

Computation and Data Committee Report to the CASCA Board, May 2016

Current Committee membership:

James Wadsley (McMaster) (Chair)	Term ends: 30 June 2016
Hugo Martel (Laval)	Term ends: 30 June 2017
J. J. Kavelaars (HIA/NRC/CADC):	Term ends: 30 June 2017
Erik Rosolowsky (Alberta)	Term ends: 30 June 2018

James Wadsley is willing to renew as Chair. We might add one more member as a transition or consider a switch of chair if someone is interested in taking over.

Computing in Canada

Compute Canada

Compute Canada (CC: <http://computeCanada.ca>) is the main source of cycles and storage for researchers at Canadian universities. It is led by CEO, Mark Dietrich, CSO (Chief Scientific Officer) Dugan O'Neil (Physics, Simon Fraser) and CTO (Chief Technical Officer) Greg Newby. The CC Advisory Council on Research (ACOR) includes James Wadsley (McMaster), of this committee and Robert Thacker (SMU). The members of CC are the universities and colleges in Canada.

CFI Funding for HPC

In 2014, CFI announced a new Cyberinfrastructure initiative to fund not only hardware but also associated software development. A particular focus was software infrastructure to deal with the surge in research data such as that generated by astronomy (both observational and simulations). The initiative originally called "research data" challenge 1 (with two competitions, one in 2015 and one in 2016) as described here:

<https://www.innovation.ca/en/OurFunds/CFIFunds/cyberinfrastructure-initiative>

The long overdue hardware refresh was called challenge 2 but its timeline completed earlier than Challenge 1 and equipment funding was first awarded in 2015. A second hardware competition was scheduled for 2016 whose purpose was partly to provide equipment (e.g. storage) in support of the funded research data projects. The first hardware refresh was decided in June 2015 with substantial investments at 4 sites: SFU (GPU focussed), UVic (cloud focussed), Waterloo and Toronto (A large parallel system). This investment would result in only a small increase over the existing but quite old infrastructure and thus leave Canada well behind our competitors (e.g. G8 or G20).

Several factors led CFI to drastically modify this plan in Fall 2015. Due to criticisms and negative comparisons with competitors and concerns that the then Harper government might claw back unspent HPC funds, CFI suddenly announced it would double the HPC funding per site. Not all sites were able to double and not all provincial governments were happy to double their matching contribution. None-the-less the equipment purchases are going ahead and the first systems are expected to become available in late 2016. Most existing systems will be decommissioned at part of this process. This will probably cause a lot of disruption, particularly for large parallel users as Scinet (Toronto) will have to decommission its current large system first.

CFI continued with a second hardware competition with final submissions due in May 2016. Compute Canada took submissions from individual members (universities and colleges). In some cases, regions vetted these proposals and submitted a combined regional proposal. The original deadlines were pushed later slightly as a result of CFI's changes in 2015. As noted below, it seems likely that this

competition will be substantially re-thought with some attendant delays.

Compute Canada Governance Changes

A long-standing concern has been how infrequently ACOR meets and a general lack of consultation by Compute Canada with stakeholders. Ontario created Compute Ontario, led by a new CEO Nizar Ladak, partly in response to this. Issues came to a head very recently when Compute Canada decided to award yet more hardware money to SFU (and UBC) via what was widely regarded as a flawed internal process. CC was due to give the results to CFI on May 24th. Instead the membership of CC (VPRs and other representatives of Canada's universities and colleges) called a special general meeting of Compute Canada (with CFI present) and threw out the results with the intent of redoing the process. CC is apparently scrambling to put together a new hardware proposal with more direct regional input but it seems likely that the whole process will instead be suspended for several months. A key question that also affects astronomy is how to handle special purpose hardware and storage requests. For example, the SFU request and several other expensive investments by Compute Canada are driven primarily for the benefit of ATLAS users. This is a small subset of Canada's particle physics community. ATLAS requires very high I/O rates making the equipment very expensive compared to the more general purpose purchases of the past. There is of course the bad optics of CC devoting massive resources to the research area of its scientific director. Astronomy also has large storage and bandwidth needs. Whatever new process CC comes up with to select hardware needs to better quantify the extent of the demand and prospective users to justify high cost purchases.

An additional outcome of the CC members meeting was an agreement between the members and CFI (and Industry Canada) to completely rework the governance model of Compute Canada over a 6-18 month period. CFI will continue to fund CC (and most importantly the staff at Canada's HPC centres) at the current level for ~ 18 months while this transition is in progress and there is no anticipated interruption in services. The current hardware installations will also continue as planned. The original governance model nominally put the members in control but they only had influence once a year at the AGM. In practice the board and the executive exercised tight control over Compute Canada and were widely felt to be unresponsive to members, regions, and researchers. A primary goal of the new model is to provide a way for researchers to provide more direct input and also to dramatically increase the levels of transparency and consultation. A key suggestion is to have a researcher representative at every campus to act as a conduit for issues affecting researchers. Such roles already exist in some regions (e.g., Researcher "site-leaders" at every institution within SHARCNET in Ontario) and could be connected to the ACOR committee already in Compute Canada or some similar committee that would supercede it.

CFI Research Data

CFI Challenge 1 was to fund research data infrastructure projects – software and expertise rather than hardware. At the NOI stage there were two astronomy-related proposals submitted. The first astro-related NOI with PI Chris Pritchard was centred around CANFAR: migrating its data to Compute Canada and upgrading that platform. The migration part of the process already had significant financial support from NRC and the cooperation of Compute Canada. The CFI proposal would have extended CANFAR use cases to include simulated data, to be able to handle (store, distribute and analyse) larger datasets and novel processing modes including common pipelines for simulated and observed data. CFI rejected the application as not eligible based on the connection to NRC and CANFAR. There has been no feedback from CFI to justify this position and the NOI was never evaluated on the basis of research or deliverables.

The second astro-related proposal had as lead institution the University of Toronto with PI Bryan Gaensler. It was to develop several data processing pipelines with science targets such as climate and planets, large scale structure and big data, gravitation and rapid signal processing. CFI invited this proposal for a full application. The review committee was fairly critical of the application and it was not funded. The reviewers did not see strong links between the science areas and wanted a more detailed plan for the software infrastructure that would be developed. A primary criticism was that the applicants were not working with CANFAR/CADC who has a track record in this area.

The Catch 22 of both requiring and not requiring the involvement of the CADC is a serious problem for Canadian astronomy and must be resolved before any future competitions. The best route is probably direct consultation with CFI. Compute Canada, and organizations such as CASCA might also make submission in support of resolving this in a useful way. The highest impact might be if VPRs (e.g., the U15 research intensive universities group) were to write a letter of support as it would avoid any perceived self-interest from computing organizations. The CASCA Board is well positioned to coordinate a consultation with CFI, and clarification of this matter over the next few months is critical to respond to future CFI solicitations.

The second Research Data competitions deadline has been pushed later and no new deadlines have been announced yet. Compute Ontario is estimating a deadline of November 2016 for NOI and that there would be pressure to fund areas neglected in the first competition. There has also been informal discussion among people involved with the two astronomy related-proposals from the first round with a general consensus in favour of a single astronomy proposal next time.

CANFAR, CADC and CANARIE

CANFAR and CADC are proceeding with the plan of pushing all infrastructure into CC operations, this is called the 'C3TP' internally. The plan includes NRC tapering its data collection with CC storage supporting CANFAR storage and processing. Development activity required to fund this comes from NRC and the full phasing in should take 3 years, with noticeable impact starting in mid 2016. Current CANFAR storage physically located at USASK will be moved to SFU. Compute Canada has limited project management capability and NRC is assisting to keep things on track.

Part of CC getting development money from NRC includes them developing some process where CANFAR can apply for and receive increased storage and processing allocations, when the infrastructure to support that exists. Since CANFAR is becoming central infrastructure for Canadian astronomy, securing its long term future remains a major concern. The Committee sees good opportunities for CANFAR to grow into the national platform for astronomical data and computing. However, this will only become possible with more communication about the activities carried out at CADC and clearer dialogue about the needs of the community. Being able to demonstrate that CANFAR is a strategic, national priority both within NRC and across the universities will be important to maintain long-term support. There are proposals to rework how CANFAR operates to create broader community engagement that will be discussed at CASCA 2016.

CANFAR and CANARIE continue to work well together.