

Policy Guidelines for Yeti Use for Teaching

SRCPAC Education Subcommittee

Version 1, December 2, 2014

In order to educate members of the Columbia community in the use of computing in research, students and faculty in courses with a substantial emphasis on computing should have access to the Yeti cluster and its descendents. It is expected that such access will require only minimal amounts of actual calculational time. It is the ready availability of a high performance cluster with a modern software stack that is important for education, so students can briefly run programs, establish benchmarks and test algorithmic ideas. For such tests to be effective, a short duration queue, with a time limit of 15 minutes, and rapid turn around time should be available. For testing and benchmarking parallel codes, it is important that students be able to run on a larger number of nodes, again for short durations.

When possible, basic code development need not be done on Yeti, unless specialized software libraries are required. Current experience shows that most students with Macintosh or Linux laptops can easily do code development on them. Inexperienced programmers using Windows will likely need a single node of a computer comparable to Yeti for code development.

The required software on Yeti should include a standard Linux software development environment with up-to-date compilers, including for example C++11. Publicly available GNU and/or LLVM compilers are generally sufficient, as maximal performance is likely not of primary importance. However, access to more hardware specific compilers, such as those from Intel, is useful. A selection of publicly available packages, such as linear algebra BLAS routines, should be installed. More specialized, publicly available packages, may be installed at the request of the instructor, as long as sufficient notification is given. Many research groups use large, specialized software packages and these are likely not needed for general educational purposes. Some of these may be installed, upon request, if there is sufficient support staff to do this.

As a guideline, it is expected that cumulative educational usage of Yeti would be capped at 5% of the available machine time, averaged over a year. Educational usage may peak when assignments are due, or final projects are being done.

Procedures for Applying for Yeti Use

- Applications to use Yeti for courses should be sent to the chair of SRCPAC, or their delegate. If the request is approved, the SRCPAC chair should inform CUIT staff of the approval and forward a copy of the application to CUIT.
- The request for Yeti access should demonstrate that computing plays a major role in the course and that the instructor is prepared to teach the necessary computing skills required by the course.
- Students should be encouraged to develop codes on their laptops, where possible, to reserve Yeti usage for running final jobs.
- The initial request should include information about the enrollment, an estimate of the CPU time needed for the assignments in the course and an estimate of the disk space needed, particularly if the course involves large data sets.
- CUIT staff will be responsible for the system administration required for access to Yeti and will not play a teaching role in the course.
- Help requests regarding system administration and software should be sent to CUIT from the instructor.
- The initial request should detail any required software for the course. Standard, publicly available software packages may be installed upon the request of the instructor. CUIT will install non-standard software on a best effort basis.
- Course related jobs will be run primarily on the older partitions of Yeti, with access to high-performance parallel nodes possible if need is demonstrated.
- A student could continue usage on the free tier, after the course is completed, if approved by the instructor.

Yeti Application for Course Usage

- Course:
- Semester:
- Instructor:
- Estimated enrollment:
- Required software:
- Estimated CPU usage and storage:
- Brief description of the role of computing in the course:
- Course staff computing experience: