



Budget Submission 2017 - Compute Canada

Advanced research computing (ARC) infrastructure and services are essential requirements for personalized medicine, advanced materials, genomics, engineering, environmental science and many other key strategic areas for growth in Canada. Nations around the world are creating strategic policies and support to ensure they are leaders in future scientific discoveries which are increasingly fuelled by big data analytics and scientific computing. Like the Trans-Canada Highway, this infrastructure requires a national strategy and leadership to be successful. It is larger than any one province or region can support, and it is fundamental to becoming an innovative nation. This is the infrastructure that is being used in Canada to develop new alloys for next generation aircraft, mapping environmental systems and infectious disease patterns and creating a platform for innovative new products. Many of these needs cannot be met by current commercial service providers. Like other countries, Canada needs to build its own advanced research computing delivery model to ensure the right systems are in place for the research community so that we can achieve our goals in research and innovation.

Who We Are

Compute Canada (CC), a national non-profit organization incorporated in 2012, leads the development of and coordinates the operation of pan-Canadian advanced research computing resources used for simulation, big data analysis, visualization, data storage, software, portals and platforms for research computing serving the majority of Canadian academic and research institutes. The Compute Canada federation is composed of 35 institutions who deliver advanced research computing (ARC) and research support services. Compute Canada works through partner Regional Organizations - ACENET, Calcul Quebec, Compute Ontario and Westgrid - that have been assembled from long-standing institutional consortia. The Compute Canada federation is funded by a combination of grants from the Canada Foundation for Innovation (CFI), matching grants from some provinces and funding from the 35 institutions.

The Compute Canada federation serves more than 70 institutions and 10,000 researchers across the country, including 3,000 faculty. Together the federation forms an essential large scale national facility that requires predictable, sustained operating and capital funding in order to serve the broader research community. The facility enables world-leading research in many disciplines, including digital humanities, engineering, computer science, physics, astronomy, chemistry, neuroscience, bioinformatics, and mathematics. The facility also supports

researchers from large and small research institutions from coast-to-coast, across many sectors. We enable some of the fastest growing foundational and applied research areas in Canada: environmental science, deep learning, advanced manufacturing and materials, life sciences and genomics. These sectors all require access to computational resources that are too extensive and costly for our institutions to maintain on their own. The breadth and depth of the community we serve puts us in a unique position to quantify both the costs and the benefits of fully meeting the growing ARC needs of Canada's diverse research community.

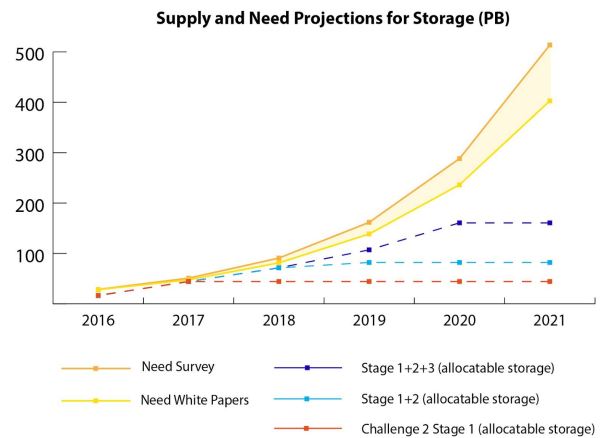
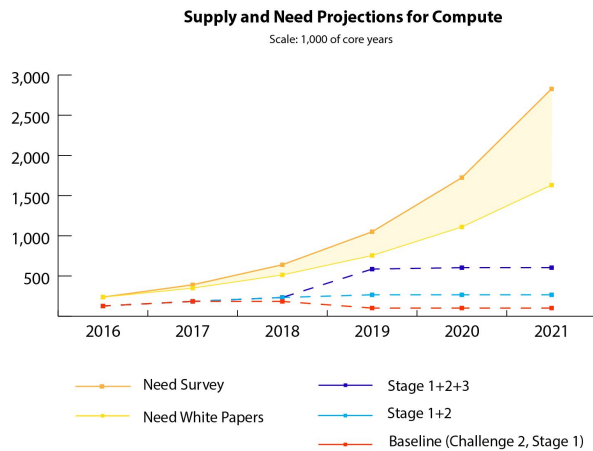
The Breadth and Depth of Canada's ARC Challenge

The diversity and growth of Compute Canada's user base demonstrates that advanced research computing is critical for a growing number of research endeavors in both the academic world and elsewhere. An analysis of the research enabled through the use of our ARC resources documents that this enabled research is high quality, generating significantly higher impact than other research conducted in Canada. This underscores how important access to ARC resources is for Canadian research excellence and global competitiveness.

At the same time, serving the needs of this broad community represents a challenge. Investment in ARC has been concentrated through Compute Canada in order to optimize the resources that can be delivered to researchers and provide greater value for money for Canadian taxpayers, yet that investment has been intermittent, and our resources are still insufficient to meet the needs of the research community. In 2016, the Compute Canada federation could only meet 57% of the technically reviewed and validated resource requests from academically funded researchers received in that year's competition (down from 84% in 2012). This represents a serious shortfall that inhibits the success of projects that have already been funded as part of the Government of Canada's investment of more than \$2.5B per year in research through grant-based funding programs.

Recently, through funding from the Canada Foundation for Innovation and matching funds from some provinces and institutions, we have been able to put in place plans to replace outmoded capacity and to provide some increases. However, our demand analysis shows that the new capacity will be insufficient to serve overall demand and we will still not be able to serve those with extremely large data analytics and computing requirements.

The need for ARC is projected to grow exponentially for the foreseeable future. Compute Canada periodically conducts a roadmapping exercise, Sustainable Planning for Advanced Research Computing (SPARC), that integrates forecasts for ARC resources prepared by the key disciplines that need those resources. **In the next 5 years the use of ARC is expected to grow seven fold in computing and 15 fold for storage and data management across a wide span of disciplines and sectors, driven by improved scientific instrumentation, and increased reliance on both high-resolution simulation and the analysis of increasingly large datasets.**



ARC is essential not only for the broader research community, but also some of Canada’s “deepest” research communities, the “big science” projects that stand out across Canada.

Our ARC resources enable the following Canadian and international collaborations:

- [ATLAS](#): We provide “Tier-2” computing and storage to the more than 150 Canadian members of the ATLAS experiment at the CERN Large Hadron Collider, an essential contribution to a global collaboration (with more than 3,000 scientists worldwide).
- [Canadian Light Source \(CLS\)](#): By storing data acquired from the CLS’ Biomedical Imaging and Therapy Beamlines, we allow CLS to focus its resources on its own areas of expertise and value-add.
- [Canadian Advanced Network for Astronomical Research \(CANFAR\)](#): We provide the primary computational platform for scientists analyzing Canadian astronomy data through CANFAR’s data portal, serving more than 4,000 international users in 2015.
- [CBRAIN](#): We operate the seven largest computing platforms accessible through the international CBRAIN project, which makes brain images and associated computational resources available to researchers around the world.
- [IceCube](#): Compute (including GPU) and data storage resources provided by the Compute Canada federation represent a significant portion of Canada’s contribution to the IceCube Neutrino Observatory.
- [International Human Epigenomics Consortium \(IHEC\)](#): We host the international data portal for high-resolution human epigenomic maps for normal and disease cell-types, with 2600 international visitors in 2015.
- [LIGO](#): Our resources were used to run simulations that verified the experimental detection of gravitational waves by the Laser Interferometer Gravitational Observatory.
- [Ocean Networks Canada \(ONC\)](#): By providing long term data storage for ONC, we allow ONC to focus its resources on its own areas of expertise and value-add.
- [SNOLAB](#): We support data analysis for several major experiments at the SNOLAB underground laboratory, including the SNO+ and DEAP experiments.

- [TRIUMF](#): We support data analysis for several initiatives at TRIUMF, Canada’s national lab for nuclear and particle physics, including the GRIFFIN, Tigress and PiENU experiments.
- [T2K](#), [Belle-2](#): We provides significant compute and storage resources in support of these two major international particle physics collaborations, based out of Japan.

These are just a few of the large scale national facilities that themselves depend on advanced research computing, and the Compute Canada federation in particular, to fulfill their objectives.

Key Recommendations

“We need Infrastructure that supports change” – PM Trudeau, Davos, January 20, 2016

In order to increase our global competitiveness and to meet the growing needs of the Canadian research community, we need to treat advanced research computing (ARC) as an essential large scale national facility, the way advanced networking was supported a generation ago. Investment and operation of such a facility should be guided by the Government of Canada’s planned strategy for digital research infrastructure.

Based on the needs analysis described above, supporting excellent academic research in Canada would require approximate ongoing investments of

- \$45 million annually for operations, including the development of new services that improve the productivity and impact of Canada’s researchers,
- \$60 million annually for capital for both compute and storage infrastructure.

These investments may need to be adjusted to reflect the objectives of a national strategy. They need to be planned and coordinated at a federal level in partnership with provincial counterparts, particularly to the extent provincial and institutional resources will be drawn upon. Planning and coordination is also needed with other federal research investments, especially with the other large scale national facilities that crucially depend on ARC resources to fulfill their missions.

Benefits to Canada

Ready access by Canada’s most competitive researchers and innovators to advanced research computing resources is certain to enable increased impact from that research. Whether modeling combustion in a jet engine, the movement of drugs and other molecules through biological environments, the effects of climate change on the ocean and atmosphere, or the collision of two black holes, access to state of the art ARC resources will directly enable innovation in advanced manufacturing, healthcare, agriculture, green technology and other critical economic sectors. Canada will be better equipped to innovate and compete globally, and Canadians will benefit from improved health while reducing negative impacts on our environment.