Computation and Data Committee Report to the CASCA Board, May 2011

Current Committee membership:

James Wadsley (McMaster U.) (Chair)	Term ends: 31 June 2013
Jonathan Dursi (CITA/SciNet)	Term ends: 31 June 2011
J. J. Kavelaars (HIA/NRC/CADC):	Term ends: 31 June 2011
Jason Fiege (U. Manitoba)	Term ends: 31 June 2013

Both Jonathan Dursi and J.J. Kavelaars are willing to continue for another term with the full support of the committee.

1. CDC Recommendations for Action by CASCA

As detailed in section 2, the situation for Canadian computing and networks is precarious. Over the past year, several issues have become critical and CASCA should step forward to help resolve them or risk experiencing severe deficiencies in infrastructure essential to Canadian Astronomy.

1.1. Lobbying for sustained HPC infrastructure

CASCA must take on a more direct role in lobbying on behalf of HPC and related facilities in Canada. Computing is not able to take care of itself. Astronomers take up 1% of discovery grants but consume 11% of Compute Canada resources. It will require a sustained effort by disciplines like Astronomy that make disproportionate and intensive use of these facilities to re-establish the importance of HPC and get Canada on a sustained and competitive funding track. Compute Canada does not currently have an effective lobbying effort under way. Securing funding will require a concerted effort across many disciplines and involving many partners, to convince government.

1.2. Addressing critical network bottlenecks

CASCA should work to raise awareness of the problem of inadequate networks and assist the various partner organizations in making the case for uniform networks at the 10 Gbs backbone speed. Local networks are typically much slower than the backbone and create severe bottlenecks. We are simply unprepared for upcoming astronomical data transfer volumes. This is a short term, high priority with respect to the overall need for a data management policy recommended in the LRP (listed in section 3).

1.3. Participation in the development of a Cyber-Infrastructure Plan for Canada

Future success with sustained funding for computing, storage and networks is likely to involve a much closer partnership between Compute Canada and Canarie as part of a Cyber-Infrastructure initiative. A key buzz-phrase in Cyber-Infrastructure is *cloud computing*, a new mechanism for accessing existing computing resources, as demonstrated by CADC's CANFAR project. By making access easier, this is likely to add many Astronomers to the HPC user base. This will accelerate the need for new HPC infrastructure necessary to process astronomical data coming from the new world observatories. Astronomers must work to make funders understand that cloud computing, rather than reducing demand or providing an alternative to traditional HPC investment, is actually an argument for a more integrated approach and increased, long-term funding for Compute Canada.

2. State of Computing, Data and Networks in Canada in 2011

High Performance Computing (HPC) is loosely defined as an order of magnitude beyond what you can do on a desktop all the way up to computing on 100,000's of processors (today's top systems). Recently CFI and NSERC have directed the majority of medium to large scale computing funding to computing consortia established over the past 10-15 years. These consortia are comprised of geographic collections of universities, colleges and other research sites (e.g. hospitals). They currently include Westgrid representing B.C., Alberta, Manitoba and Saskatchewan; SHARCNET, SciNet and HPCVL representing Western Ontario, Toronto and Eastern Ontario, respectively, RQCHP and CLUMEQ representing Quebec and ACEnet representing Nova Scotia, New

Brunswick, and Newfoundland and Labrador. There is pressure within Ontario and Quebec to consolidate into a single provincial consortium in each case and discussions are ongoing. In Ontario, consolidation is a probable requirement for continued operating support beyond 2011.

Historically, consortia would manage many smaller systems (~128 processors) that were distributed over all their sites. Increasingly, consortia are consolidating their computing into large clusters (1000+ processors) which may be off-site. There are systems are designed to meet both large and medium scale tasks. Medium scale needs include serial farming (e.g. data processing), non-parallel computing, and small scale parallel computing . Cloud computing comes under this model. Note that systems for medium scale needs may be still be large, just designed for many separate users and tasks. Elsewhere in the world, medium scale needs are more often satisfied with timely investments in research group or university in-house systems. Canadian governments and agencies created this difference and show increasing reluctance to fund computing hardware outside the consortia, insisting that showing that need for the system cannot be met by existing systems as a condition of funding. The primary drivers seem to be limited budgets and seeking of cost savings. Even though consortia are efficient managers and operators of equipment generally, (at least as efficient as commercial providers such as Amazon or Google) it is not clear that this approach provides the best solution for specialized, intensive research. For example, Consortia have been uneven in supporting data intensive needs and have tended to favour short term storage associated with computing (e.g. scratch disks with a several month expiry time) with some notable exceptions such a 0.5 PB allocation for CADC on the storage grid maintained by Westgrid (U Sask.). Data volumes are increasing beyond the resources of individual research groups and long term data storage is a key issue that HPC providers are now grappling with.

As of 2006, HPC in Canada has been collected under a single banner: Compute/Calcul Canada (CC). This was a required step as part of the new CFI funding model: the National Platforms Fund (NPF). The NPF was a direct result of the C3.ca Long Range Plan for High Performance Computing (2005) and intensive lobbying efforts by the C3 ca organization and the 7 CFI funded computing consortia. C3.ca was folded into Compute Canada becoming part of the CPAC committee. Compute Canada is structured with a governing board composed of VPs Research (or equivalent) from the universities and colleges that collectively compose the consortia, an executive director, Susan Baldwin, and various committees that provide advice to the board and director. Under the NPF model the consortia were required to make a single application for funds. However, the funds allocated to each consortium were predetermined and each consortium directly manages procurement and operation of equipment. Neither Compute Canada nor the consortia are incorporated, thus they are unable to apply for funds as distinct entities. Funds initially flow to specific universities designated by the consortia. Hardware is primarily funded through CFI (40%) with a provincial match (40%) and an industry contribution (20%) (primarily a vendor discount). Operating costs come from the CFI overhead (IOF), NSERC MRS and provincial funds where available (primarily Ontario and Quebec). Universities contribute substantial in-kind in the form of space, power and some staff costs (depending on the consortium). The new CFI MSI program may provide for some operating costs. In practice, staff funding is strongly tied to equipment funding. As a result consortia that did not receive substantial funds in the NPF allocation are experiencing severe difficulties with staff and operating costs. There is a serious risk that we will start to lose valuable technical HQP who have been training and supporting HQP in astronomy. An effective long term solution for this situation would be predictable funding on a regular 2-3 year time scale with associated IOF or equivalent.

Support for high performance computing is set to fall precipitously over the next few years. CFI last announced meaningful new funding opportunities in 2005 (in the form of the NPF1 awarded in 2006) which resulted in generally available new facilities in 2009 (e.g. SciNet's international top 20 GPC cluster) with the last major installations coming on line now (in Quebec). Thus even if new funding was fast tracked and an NPF2 competition began in 2012, the likely time scale for new facilities is 2015. It is important to note that CFI has not confirmed that there will be an NPF2 and has suggested that future funding for HPC would be competitive rather than a guaranteed component. As stated in the astronomy LRP, this will leave many regions with equipment that is 8 years old at minimum and even the newest installations will be 4 years old. This contrasts with the computing LRP (2005) recommendations for a sustained computing infrastructure with rolling refreshes such that each major centre/consortium is refreshed on 2-3 year time scales in keeping with rapid ageing of HPC facilities. These issues are discussed in the astronomy LRP, emphasizing that 2009 was a transient high point for Canada with regard to HPC capability. We are now heading into a 13 year low. Unfortunately, little progress has been made to secure new funding, let alone sustained long-term funding. In addition to the clear difficulties for processing observational data, this creates the serious prospect of an exodus of HQP who are dependent on computing and simulation capability from Canada and Canadian astronomy, including post-docs and faculty.

The C3 Long Range Plan of 2005 was a considerable success, catching the attention of government and leading directly to the creation of the National Platforms Fund. The LRP document was published in colour and was crafted to be highly accessible (high impact case studies with illustrations) and heavily focussed on the broad utility of computing across disciplines and its critical role in enabling research. The plan was heavily promoted via direct lobbying by prominent researchers and consortium directors to CFI, NSERC, Industry Canada and to ministers. C3 put together a substantial budget in support of this effort (~ $100K). In contrast, recent lobbying on behalf of HPC has not been particularly effective due to a severely under-resourced central office at Compute Canada, a less effective organizational structure and much reduced receptiveness from relevant parties including NSERC, CFI and Industry Canada. CFI has been problematic in many ways, not least through requests that CC not engage Industry Canada. Compute Canada has also had to deal with perceptions that HPC is driven by consortium interests that are somehow distinct from grass-roots demand.

Canarie, the Canadian research networks organization, has been seeking a renewed mandate in order to maintain its own long term funding. As part of this process it has engaged with Compute Canada and other interested parties with the aim of developing a Cyber-Infrastructure plan for Canada that would incorporate both networking and HPC hardware and serve needs across the spectrum from data to simulations. Some versions of the plan would merge Compute Canada and Canarie. Discussions of this nature have been going on since 2005. Points of friction in these discussions include the tendency for discussion papers to subordinate HPC as one of a dozen or so aspects of Cyber-Infrastructure while giving the central role to the networks. An additional issue has been the emphasizing of grid and cloud computing at the perceived expense of traditional HPC. Many applications do not work well in the cloud model and converting applications to the cloud/grid model requires a large investment of technical staff time. Cloud computing is chiefly an alternative mechanism to access the type of HPC known as serial farming. Thus cloud computing creates new HPC users and a need for more HPC infrastructure of the type that makes up the majority of Compute Canada's systems. Cloud computing is in no way an alternative to direct investment in HPC hardware. Increased interest in cloud computing can help make the case for increasing HPC funding closer to G8 norms (currently at about 1/3 or less).

In this context, in is important to highlight the current involvement of the CADC with Canarie in the CANFAR project, a virtual machine (VM)/cloud computing layer for data reduction on the CC infrastructure and currently using HPC systems within Westgrid. The project is nearing completion and appears to be working for real-world examples. The goal is to provide observational astronomers with an easy, configurable road into HPC. CANFAR has built VMs that run their special variants of various analysis packages and with little effort on the part of users, so the project is successful. We expect that, once the CANFAR cloud system is up and networking issues to storage outside CADC are resolved, we will have many researchers requesting substantial data-analysis jobs and CADC will then need a larger allocation from CC for computing. Right now CADC has a small (500 core) allocation for testing and general assurance from CC that this will be expanded if the need is there.

Canarie maintains an excellent network with respect to the backbone but the throughput gets progressively worse in the "last mile" leading into individual institutions. Thus most users are experiencing typical transfer rates substantially below the 10 Gbs Canarie backbone, as documented at CADC. Much of this poor performance may be associated with routers and gateways rather than the nominal network speeds though in many cases on campus networks are quite old as well. We could substantially benefit from a coordinated push to educate network groups on how to tune and maintain high performance networks. At this time CADC is saturating its 1 Gbs network connection during peak periods and running near 50% capacity most of the time . CADC has been given indications from NRC that they will upgrade the connection to 10Gbs in the near future and a final decision on this should be made after the June budget. If SCUBA2 had arrived on time we would currently be in poor shape and CADC would likely have started to shape network traffic.

3. Reference: LRP 2010 Recommendations related to Computation and Data

In the LRP 2010 executive summary the computing and networking needs are given as follows:

"Equally importantly the LRPP is concerned that computing infrastructure, as funded through Compute Canada (CC), will not meet the needs of astronomy researchers in the next decade. To analyze petabytes of data, petaflops of computing are required. This lack of resources is a critical concern for the CADC, which plays a pivotal role in Canadian astronomy by providing archived data and advanced data products to the community, as it transitions to a reliance on CC hardware. Thus the LRPP strongly recommends an increase in funding to CC comparable to that of other countries in the G8. The LRPP also reiterates the need for a Canadian data management policy. As data volumes continue to grow it is important to the future planning of the CADC to establish how these will be managed."

LRP Section 3.4, Computational Resources, CADC and Networks provides details of recent trends and potential impacts for both observational and theoretical astronomy. The key LRP recommendations put forward to address these issues were.

Recommendation 10 Compute Canada funding should be doubled to bring us up to at least 2/3 of the G8 average HPC funding per GDP. At least 1/5 of this funding should go towards encouraging user innovation through research support, and to the provision of HPC consultants. Compute Canada should also budget funds to ensure a "Tier 1" facility is available to researchers.

Recommendation 11 The LRPP recommends that Compute Canada move to fully support users with cloud computing requirements

Recommendation 12 The LRPP reiterates the need for a coherent Canadian Data Management Policy. As observatories become more dependent on data analysis, end-to-end management of data, including decommissioning archiving, needs to be a critical part of project development. A working group from the NRC (CADC), CSA, and CASCA should be formed to address this point, particularly focusing on the needs of the "world observatories".