

Attendance for SRCPAC meeting of May 22, 2014

Name

Breck Witte
Mark Newton
Alan Crosswell
Kathryn Johnston
Wojciech Kopczuk
Robert Cartolano

Victoria Hamilton
David Madigan

Greg Bryan
Chris Marianetti
Andreas Hielscher
Donald Lemma
Robert Mawhinney
Rajendra Bose
David Greenberg
Alex Bergier
Abhishek Joshi
George Garrett
Robert Lane
Mahdad Parsi
Hatim Diab

Draft Minutes for SRCPAC Meeting of Thursday, May 22, 2014

Call to Order - David Madigan

Chair David Madigan called the meeting to order at noon. He reminded the Committee that it had selected Prof. Kathryn Johnston as the Vice-Chair in the Fall, anticipating that she would succeed David Madigan so this would be his last meeting. The Chair mentioned that Rob Lane had been promoted to manager of Research Computing Services to succeed Raj Bose who has assumed responsibilities for ZMBBI Computing Services. The Chair also introduced George Garrett as one of the RCS staff. <http://rcs.columbia.edu/>. The Chair indicated that normally he would call for approval of the Minutes of the September 5, 2013 meeting, but asked to postpone that motion until other agenda items were discussed that might help inform any discussion.

RCEC – Kathryn Johnston

The Chair asked the Vice-Chair to report on their annual review with the Research Computing Executive Committee (RCEC). The Vice-Chair reported that the RCEC had asked SRCPAC to reinstate two committees, namely the Intercampus Committee, with particular concern about including Manhattanville, and the Cloud committee. The Vice-Chair asked Raj Bose to say a few words about his mandate for computing for ZMBBI, which he proceeded to do. The Vice-Chair also called on Don Lemma, who as CIO of the Business School which will be building on the Manhattanville campus, also spoke to the potential for synergies. The Vice-Chair asked Bose, Lemma and Alan Crosswell to serve on the Intercampus Committee, and then called for other volunteers. The Vice-Chair then asked for comments about whether SRCPAC should resurrect the Cloud Committee. Vigorous discussion ensued about public, private, and government options beyond the campus. It was determined that it made sense to create a committee to keep an eye on this as well. Bob Mawhinney and Alan Crosswell volunteered. Again, the Vice Chair asked for additional volunteers to contact her.

Update on Research Storage Pilot

The Chair then asked Rob Lane to deliver Tom DiPrete's report on the Research Storage Pilot. Lane indicated that the good news is that the current pilot users were happy. The bad news is that there were only three of them. He quoted from an email from one user:

Julien Teitler: It does what I want from it, which is to provide a seamless way to share files with collaborators, within the Columbia system. For security reasons, there are many files we don't want to store in Dropbox or other external systems.

Discussion ensued about pricing and caps. There was a call to investigate whether there was a sustainable funding model or whether the storage capacity should be redeployed in another way, for example, added to *Yeti*.

HPC Update and Proposed new HPC purchase cycle

The Chair then asked Victoria Hamilton to report on Hotfoot and Yeti usage. Hamilton displayed several graphs showing the usage of each cluster. The 200 users of Hotfoot keep it fairly well used, with a dip at Christmas time although the usage of different groups fluctuate, which is one of the advantage of a shared cluster. The Yeti usage graph showed an initial slow start followed by a continued upward ramp. A question was asked about capacity utilization, and the response was that usage was roughly half in April, although it was pointed out that would be close to 100% of the capacity purchased by the users as the NYSTEM grant approximately doubled the size of the cluster available. Another SRCPAC member recommended that usage rules and queuing priorities be reviewed to encourage fuller usage of the machine.

Vice-Chair Kathryn Johnston then asked Hamilton to discuss the possibilities of an incremental purchase round. Hamilton reported that there appeared to be interest from new faculty, particularly given the push from the Data Sciences Institute. She added that they had been weighing whether it made sense to push for a speedy decision on the part of purchasers in order to make sure the incremental capacity was ordered, installed and tested by the start of the semester, or better to wait to let new faculty determine how best to use start-up packages. Fortunately, the *Yeti Operating Committee* had resolved the issue by agreeing to allow PIs to join *Yeti* as soon as they provided a chart-string for purchase. SRCPAC agreed that was a good use of the current unused capacity.

Proposed Trial-with-Intent-to-Purchase (TIP)

The Chair then asked the Vice-Chair to lead the discussion of the proposed **Trial-with-Intent-to-Purchase** or **TIP** program, an outline of which had accompanied the notice of the SRCPAC meeting. The Vice-Chair indicated that the goal of TIP was to enable PIs with serious interest to purchase to have a brief trial before committing. Unlike a pure free tier, the TIP would provide full service and priority above the free tier. The TIP would be capped to ensure a good experience for TIP participants and current users (*see attached for specifics*). Various comments and questions followed. The Vice-Chair indicated that the TIP program would be introduced as soon as feasible to give interested parties a chance to trial *Yeti* before the anticipated purchase date in August or September. Participants in the TIP program who decided not to purchase or rent would convert to normal Free Tier after the TIP trial period.

The Chair then asked several members of SRCPAC to discuss concerns about the purchasing process. Prof. Bob Mawhinney asked if non-conventional types of machines were explored. Prof. Chris Marianetti also indicated that the absence of certainty about what machines and the queueing protocols had squelched enthusiasm among some researchers. CUIT explained there had been a normal, required bid process and only three vendors had responded. Discussion during the meeting and after also focused on the pros and cons of low-power nodes, Infiniband and particular choices of software. Queueing protocols are dictated by the Operating Committee itself, so can't, by definition, be determined until the operating committee is formed.

Another SRCPAC member asked about more methods for on-ramping researchers interested in using HPC. In particular, it would be useful to understand schedulers, see some recipes and templates. RCS was asked to reach out to the Psych group. Rob Cartolano of the library also indicated willingness to support such training in collaboration with CUIT.

Prof. Mawhinney then raised the question of whether it would be appropriate to use some of *Yeti's* capacity for teaching scientific computing. He had taught a small class of 10 or a dozen people last semester, and access to *Yeti* could be valuable. There seemed general consensus that this was an avenue worth exploring by SRCPAC and the *Yeti OC*.

Yeti Governance Model

The Chair returned to the issue of the Minutes. After publication of the draft minutes in the Fall, there had been comments from the members that indicated some confusion about the final decision regarding the make-up of the *Yeti OC*. The Chair indicated that the *Yeti OC* has been meeting regularly & working well. He reminded SRCPAC that the voting members of the OC consisted of 1 representative for the small toe-in-the-water purchasers, 1 representative for the mid-range purchasers, 1 representative for the renters, and currently 2 representatives for the large purchasers. However, some people had felt there should be three representatives for the large users. Some discussion ensued. The Chair also pointed out that depending on the results of the next proposed round of purchasing, the representation might have to be revisited. The group consensus was to continue the committee as currently constituted, and revisit should there be a large influx of new users.

The Chair concluded the meeting at 1pm.