license

Tufts Abaqus license has been increased by two additional network concurrent seats for Explicit, Standard and CAE.

Comsol Upgrade

Comsol recently released a major upgrade. Both version 4.0 and 4.0a are available. COMSOL Multiphysics version 4.0 is in many aspects an entirely new product. The release notes outline the differences between version 3.5a and version 4.0. It also outlines additions that Comsol intends to make in COMSOL Multiphysics version 4.0a. Comsol strived to achieve backward compatibility with the previous version and to include all functionality that is available there. However, version 4.0 is not fully backward compatible with version 3.5a. All backward compatibility issues are planned to be solved for version 4.0a unless explicitly stated.

Star-P removed from cluster

Interactive SuperComputing was purchased by Microsoft last Fall 2009. The product roadmap for the linux version of Star P is unknown and further development, releases and bugfixes has ceased.

Ansys v.12.1 and Polyflow

Ansys has been upgraded to ver. 12.1. This represents the first merger of Ansys and Fluent product lines. **Polyflow** has been licensed as well. It is normally access via the **WorkBench** interface.

IBM HS22 blades

As part of a yearly effort to enhance Tufts High Performance Cluster computing resources, TTS has added five additional compute nodes. These nodes are next generation **IBM HS22** blades with 2 sockets, 6 cores per socket and 50 gig ram per node. The new nodes yield an additional 60 compute cores and brings the total cluster core count to 468! This addition will help ease job wait times during heavy periods of demand.

New Matlab Toolboxes

Two new Matlab toolboxes have been added to our Tufts network concurrent license to support Bioinformatics and Computational Biology at Tufts. For additional information please check:

[Bioinformatics](#)
[SimBiology](#)

Recorded presentations, webinars and demos can be found [here](#).

Viswall NSF Article

Recently the National Science Foundation's online magazine Science Nation featured an article about Tufts Viswall visualization facility. This resource is in fact connected to Tufts Research Cluster to provide additional computation capabilities. You may view the article [here](#).

Fe-Safe

This is a new addition to the cluster. **Fe-Safe** is a tool for fatigue analysis of Finite Element models.

Abaqus Upgrade

Abaqus was upgraded to version 6.9-EF2. This is a major enhancement to v.6.9.

Maple Upgrade

Tufts network concurrent license has increased to 5 seats. This is reflected in the Maple version 13 upgrade which is now the default version on the cluster.

Deform Batch

Scientific Forming Technologies Corporation's add-on Batch product has been added to **Deform 3D** on the cluster. This is an internal batch facility to Deform.

Comsol MEMS

An addition to the Comsol suite now includes module **MEMS**. MEMS is now available on the cluster and can be added to a PC install as well. Our network concurrent license supports one seat.

Chaos Uncovered?

Mathematician Bruce Boghosian has shed some order on the matter of chaos. His recent experiments at supercomputing center **TAAC** is the subject of a featured article, The Skeleton of Chaos. Without the computing power of TAAC this would have been a more difficult problem. A current project in the development stage includes using the Tufts Cluster with **PetSC** software and parallel processing. Once satisfied with their code development and performance, efforts to migrate it to a larger cluster will help uncover details and complexity that can result from the promise of computational scalability.

HPC class at Tufts

How fast is the cluster anyway? Well that depends on many factors. Tufts recently offered a class in High Performance Computing(**HPC**), Comp_150, sponsored by Computer Science, Mathematics and TTS. Students of visiting Brown University Prof. Leopold Grinberg have explored the fundamentals of HPC on Tufts' cluster as well as machines at supercomputing centers **TACC** and **NICS**. TAAC's Ranger and NICS's Karken were used for class projects. Student Constantin Berzan has kindly offered his recent review of serial program benchmarking. The attached document summarizes his effort. Tufts cluster does very well, indeed!

Private nodes added to cluster

Since the Fall of 2009, Tufts Professors Khardon(CS), Miller(EE), Abriola(Civil) and Cohen(CS) have contributed compute nodes to support their research work. As explained on the cluster main page, arrangements have been made to allow their groups priority access.