Computation and Data Committee Report to the CASCA Board, May 2019

Current Committee membership:

James Wadsley (McMaster) (Chair)	Term ends: 30 June 2021 (OBO)
Pauline Barmby	Term ends: 30 June 2021
Catherine Lovekin	Term ends: 30 June 2021
J. J. Kavelaars (HIA/NRC/CADC):	Term ends: 30 June 2021
Erik Rosolowsky (Alberta)	Term ends: 30 June 2021 (OBO)

The current committee is nominally signed up for three years. However it is suggested that repeat members could rotate of earlier in a staggered way to make new membership smoother (e.g. Wadsley, Rosolowsky, Kavelaars). Note that Kavelaars is an ideally placed member for NRC input. If he chooses to rotate off then it is highly recommended that a replacement from NRC/CADC be found.

CANFAR and CADC

Starting in about 2015 the CADC began to work with Compute Canada to migrate services and storage to their facilities. This migration fell under the banner of the 'Compute Canada CADC Transition Plan' (C3TP). Under this plan, a copy of the astronomical data archive would be stored on CC hardware and managed via CC provisioned services, with the storage allocation managed under 'CANFAR'. Part of the motivation for this effort was the dwindling hardware resources within the CADC, confusion over mandates and difficulty working with the newly formed Shared Service Canada (SSC) which is responsible for all computing infrastructure inside the Government of Canada. An additional motivation, however, was the opportunity to place a complete copy of the CADC collection on hardware that is co-located with university research computing. This co-location provides a strong advantage to Canadian researchers with astronomy data intensive research programs. Although born of this somewhat existential crisis within the CADC, the co-location of archival data in CC computing centres had become the primary motivation for C3TP

The US community is also moving to models where the data is stored within a computing centre, but in those cases the computing centre is commercial cloud infrastructure. The commercial cloud provides hardware to store these datasets at minimal cost (often zero) while the archive centres maintain the collections and services that enable data discovery. Researchers pay for computing access on the commercial cloud system or pay to transfer the data off the commercial cloud. For the JWST dataset, the Space Telescope Science Institute is considering paying the computing bill for the first year of Webb operations. They are estimating the cost at about $US200k/year, a minor increment on the cost of operating JWST. At this time the Government of Canada is pursuing funding computing via national public infrastructure (i.e. Computing Canada) and is not providing access to funds to enable use of commercial cloud.

After 18 months of investigation, the development of a CC managed storage and computing service, that would enable CADC to transfer the management of storage to CC was determined to not be feasible under the current CC development plan. Both CADC and CC agree that such services will eventually be provided, but creating those services would not occur on a timescale appropriate for CADC. Simultaneous to this, CADC received funding to enable a new service agreement with SSC, under which SSC will provide a storage service layer for the CADC archive holdings. This new storage service is being housed at HAA and operated jointly by SSC and the CADC. About half of the total 4.5 PB is expected to be available from September 2019. Once the full deployment occurs, the physical

space within the computing room at HAA/DAO will be exhausted. Further expansion, required for projects SKA or an LSST public archive, will require the storage to be house in at a different location such as CC host site (e.g. U-Vic , SFU, Waterloo, etc.), a new computing room at HAA/DAO or a storage facility managed by SSC.

The C3TP project has been extended and refocused on enabling CANFAR to become a high-level science platform on Compute Canada. This work has focused on developing support for new services (JupyterHub, DOIs, mountable storage) and specific analysis package. This new layer is being developed against specific science use cases and in collaboration with members of the research community (e.g. CIRADA). The goal for is to provide a single environment where data from major astronomy projects (e.g. ALMA, JWST) will be collocated with computing and provide to that computing in ways that enable their direct and combined analysis. The C3TP project is scheduled for completion in spring of 2021.

The last few years within the CADC has been a period dominated by crisis management. However, we now have a period of relative stability. There is still concern, however, as the currently planned storage will not be sufficient to meet the needs of the research community beyond 2021. The evolving DRI funding landscape in Canada makes this a critical time for the astronomy community. The CASCA LRP2020 will need to provide a clear priority to the expansion of computing and storage resources needed for the archiving and analysis of data flowing from the many data intensive astronomical projects the community is engaging in.

Compute Canada Governance and Digital Research Infrastructure

Compute Canada (CC) has been designated as the main source of cycles, storage and HPC services for researchers at Canadian universities for the next 3years with funding provided for operations by CFI and $50M for a hardware refresh which will be distributed among existing sites to increase their capacity (in cycles and storage). After 3 years, Compute Canada is expected to cease operations in favour of a new, non-profit corporation with a wider mandate.

The federal government, through Innovation, Science and Economic Development (ISED) has decided on a major reorganization of computing in Canada. The new organization will have responsibilities for Digital Research Infrastructure (DRI). These expanded responsibilities cover not just hardware and user support but also research data and software. Networks will remain the responsibility of CANARIE. A call was made for applications for a new non-profit corporation to perform these tasks here: https://www.ic.gc.ca/eic/site/136.nsf/eng/home

The new organization is intended to address deficiencies currently associated with Compute Canada and CFI. These include a lack of consultation and coordination with researchers, inconsistent and unreliable funding cycles and difficulties in coordinating with provincial and local partners. In the new model, funding comes directly from ISED (no direct CFI involvement) at a level of ~$600M over 5 years. National systems (and system support staff) will be largely fully funded by ISED without matching requirements. However, ISED expects that provinces and universities will take over the majority of funding for technical support staff. New systems will be more national in some sense and thus may target specific kinds of research needs. Regions, provinces and/or institutions may also choose to build their own systems focused on local specific needs. Thus it may be expected that Compute Canada/Consortium staff currently performing training, user and software support roles may end up being associated with these local centres and funded differently. The expanded responsibilities will also require coordination with groups that currently curate data and security. There will a ramp-up

and consultation period over 2-3 years after which the new organization will take over operations from Compute Canada. Given the upcoming Federal election and the long ramp-up there is room for considerable changes to the organization, its funding and its mandate before this time.

The sole applicant, to our knowledge, is a consortium created by the U15 (Canadian research intensive universities). On their behalf, and in consultation with stakeholders, a working group wrote a proposal (now submitted to ISED). CASCA members and the CDC committee interacted with the working group and had some success improving the proposal, particularly with respect to researcher consultation and direct researcher involvement in governance and management. The submitted proposal is here:
http://engagedri.ca/wp-content/uploads/2019/05/Digital-Research-Infrastructure-Contribution-Program-Proposal-Complete-Submisson-May-6-2019.pdf

The proposed governance structure similar to that of Compute Canada, with a base membership consisting of Canada's universities and colleges with an extended set of stakeholders including librarians, information officers and researcher groups (such as CASCA). Compute Canada has an inactive user committee (ACOR). The new organization will have a new committee with stronger connections to the board (including contributing up to 2 user board members). ISED is expected to make a decision with respect this proposal this summer.

There are still significant concerns to be addressed. Currently, Compute Canada consults closely with regional organizations (e.g. Westgrid, Compute Ontario) who manage the systems and also represent provincial partners and institutions. These regional groups typically have strong researcher involvement (e.g. researcher directors and/or boards) and thus serve an important role collecting researcher concerns and passing them on to Compute Canada in a fairly comprehensive and representative fashion. It is unclear what kind of role these or other sub-national groups will have going forward. This may result in poor coordination with the regions and other levels below the national one. The proposed user council will not be representative in any sense. CDC communicated these concerns to the U15/authors. We believe CASCA could help fill the representation gap and help the new organization with its long term planning and coordination with Canada's major science investments.

A second concern is whether some aspects of current digitally-enabled research might fall through the cracks. For example, four key elements of DRI are described by ISED. Among these, research software (RS) has been defined as "enabling researchers to access and use data". Similarly advanced research computing (ARC) is defined as "involving super computers that allow researchers to analyze massive amounts of data". This appears to exclude simulations and modeling. It also overlooks software development and maintenance. Coupled with the reduced funding for local technical staff, this may undercut training in general. Canada's researchers have benefited greatly from training that is not typically otherwise available (e.g. in parallel computing and programming), provided by regional Compute Canada staff, particularly for students.

National Computing Hardware

A fifth new system Beluga, has been installed in Montreal, Quebec with roughly 34,000 cores, 688 GPUs and 12 PB of disk. The other four systems major systems are SFU's CEDAR (GPU focused ~ 30,000 cores, ~ 3000 GPUs), U-Vic's ARBUTUS (cloud focused ~15000 cores, 1.6 PB storage), Waterloo's GRAHAM (~30,000 cores)and Toronto's NIAGARA system (Large parallel system, 60,000 Intel Skylake cores, 300 TB RAM, 9 PB disk).