



Budget Submission 2018 - Compute Canada

Advanced research computing (ARC) is the key underlying infrastructure of the data revolution. ARC is essential infrastructure for artificial intelligence, precision agriculture, personalized medicine, advanced materials, clean energy technologies, genomics, engineering, climate modelling, drug design, criminology, neuroscience and basic research in fields such as physics and astronomy. Nations around the world are creating strategic policies and support to ensure they are leaders in future scientific discovery, which is increasingly fuelled by big data analytics and scientific computing. This infrastructure requires a national strategy and national leadership to realize its benefits. Canada's ARC infrastructure is larger than any one province or region can support, and it is fundamental to our role as an innovative nation. Like other countries, Canada needs to build and fund its ARC delivery model to ensure the right capabilities are in place for the digital transformation of research activities, to achieve our goals in research and innovation. We have a drastic shortage in resources to support both recent investments in science and innovation in such areas as artificial intelligence and precision agriculture, as well as support for small to medium size enterprises, which is currently limiting innovation and R&D based company growth in Canada.

Who We Are

Compute Canada, a national non-profit organization incorporated in 2012, leads the development of and coordinates the operation of pan-Canadian ARC resources used for simulation, big data analysis, visualization, data storage, software, portals and platforms serving the majority of Canadian academic and research institutes. The breadth and depth of the community served puts Compute Canada in a unique position to quantify the costs and the benefits of fully meeting the growing ARC needs of Canada's diverse research community.

Today Compute Canada serves almost 12,000 researchers, including almost 3,500 faculty working at more than 100 academic institutions and research institutes across the country and their industrial and international partners. This base of research users has been growing almost 10% per year over the last 5 years. In addition, today another 10,000 users access Compute Canada's national platform through researcher-designed portals, to access specialized data and computational services. Compute Canada is an essential large scale national facility that requires predictable, sustained operating and capital funding in order to serve the broader research community. The facility supports researchers from large and small research institutions from coast-to-coast, across many industrial sectors. Compute Canada enables some of the fastest growing areas of innovation in Canada: precision agriculture, artificial intelligence, 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, provinces or small enterprises to maintain on their own.

The Breadth and Depth of Canada's ARC 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 these investments are unpredictable, with decade long gaps in capital funding, and the resources remain insufficient to meet the needs of the research community. In 2017, Compute Canada met 58% of the technically reviewed and validated resource requests from federally funded academic 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.

Based on Compute Canada's most-recent needs analysis, over the next 5 years the use of ARC is expected to grow seven fold (48% annual growth) in computing and 15 fold (72% annual growth) for storage and data management across a wide span of disciplines and sectors, driven both by improved scientific instrumentation, and increased reliance on both high-resolution simulation and the analysis of increasingly large datasets. Furthermore, Canadian R&D based SME's such as Fusion Genomics, Eosense, Vox Pop Labs and others are increasingly challenged to find access to this essential infrastructure.

Examples of Major Science Initiatives Supported by Compute Canada

ARC resources enable the following Canadian science and R&D based industrial initiatives including:

- [ATLAS](#): Compute Canada provides "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, Compute Canada allows CLS to focus its resources on its own areas of expertise and value-add.
- [CBRAIN](#): Compute Canada operates 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.



- [International Human Epigenomics Consortium](#) (IHEC): Compute Canada hosts the international data portal for high-resolution human epigenomic maps for normal and disease cell-types, with 2600 international visitors in 2015.
- [Ocean Networks Canada](#) (ONC): By providing long term data storage, Compute Canada allows ONC to focus its resources on its own areas of expertise and value-add.
- [SNOLAB](#): Compute Canada supports data analysis for several major experiments at the SNOLAB underground laboratory, including the recent Nobel Prize.
- [TRIUMF](#): Compute Canada supports data analysis for several initiatives at TRIUMF, Canada's national lab for nuclear and particle physics

Examples of Small and Medium-sized Enterprises supported by Compute Canada

- Eosense is located in Dartmouth, NS. Their soil flux chambers and gas probes measure environmental gas flux and concentration. Eosense instruments are deployed around the globe offering excellent performance in challenging field conditions. The Eosense R&D team creates new ways for researchers to capture high quality data using Compute Canada resources.
- Fusion Genomics developed the first ever screening technology in the environment for viruses - only system in the world - focusing on air close to pig farms and poultry farms to detect bird/swine flu before it infects other sites. Fusion Genomics invented this quantum technology using Compute Canada resources.

Large Industry supported by Compute Canada

Industrial collaborations are active across all economic sectors including finance, advanced materials and manufacturing, artificial intelligence, drug development and aerospace, automotive and agricultural sectors with global companies such as Rolls Royce and other automotive manufacturers, Merck, Bombardier, Thomson Reuters etc.

Key Recommendations

In order to increase our global competitiveness, meet the growing needs of the Canadian research and innovation communities and continue to train and attract talent for the knowledge economy, 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 directed to support the Government of Canada's planned strategy for innovation and science.



1. Move from project-based funding to a sustained infrastructure funding model guided by a national supercomputing strategy designed to meet Canadian research and innovation priorities;
2. Support and fund appropriate national research data management efforts to protect and gain value from this national resource;
3. Engage and grow R&D based companies in Canada; currently there is no official platform for these companies which currently have limited access to ARC facilities;
4. Create common middleware software services capable of supporting innovative science of all types, optimizing investments, and helping to ensure enhanced cybersecurity.

Compute Canada is contributing significantly to the efforts of the Leadership Council for Digital Research Infrastructure, which is developing recommendations for the Government of Canada's DRI strategy.

Based on the needs analysis described above, supporting excellent academic research in Canada at a steady state of capacity growth and ongoing technology refresh requires the following investment¹:

- \$52 million annually for operations, including the development of new services that improve the productivity and impact of Canada's researchers. This annual total includes \$32 million to operate a national ARC platform adequate to serve the growing capacity needs of our research community, and \$20 million to provide end user support, outreach and training to that community across Canada;
- \$60 million annually for capital, for both national compute and storage infrastructure, starting with \$20 million in each of the next two fiscal years in addition to the CFI's existing plans for its Cyberinfrastructure Initiative.

These investments could be shared among the federal government, provincial governments and regional initiatives, and research-intensive institutions, including those hosting infrastructure.

The chart below illustrates how this investment plan would allow Compute Canada to meet projected demand for academic compute and storage resources to maintain global competitiveness.

¹ All figures based on projections for the fiscal year ending March 2023.



% of Demand Met at \$60M/year



Supporting Canada's aspirations for research and innovation leadership requires enhanced levels of investment. For example, recently announced initiatives in artificial intelligence will require additional investments in shared resources, particularly GPU devices for computation. We estimate that adequately supporting the AI initiative alone would require an additional \$10-20 million capital investment per year. More broadly, Canada's major research universities, represented by the U15, are developing strategies to bring Canada to the top 3 in OECD ratings for researcher productivity and outcomes. Further resources are needed to serve the research needs of business, especially small to medium-sized enterprises.


Meeting these aspirational levels of computational capacity and associated support would require increased levels of sustained investment:

- \$61 million annually for operations, including the development of new services that improve the productivity and impact of Canada's researchers. This annual amount includes \$41 million to operate an internationally competitive national ARC platform and \$20 million to provide end user support, outreach and training across Canada,
- \$90 million annually for capital, for both national compute and storage infrastructure, starting with \$20 million in each of the next two fiscal years in addition to the CFI's existing plans for its Cyberinfrastructure Initiative.



As above, these investments could be shared among the federal government, provincial governments and regional initiatives, and research-intensive institutions, including those hosting infrastructure.


The chart below illustrates how this enhanced investment plan will allow Compute Canada to meet projected demand for academic compute and storage resources to maintain global competitiveness.





Depending on the level of investment, Canada is projected to reach either #6 or #3 in the world based on aggregate “Top 500” gigaflops/researcher, as illustrated below:

Canadian Investment Scenarios vs. Comparator Countries



Benefits to Canada

Canada’s Innovation strategy cannot be realized without the appropriate digital infrastructure on which it relies. Ready access by Canada’s most competitive researchers and innovators to advanced research computing resources is essential to our future as leaders in the knowledge economy, to build and train our experts in information technology and machine learning and to ensure diversity and engagement in these fields. 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, 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. We are currently unable to serve the growing requests from small to medium R&D firms in Canada.



Conclusion

We are missing a key component in the discussion around digital infrastructure. We talk of networking to transfer big data, secure places to store it, govern it, archive it and manage it. A critical step, in order to extract value, is to compute with it. Supercomputers are the modern day workhorse where discovery and innovation occur and where research software runs. We have the opportunity to realize the benefits here in Canada with a national supercomputing strategy.

Appropriate digital research infrastructure to enable an innovation nation is critical to Canada's success. The costs of delivering this capability exceed the capacity of any one province, institution or federal department, so we need to create shared infrastructure, delivering services nationally, with funding shared appropriately among many stakeholders. Such common infrastructure can enable all sectors and disciplines, across the full pipeline of innovation from invention to invoice.

Compute Canada Calcul Canada
155 University Avenue, Suite 302
Toronto, ON M5H 3B7
(416) 228-1234