

SRCPAC Fall Meeting

Wednesday, December 3, 2014

12:30 p.m. – 2:00 p.m., 523 Butler Library

Minutes Report

Attendance

Kathryn Johnston, Astronomy (Chair of SRCPAC)

Rajendra Bose, Zuckerman Mind Brain Behavior Institute (Chair of Intercampus Subcommittee)

Greg Bryan, Astronomy (Chair of Cloud Subcommittee & Yeti Operating Committee)

Bob Mawhinney, Physics (Chair of Education Subcommittee)

Ryan Abernathy, Earth & Environmental Studies

Alexander Antoniadis, CUIT

Marley Bauce, Research Initiatives

Steve Bellovin, Computer Science

Jerri Bland, CUIT

Michael Burke, Mechanical Engineering

Harmen Bussemaker, Biological Sciences (*Phone*)

Robert Cartolano, Libraries & Information Sciences

Alan Crosswell, CUIT

George Garrett, CUIT

Victoria Hamilton, Research Initiatives

Halayn Hescoek, CUIT

Jeffrey Lancaster, Libraries & Information Sciences

Rob Lane, CUIT

Don Lemma, Business School

W. Bentley MacLeod, Economics

Chris Marianetti, APAM

Amy Nurnberger, Libraries & Information Sciences

Mahdad Parsi, Lamont

Ingrid Richter, Psychology

William Vanti, Libraries & Information Sciences

Jochen Weber, Psychology

Breck Witte, Libraries & Information Sciences

Tian Zheng, Statistics

Summary

At the Fall 2014 SRCPAC Committee meeting, members voted affirmatively on three new policies:

1. The creation of a SRCPAC Executive Committee that will meet as necessary to make decisions between larger SRCPAC meetings.
2. Allocation of up to 5% of Yeti for use in courses with a substantial emphasis on the use of computing in research, following the policy guidelines presented by the Education Subcommittee.
3. Restructuring the Yeti Operating Committee to include the Chair of the Yeti Operating Committee, four members representing the Top 50% of users, and four members representing the Bottom 50% of users (as defined by the number of nodes purchased). The Chair will have special responsibility to represent renters and free tier users. Issues unresolved by the Yeti Operating Committee will be referred to SRCPAC. The revised governance structure will be reviewed at the Spring meeting of Yeti.

All three items unanimously passed by voice-vote without any objecting or dissenting committee members.

Introduction

Chair Kathryn Johnston called the meeting to order at 12:30 by asking all meeting participants to introduce themselves. As the entire SRCPAC committee only meets once per semester, the Chair introduced the prospect of creating a SRCPAC Executive Committee – composed of the SRCPAC Chair and the Chairs of the various SRCPAC Subcommittees – in order to make time-sensitive decisions that do not seem to merit requiring a special full SRCPAC meeting. The committee unanimously agreed to create this Executive Committee, effective immediately.

The Chair then reminded the Committee that SRCPAC advises and takes direction from the Research Computing Executive Committee (RCEC). Last summer, the RCEC requested the creation of three new SRCPAC Subcommittees; a primary agenda item of the current meeting would be to review progress for each Subcommittee. The Chair noted that a fourth Subcommittee – the Storage Subcommittee, chaired by Tom DiPrete, which generated the Research Storage Pilot – has been disbanded.

Cloud Subcommittee

Greg Bryan, Chair of the Cloud Subcommittee, began by presenting his Subcommittee's charge of identifying which communities are using Cloud services, which communities are not but *should* be using Cloud services, and how to best mobilize university resources in order to facilitate researchers' transition over to cost- and space-effective computing services.

To this end, the Cloud Subcommittee met twice over the semester and identified three recommendations: 1) Purchasing of services is difficult via the university's P-Card system (as there is a limit on spending amount), which can be circumvented if we finalize an enterprise purchasing agreement with Amazon; 2) The Office of the Executive Vice President for Research should develop a public website that identifies the rules of Federal grants regarding purchasing Cloud services; and 3) SRCPAC will solicit volunteers to participate in a Cloud Pilot Project, whereby a "modest" amount of cloud resources will be purchased, and the SRCPAC committee will identify, develop, and document appropriate technical resources to aid in the transition from Yeti to the Cloud. The Cloud Subcommittee fielded questions about the definition of "modest," and of the degree of financial support required for such a pilot, though the Subcommittee and SRCPAC chairs replied that this term has yet to be defined, that the Executive Committee can begin discussing this, and that there will be a progress report provided

at the Spring 2015 SRCPAC meeting. It was also noted that the cost, size, and scope of support services will depend on the composition of the pilot program.

Slides presented attached as [Appendix I](#).

Intercampus Subcommittee

Rajendra (Raj) Bose, Chair of the Intercampus Subcommittee, briefly summarized his Subcommittee's goal of identifying potential commonalities, overlaps, and synergies among the research computing activities happening on the Morningside, Manhattanville, Lamont, and Nevis campuses. The Subcommittee currently does not have CUMC representation, though there are plans to have a CUMC IT representative join the committee for the Spring 2015 semester. Leadership representatives from A&S and SEAS will also be invited to join the Intercampus Subcommittee for the Spring 2015 semester.

The Intercampus Subcommittee concluded its two Fall 2014 meetings with multiple policy recommendations. It supports the Cloud Subcommittee's proposal for an enterprise purchasing agreement with Amazon. It also recommended a survey of the computational needs of recent Data Science Institute faculty recruits in order to project potential computing needs for the university's future strategic recruitment initiatives. To ensure that multiple academic units and individual PIs don't reproduce each other's work, the Subcommittee proposes the hiring of a full-time employee (potentially housed within CUIT) to research external computing resources – be those resources facilities, vendors, partnerships, etc. The Subcommittee notes its investigation of a proposed Connecticut regional Data Center modeled after the Massachusetts Fall River Data Center, and recommended that CUIT continue to investigate regional facilities as an option for external resources.

The Subcommittee proposes establishing a formal Research Computing & IT community containing representatives from all Columbia campuses, with a rotating group of three representatives charged with communicating policy recommendations to SRCPAC and the RCEC.

Finally, the Subcommittee advocates that a CUIT staff member attend an upcoming Basic Science & Engineering Chairs meeting (hosted by the Office of the Executive Vice President for Research) in order to encourage faculty to communicate with CUIT when they are preparing grant proposals for resource-intensive research projects. CUIT can do this by providing case study examples of successful collaborations that involved CUIT at the initial stages of proposal development, as well as examples when CUIT was apprised of a proposal midway through its development and was able to intervene to provide adequate support services but perhaps in less than optimal ways.

Slides attached as [Appendix II](#).

Education Subcommittee

Bob Mawhinney, Chair of the Education Subcommittee, reviewed the *Policy Guidelines for Yeti Use for Teaching* document generated by his Subcommittee, which met twice during the semester. The Committee proffered a general philosophy of fostering the education of the Columbia community in the use of computing in research by providing access to the Yeti cluster and its descendants by students and faculty in courses with a substantial emphasis on computing. A list of practical steps and considerations for granting students access were presented. The Education Subcommittee proposes that educational use of the Yeti be capped at a total of 5%, averaged over the Academic Year. To the extent feasible, certain primarily publicly available software should be made available such as GNU and/or LLVM, and up-to-date compilers such as C++11. The Subcommittee understands that there are multiple Columbia courses that have substantial foci on computational methods, where students would be advantaged by the opportunities to run benchmarks and test algorithms, which requires that they utilize a large number of nodes, even if only for a brief amount of time. .

Specific details of the step-by-step procedures for granting professors and students access to Yeti are contained in the *Policy Guidelines for Yeti Use for Teaching* document attached as Appendix III, and are posted to the SRCPAC website.

The SRCPAC committee questioned the plausibility of building an environment – accessible on student laptops – that closely matches the Yeti interface, thereby localizing the development cycle as much as possible. CUIT representatives on the committee replied that this was feasible to accomplish, depending on the capabilities of individual laptops. For most students, an environment can be built that is *very close* to Yeti, and they should be encouraged to do most of their work on laptops before testing on Yeti.

The committee then discussed whether more focus should be placed on running code via the browser (such as Python), though some committee members did not think this to be the best solution for hard science courses (which typically encourage students to adjust to massive computing environments that use traditionally old tools), and may be more appropriate for the social sciences and humanities.

The Chair and the Committee agreed to table the discussion until the Spring 2015 SRCPAC meeting, once the Education Subcommittee has had more time to deliberate and define additional policies and procedures. The Subcommittee further noted the benefit of establishing stronger ties between the Subcommittee, the Computer Science Department, Libraries & Information Sciences, and various departments in the social sciences and humanities. The larger Committee additionally proposed that individual project courses (independent studies) also be granted access via the policy procedures identified for traditional courses.

Upon conclusion of discussion of educational access to Yeti, the committee unanimously approved the Education Subcommittee’s proposals, with no objections.

The issue arose of providing curious students who were no longer enrolled in a class (or had never attended such a class in the first place) a place to experiment. The SRCPAC Chair indicated that the Free Tier would soon be made available, and that, with the endorsement of a faculty member, students would have access to that Free Tier. The Guidelines for student use of the Free Tier will also be made available online.

The Education Subcommittee briefly reviewed Fall 2014 activities occurring in collaboration with CUIT and the Libraries, which offered several training workshops such as Introduction to Linux and Introduction to HPC, both of which were received positively. The Libraries noted that their contribution was to plan and staff events, though CUIT contributed all content and instruction. Discussion also focused on how to leverage online resources to teach some basic tools. The committee noted a number of external resources already exist, such as Lynda.com (http://ccnmtl.columbia.edu/enhanced/noted/new_subscription_lynda.html). SRCPAC should consider how to leverage these types of educational activities, thereby concentrating efforts onto *new* methods and resources instead of duplicating preexisting resources.

Slides attached as [Appendix III](#).

Yeti Round 2 Expansion

Next, Rob Lane, who staffs the Yeti Operating Committee, outlined major updates in the recent Round 2 expansion processes of Yeti, which will have new systems in production by January 2015.

This recent expansion added users from 15 departments or units across the university, in contrast to the 10 in the original Yeti cohort (although one group is a repeat buyer). Many new users are from SEAS, which incited users to join Yeti and use the SRCPAC facility. Researchers purchased 66 new nodes, compared to the 49 purchased in the original round (augmented originally by the 52 purchased through the NYS matching grant). In particular, there was a notable increase in the number of Infiniband purchases (48... up from 16 from the first round). Since its initial launch, Yeti usage has trended upwards; a distributed graph illustrated growing usage. The “theoretical” limit on Yeti is approximately 48,000 core days per month (with 1,616 cores on the cluster). There is currently ample room, due in part

to the contribution of capacity from the NYS Match. Rob Lane commented that he has received assurance from users that this limit will be met!

The Committee briefly discussed the future financial model and spoke about contributions from the University. The SRCPAC Chair interjected the need for appreciating the university's commitment to this resource by providing the electrical power, physical space, and contributions towards administrative staffing that have been necessary to ensure past, current, and future successes. The committee noted the potential for suggesting that the university additionally contribute minor funds for Yeti nodes exclusively dedicated for educational purposes as a way of ensuring continued financial sustainability.

Slides attached as [Appendix IV](#).

Yeti Governance Committee

The Chair took up the matter of altering the Yeti Operating Committee to recognize and incorporate new participants. The Chair asked Rob Lane to outline the current structure and practices of the Yeti Governance Committee. He indicated that the Yeti OC meets several times per year, primarily to make collective decisions on maximizing the resource's use, and ensuring fair treatment (given that multiple users have varying levels of system access and have contributed varying amounts of funds). The format of the Governance Committee meeting has the first part open to any and all Yeti users for informational purposes and to solicit input. Afterwards, a private Executive Committee of those empowered to vote deliberates and decides upon action items. Issues which cannot be reconciled are brought to the next SRCPAC meeting.

The current Governance Committee's composition includes six executive committee members: two representing the large groups, one representing the medium groups, one representing the small groups, and the Chair representing the rental and free tiers. The SRCPAC Chair then proposed a reconfiguration of the Yeti Governance Committee to include nine members: Four users representing the seven groups who collectively purchased 50% of the hardware, four users representing the seventeen groups whose smaller purchases totaled the other 50%, and Greg Bryan, the Yeti Governance Committee Chair, again representing the rental and free tier users. The Chair indicated that SRCPAC would review the success of this new governance structure at the Spring 2015 meeting.

The SRCPAC committee unanimously voted to approve this Governance Committee composition. The SRCPAC Chair will subsequently email all SRCPAC members soliciting volunteers to join the Governance Committee.

The meeting concluded with the thanks of the Chair for all the contributions. The next SRCPAC meeting will be held in April or May 2015.