

StarHPC – Elastic Computing in the Classroom

By Justin Riley , Programmer Analyst, OEIT jtriley@mit.edu

StarHPC provides a platform for students at MIT to learn how to program in parallel environments. It consists of a 4 node compute cluster hosted on Amazon's Elastic Compute Cloud (EC2) web service. The cluster comes equipped with the Eclipse Parallel Tools Platform IDE for developing with OpenMPI (Open Message Passing Interface) and the Sun Studio compilers for developing with OpenMP (Open Multi-Processing).

For OEIT, StarHPC also serves as a first look at using Amazon's web services in the classroom as a way to reduce both administrative overhead and housing costs associated with owning a compute cluster. With Amazon's new EC2 web service, one can request a number of virtual machines (up to 20 by default) running in Amazon's data centers and pay only for what they use at \$0.10/hr, per virtual machine. Comparing this to the overhead of purchasing servers, providing housing and cooling, and administering and maintaining the cluster this becomes a very appealing service with an affordable pricing model.

What's even more appealing is that there is an API for managing the virtual machines programmatically. This allows for automating the administrative overhead associated with configuring the servers, configuring software, and setting up user accounts. It also allows for the creation of a dynamic pool of resources that is resizable in terms of current needs. The vision here is that departments eventually could sign up for this service, download a set of scripts from the STAR project, fill in their account info, and have an up and running system in a matter of minutes that serves a variety of needs. The department would then only have to foot the cost of what they've used and not the cost associated with storage, cooling, server maintenance and configuration.

During this year's Independent Activities Period StarHPC had its first use in the classroom in Dr. Constantinos Evangelinos' IAP 12.950: *Parallel Programming for Multicore Machines Using OpenMP and MPI* course based out of the department of Earth, Atmospheric, and Planetary Sciences. The class was meant to be a crash course in both MPI and OpenMP programming techniques for writing parallel codes to run on high performance computing systems. Throughout the course, students were assigned a set of evolving parallel codes as homework. In order to do the homework, students logged in remotely to the StarHPC cluster to develop, compile, and run their parallel codes from virtually anywhere; even from a web browser if necessary. The course lasted two weeks and had a total of ten users actively developing on the cluster. The estimated cost for using the StarHPC hosted on Amazon's EC2 web service is around ~150USD. Comparing this to the cost of a physical cluster, EC2 is an attractive way to support peak period computing needs in the classroom.

In addition to the cluster, StarHPC also provides a standalone virtual machine image which can be used locally on a student's machine to develop with the same set of development tools provided by the cluster. Eventually, as the Eclipse Parallel Tools Platform matures, we hope to be able to connect the local virtual machine image to the remote cluster. This would allow application development to happen on the student's local machine and parallel runs to be submitted to the remote cluster on Amazon's EC2.

To learn more about StarHPC, visit the project web page at <http://web.mit.edu/star/hpc>. There you will find documentation and links to download the standalone virtual machine image. At this time the cluster is no longer active as the IAP course is finished. Those interested may contact the STAR group directly at star@mit.edu.