

## Computation and Data Committee Report to the CASCA Board, May 2013

Current Committee membership:

James Wadsley (McMaster U.) (Chair)	Term ends: 31 June 2013 willing to renew to 2016
Jonathan Dursi (CITA/SciNet)	Term ends: 31 June 2014
J. J. Kavelaars (HIA/NRC/CADC):	Term ends: 31 June 2014
Erik Rosolowsky (UBC/ UAlberta)	Term ends: 31 June 2015

### Status of Computing in Canada

Compute Canada (CC) is a national organization tasked with coordinating the delivery of HPC services in Canada. Canada's model is different to most countries in that all HPC (anything bigger than desktop computers, not just high-end super computing), is provided through CFI and thus Compute Canada. NSERC has largely stopped funding equipment and CFI only funds it for small groups under special circumstances (e.g. LIF) and also requires that Compute Canada manage the hardware.

In the past Compute Canada was an umbrella for regionally based consortia: Westgrid, SHARCNET, SciNET, HPCVL, Calcul Quebec and ACENET. The component institutions owned HPC equipment and the consortia managed it and represented provincial interests, including the interests of the specific institutions and their researchers. The consortia coordinated necessary matching of in-kind (institutional/vendor contributions) and provincial matches to CFI grants for both infrastructure (e.g. National Platforms Fund, NPF) and overhead/operations (Major Science Initiatives, MSI). The last major investment by CFI was the one-of NPF program of 2006 (~ \$78M) and this money (with matching and vendor contributions) was translated into around \$180M worth of equipment over the period 2007-2009. Since computers depreciate rapidly (halving in value each year with maximum lifetimes of 3-5 years), many regions of Canada will soon have no significant HPC hardware in place. One source of conflict is that CFI sees little problem with this but regions/provinces want local infrastructure and associated staff support.

There has been some pressure on CFI to provide a new NPF or similar new funding but it has been disorganized and unsuccessful. Compute Canada was meant to organize this effort but failed to be effective for many years. Much of the pressure was applied through the consortia. However, CFI has tended to express concerns about local self-interest rather than a real need. However, we also know CFI has been uncertain about how much money it would have to give. CFI also felt that there was too little national coordination and possible duplication. As a condition for funding, CFI sought a new structure for Compute Canada. A new Compute/Calcul Canada was incorporated as a not-for profit in Fall 2012 with a new board comprised of a mix of business and academic sector members (see: <https://computecanada.ca/index.php/en/about-us/governance/board>). This board severely lacks HPC experience. The board appointed a CEO, Bill Appelbe, who started in January 2013. At that time CASCA approached the board (James Wadsley and Jonathan Dursi met with the chair, Don Hathaway) and presented a white paper outlining Astronomy's major upcoming needs with respect to computing (both cycles and storage). CASCA requested the opportunity to play a role in strategic planning and emphasized the urgent need to replace Canada's aging HPC hardware and expand in areas such as providing storage and data analysis capabilities. On CASCA's behalf we expressed a willingness to participate in forums such as the Research Advisory Council planned for CC. The CEO also visited several major sites and spoke with key researchers.

The board fired the CEO with relatively little explanation in April 2013. This was one of many

circumstances that lowered confidence for all stakeholders in Compute Canada – including that of VPR/institutions, researchers/faculty and funders (Provinces/Federal Government/CFI). Specific problems include making no effort to engage researchers in their planning activities (no Chief Science Officer hired or Research Advisory Group established despite their own bylaws) and passing resolutions to stop funding consortia in favour of directly funding operations on an institution by institution basis. Many institutions rely on consortia to coordinate/share staff and operational activities and do not want to manage money directly. There are strong indications that the former CEO felt working with regions and consortia was an effective approach and influential elements in the board disagreed. A serious concern is that the board was not taking information from the community or any other stakeholders whereas the CEO had been consulting with researchers and institutions (VPR). Another key concern is the lack of transparency by the board. Many of their key decisions and ideas regarding directions for HPC are deliberately kept secret. When the CEO was fired, many key functions of Compute Canada were disrupted with no interim plan. Another issue is that the board has spent quite a bit of money on consultants, head-hunters, lawyers and the CEO with essentially nothing to show for it. Everything the CC board has claimed to have done so far has actually been done by staff working for consortia based on long standing practice predating the new board. These events raised serious issues regarding the competence of the board and its understanding of the important issues facing HPC. As a result the research community generated a petition with a letter dated May 10, 2013 to refocus Compute Canada on the major issues we face. One of us (James Wadsley) was central to promoting the letter via emails to researchers including CASCA members and large HPC users: <http://www.ipetitions.com/petition/restore-confidence-in-compute-canada/>

The letter has now been signed by 228 researchers. It was quite successful in terms of attracting attention to the problem, including the attention of the VPR at the U15 research intensive Canadian universities, VPR's generally, CFI and Compute Canada itself. Even prior to the letter, we heard that the federal cabinet was aware of the firing and implied board problems and asked the STIC (Science, Technology and Innovation Council <http://www.stic-csti.ca>) to rapidly produce a report on the status of HPC. SHARCNET researchers submitted a brief paper to STIC with similar concerns to the letter. We have no idea what cabinet is thinking but we hope that things don't get perturbed in a more business oriented direction or that researchers are further marginalized. In the week following the letter, an interim CEO, Jill Kowalchuk was appointed. Jill has experience managing Westgrid and was a previous interim CC CEO . We have received some input from her and through VPR regarding other outcomes. We were pleased to hear that the CC board felt the letter was constructive, reasonable and needed to be responded to.

A key concern remains that CC needs to pay more than lip service to the idea of making researchers part of the planning process. The former HPC LRP of 2005 was written by respected researchers/academics. This was supposed to be the function of the Researcher Advisory Council in the initial plans for the incorporated CC. In order to address the concerns of the letter, CC apparently now intends to finally constitute such a group. However, they wrote new terms of reference for that group removing its responsibility for strategic planning and to allow Compute Canada to plan without any input from researchers. All the current equipment resulted from grants and planning by researchers/academics within consortia. This was the function of the National Initiatives Committee of Compute Canada which was immediately disbanded by the board in 2012. CC replaced it with a regional committee that only serves to distribute information from CC. This still leaves researchers, consortia and provinces with no formal role in the most important aspects of planning for HPC as research infrastructure.

Compute Canada has organized a strategic planning meeting for June 6<sup>th</sup> in Ottawa. It was invitation

only and quite secretive. As a result of the letter, we heard that 2 researchers per region will now be invited but we have no details yet. We also heard this week that key groups such as CASCA and ATLAS may be approached. It is not yet clear if that means CASCA may be approached to be part of the Research Advisory Council or the strategic planning meeting. We have heard that David Schade will be working with planners at NRC and acting to provide advice to the CC board. Even though these things seem positive, historical experience indicates that strategic plans for research infrastructure should be authored by expert researchers. Anything short of that is unlikely to be effective or useful. We know that Compute Canada continues to spend quite a bit of money on consultants. Outside consultants have produced quite a few failed strategic plans over the years at considerable expense. Continued pressure on Compute Canada will be necessary until we ensure that well connected, experienced researchers will be authoring strategic planning documents.

The federal government recently allocated around \$50 million dollars for digital infrastructure. If this goes through CFI it would presumably require an equal provincial match. Even though this is basically the minimum amount required to keep HPC going at any reasonable level for the next 5 years or so, we cannot assume that is how it will be used. We know several interest groups (such as CANARIE) will be angling for this money. There are working groups writing national plans for “cyberinfrastructure” that may also play a role. Some provinces such as Ontario have also put aside money (\$24 M). Ontario intends to incorporate High Performance and Research Computing Ontario (HPRCO) this summer – essentially Compute Ontario. Ontario has also indicated that it wants more say if it is to continue matching CFI funds. This money is overdue and the time frames are unclear but it seems that funds may finally be coming. This is why it is so critical to restore confidence in Compute Canada as credible partner for CFI and the provinces.

We have also been in contact with CAUT and indirectly with other professional societies who may become usefully engaged. One product of this was a CAUT Bulletin article outlining the critical state of HPC and its role in supporting Canadian research and innovation ([www.cautbulletin.ca/en\\_article.asp?articleid=3624/](http://www.cautbulletin.ca/en_article.asp?articleid=3624/)) The HPC community has never been as cohesive as traditional disciplines and most researchers look first to their primary area, such as astronomy. Many fields, including astronomy, are increasingly dependent on HPC. It is now essential infrastructure. If we can engage researchers and discipline based groups and educate them about the need to get involved it will help us lobby for the necessary HPC support. This remains a critical issue to support the infrastructure described in the Astronomy LRP.

The HPC LRP ([https://computecanada.ca/cc\\_files/publications/lrp/LRP.pdf](https://computecanada.ca/cc_files/publications/lrp/LRP.pdf)) will be 10 years old in 2015, perfect timing for a new LRP that should get under way now. However, over the last several years researcher participation in HPC organizations and events (e.g. HPCS the HPC conference) has waned. The 2005 HPC LRP grew out of c3, a researcher driven national HPC organization that was disbanded around the time of the 2006 NPF in favour of Compute Canada. We feel that CASCA, through astronomers on the panel and/or explicitly through submissions, should help facilitate a new HPC LRP. Researchers within the HPC community are quite concerned about re-establishing the researcher voice. It isn't clear that Compute Canada is planning to allow researchers to have the central role that is required. We may need to go higher, to NRC, CFI or industry Canada to find sponsors for a new LRP/strategic planning process that will appropriately involve researchers. There may also be opportunities at the provincial level for researchers to get engaged. It may be useful for CASCA to throw its weight behind initiatives of this nature. This committee will keep CASCA informed regarding such developments.

## Recommendations

That CASCA strongly support a central role for researchers in all strategic planning for Compute Canada (and other Canadian research infrastructure). This committee and the HPC community feel that restoring strategic planning to researcher control is an essential step to restore confidence in Compute Canada. In this sense, the concerns that drove the petition and letter have not yet been addressed.

That CASCA seek to be involved with the June 6 CC planning exercise through a representative who has the confidence of the CASCA board and that CASCA contribute the January white paper on the needs of Canadian Astronomy with respect to HPC as a submission to the planning process.

That CASCA seek to be involved on the researcher side of Compute Canada through the proposed Research Advisory Council, the choice of the Research Officer and any similar or alternative committees and roles that are proposed. In addition, CASCA should insist that these aspects be representative, researcher driven and hold primary responsibility for strategic planning within Compute Canada.

That CASCA support the development of a new LRP for HPC in Canada with the goal of ensuring that Astronomy's HPC needs will be met and that a stable HPC infrastructure will be in place to support Canadian Astronomers. CASCA should help to enlist the support of other discipline-based societies and other stakeholders such as CFI, NSERC and NRC.

That CASCA engage in cooperative dialogues with regional and provincial HPC providers/consortia to help guide provincial planning and the provision of necessary resources to support Astronomy both at a regional level and for national projects managed by a specific region or consortium.

That CASCA continue to communicate to CC that astronomical data storage and much observational data analysis is heavily dependent on WestGrid resources, via a CC allocation. The current allocation process is annual in nature and, obviously, this doesn't make for stable archive process where decades of data are concerned. A different process for storage allocations seems appropriate as these are more consumables rather than recyclable.