compute
canada

Part A: Canadian Foundation of Innovation (CFI) Major Science Initiative

2017-2022 | Proposal

Contents

155 University Ave, Suite 302 Toronto, Ontario, Canada M5H 3B7
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

Compute Canada (CC) is simultaneously the national advanced research computing (ARC) facility of Canada, a Major Science Initiatives (MSI) project funded by CFI, a not-for-profit corporation formed to manage that project, and a federation of 35 institutions who are members of that corporation. The 35 institutions are partners in the MSI project and collectively own the infrastructure and employ the skilled personnel that comprise the Facility. CC's skilled personnel are themselves an essential resource for Canada, working together to help CC users accelerate and amplify their own research achievements. The federated CC Team has been assembled from long-standing institutional consortia that now participate in CC as partner Regional Organizations: ACENET, Calcul Quebec, Compute Ontario and WestGrid.

155 University Ave, Suite 302 Toronto, Ontario, Canada M5H 3B7
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

1. Scientific Excellence

1.1 The Scope of Research Enabled by the Facility

As described in the Compute Canada Strategic Plan, it is part of the CC mandate to serve ARC needs at any scale, for any discipline, for all of Canada. CC supports project scales ranging from a single faculty member, up to the largest "Big Science" projects in the country. 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. The diverse community needs are met through CC's delivery of services, which include traditional tightly coupled High Performance Computing (HPC), High Throughput Computing (HTC), serial computing, cloud computing and storage, visualization and other technology solutions as required.

CC services are critical enablers of Big Science in Canada, including within international collaborations. This includes several other MSI-funded facilities as well as major facilities funded outside MSI:

- ATLAS: CC provides "Tier-2" computing and storage to the more than 150 Canadian members of the ATLAS experiment at the CERN Large Hadron Collider. This contribution serves the full international collaboration (more than 3,000 scientists).
- Canadian Light Source **(CLS)**: CC stores the data for Biomedical Imaging and Therapy Beamlines.
- Canadian Advanced Network for Astronomical Research **(CANFAR)**: CC serves as the primary computational platform for scientists analyzing Canadian astronomy data. The data portal associated with CANFAR had more than 4,000 unique international users in 2015.
- CBRAIN: CC provides the seven largest computing platforms for the international CBRAIN project, which makes brain images and associated computational resources available to researchers around the world.
- IceCube: Canada's contribution to the IceCube Neutrino Observatory includes significant compute (including GPU) and storage resources from CC.
- International Human Epigenomics Consortium **(IHEC)**: CC hosts the data portal which provides high-resolution human epigenomic maps for normal and disease cell-types for this international consortium. The site had 2600 unique international visitors in 2015.
- LIGO: Canadian participants in the Laser Interferometer Gravitational Observatory rely on CC resources for detailed simulations.
- Ocean Networks Canada **(ONC)**: CC provides long-term storage for all ONC data.
- SNOLAB: CC supports data analysis for several major experiments at the SNOLAB underground laboratory, including the SNO+ and DEAP experiments.
- TRIUMF: CC supports 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: CC provides significant compute and storage resources in support of these two major international particle physics collaborations, based out of Japan.

155 University Ave, Suite 302 Toronto, Ontario, Canada M5H 3B7
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

CC also supports numerous platforms and partnerships which span the country and are integral components of international networks in tracking infectious diseases, environmental monitoring, climate modelling and other globally significant areas. Furthermore, CC already provides resources for prototyping future large-scale international projects, including the Square Kilometre Array (SKA).

Thanks to improved instrumentation, observational and experimental datasets will continue to grow at a tremendous pace. The challenge to store and manage complex datasets, once solely the domain of the aforementioned Big Science projects, has now become applicable to many individual researchers. The concomitant need for computation to enable data processing is also extending to an ever-broadening community of researchers. For example, the reduction in cost for next-generation DNA sequencers (NGS) is causing rapid expansion in the types of problems that can be considered in a variety of fields, from agriculture, to natural resources, to human health. Within the field of artificial intelligence (AI), the rise of deep learning has been enabled by growth in computational power. As AI techniques are applied to a wider variety of problems in a variety of fields, the need for ARC will grow.

At the same time, the computing power required to model complex systems continues to increase. 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, the leading Canadian researchers require increasing computational resources.

The rise of cloud computing has also begun to bring the power of ARC to a broader research audience. The ability to provision a virtual machine and give the user control over his or her software environment is changing the way researchers work and the way CC administers systems. As cloud technologies advance, we expect the line between cloud and other ARC services to blur. For example, research projects in the digital humanities, which today use cloud to text-mine single articles, may wish to analyze a collection of 10,000 articles by launching the tool-set on thousands of cores in parallel. Researchers on the cutting edge will need access to facilities and services which allow them to experiment with HPC, virtualization, containerization and emerging technologies, with support from CC's expert staff.

As a reflection of its national role, in the last two years CC has signed or is in the process of negotiating partnership agreements with several national and international organizations:

- Canadian Association of Research Libraries (CARL): CC and CARL signed an MOA to work on a joint Research Data Management (RDM) solution for Canadian researchers.
- Mitacs: CC and Mitacs signed an MOU to promote Mitacs innovation training and fellowship initiatives through the CC network of sites and users.
- European Grid Infrastructure (EGI): CC and EGI signed an MOU to enable work towards platform interoperability between Canadian and European e-infrastructures.
- Software Carpentry: CC is the national partner in Canada for Software Carpentry, an international organization focused on training researchers with basic skills to accelerate their use of ARC.
- Pacific Institute of Mathematical Sciences (PIMS): CC and PIMS are negotiating an agreement to make a PIMS-created (and maintained) platform available to researchers across Canada.
- Women in HPC: CC is a founding member and has signed an MOU with Women in HPC.

In the next 5 years the use of ARC is expected to grow across a wide span of disciplines and sectors, driven both by improved scientific instrumentation, and improved ARC technologies.

155 University Ave, Suite 302 Toronto, Ontario, Canada M5H 3B7
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

1.2 The Community of Facility Users

The CC user base, as measured by the number of active faculty research groups, has grown substantially over time. It has increased from 1,250 faculty members at the time of the original MSI proposal in 2012 to nearly 3,200 as of early 2016. In all, more than 10,000 researchers, in a wide range of disciplines, are using CC resources as an essential tool.

As illustrated in the charts below, the CC platform serves a diverse and comprehensive community, from traditional HPC disciplines such as engineering, physics and chemistry, to more emergent ARC communities such as medical sciences and humanities. Researchers are served from coast-to-coast and at all phases of their career from undergraduate to graduate student to postdoc to faculty member.

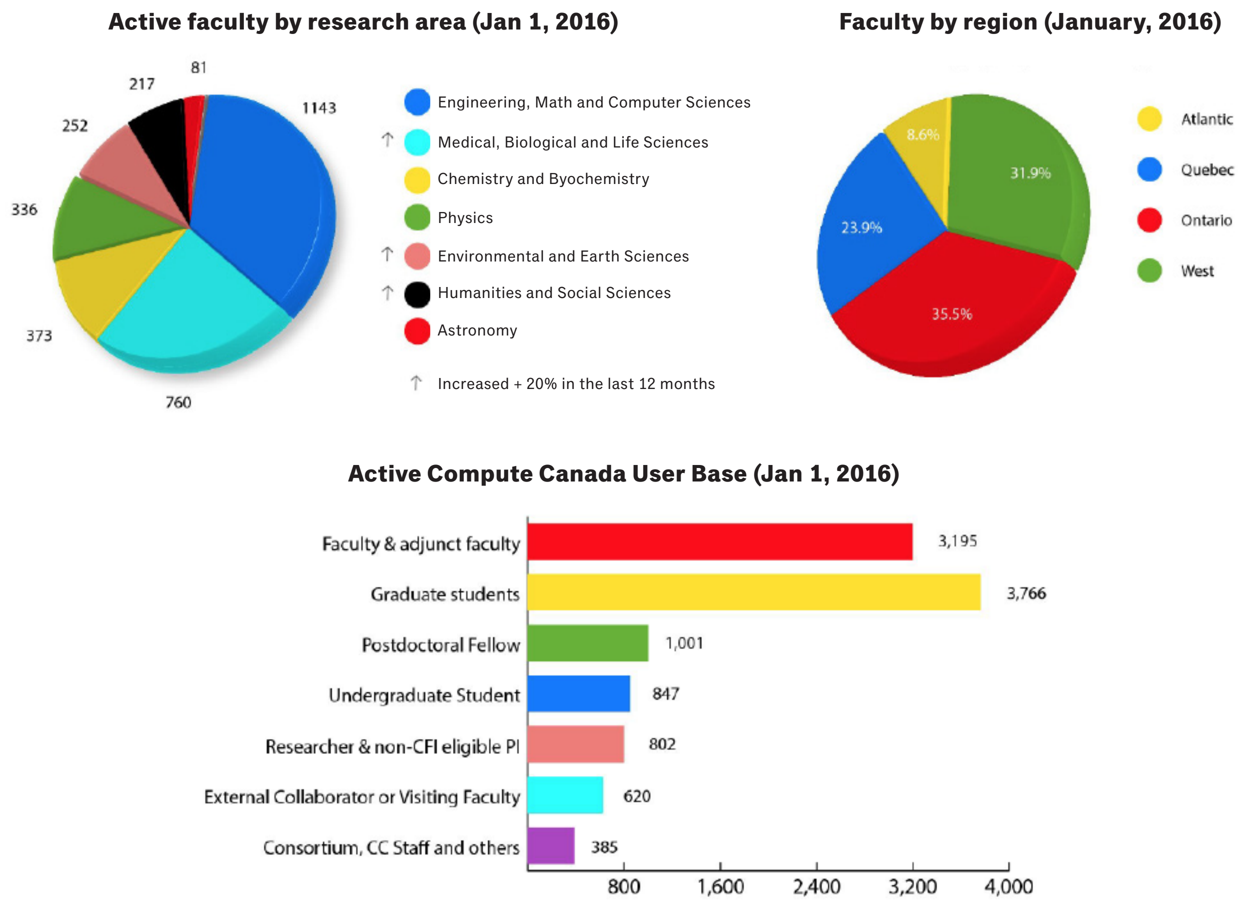

Figure 1.1: Compute Canada active researchers as of January 1, 2016.

From January 2015 to January 2016, CC saw roughly 16% growth in active faculty accounts. Year-over-year growth is highest in Environmental and Earth Sciences (36%), Medical Biological and Life Sciences (26%), and Humanities and Social Sciences (20%).

The introduction of the Research Platforms and Portals (RPP) competition in 2014 has promoted use of the CC platform to serve shared datasets and shared toolsets to the world. Thousands of users from around the world now access CC-supported platforms such as CANFAR, CBRAIN, GenAP, and IReceptor.

155 University Ave, Suite 302 Toronto, Ontario, Canada M5H 3B7
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

1.3 The Transformative Research Enabled by the Facility

Many high-profile Canadian researchers rely on CC. CC's users include 46% of the prestigious Canada Excellence Research Chairs, 32% of NSERC Canada Research Chairs, and more than 25% of Canada's Highly Cited Researchers (all data from Thomson Reuters, 2014). Consider the winners of E.W.R. Steacie Fellowships, which are prestigious awards for young faculty members in Canada. In each of the last four years, at least one fellowship winner has been a significant user of CC computational resources to support his or her work. Another example is the Herzberg Medal, the highest honour awarded by NSERC, which recognizes a body of work spanning a career. Four of the last five Herzberg Medal winners have been CC users, including the current winner Victoria Kaspi, the first female to win this award. Two Canadian Nobel Prize winners (John Polanyi and Arthur MacDonald) hold current CC resource allocations.

1.3.1 Analysis of Science Impact of the Facility

CC has collected Canadian Common CVs (CCVs) from more than 2,300 of its Canadian faculty users. This exercise produced a list of more than 70,000 research contributions since 2010. The researchers identified more than 33,000 CC-enabled contributions from this list, including more than 25,000 journal articles and more than 8,000 conference publications.

An analysis has been performed on 16,387 CC-enabled articles for which a Digital Object Identifier (DOI) could be obtained. The chart below provides a breakdown of those articles by field.

Publications enabled by Compute Canada 2010 - 2016 binned by Discipline

Cut-off date: March 21, 2016

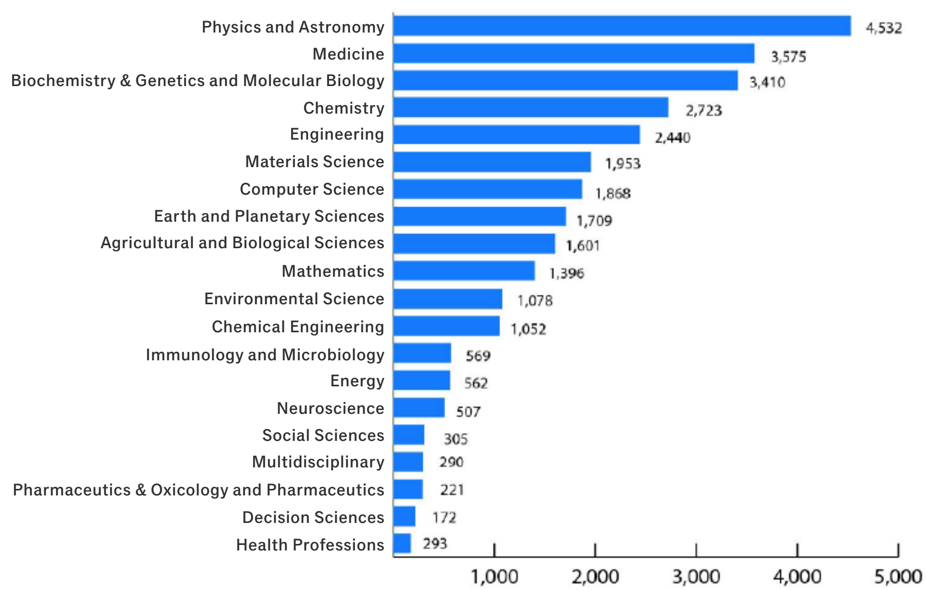

Figure 1.2: CC-enabled publications by discipline.

The large range of publication and citation rates across the many disciplines served by CC limits the value of raw article or citation counts. To assess impact consistently across disciplines, we use the Field Weighted Citation Impact (FWCI), which is the ratio of the number of citations obtained by an article to the world-average in the same discipline. The world-average in each discipline is then set to 1.0, and the FWCI is quoted relative to that baseline.

155 University Ave, Suite 302 Toronto, Ontario, Canada M5H 3B7
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

In the chart below, horizontal bars show both the CC-enabled FWCI and the Canadian average FWCI, over a number of disciplines in which at least 100 CC-enabled articles were identified. In every case, the CC-enabled publications have impacts above the Canadian average. The red bar represents the publication-weighted average across all disciplines, and the overall FWCI is double the world-average (32% above the Canadian average).

The chart's ranking of comparative impact illustrates that the largest effects of CC support relative to the Canadian average are in emerging disciplines in relatively early stages of ARC adoption. One interpretation of this trend is that early adopters are seeing large benefits in impact, while in traditional disciplines, with high ARC penetration, ARC-enabled publications are then naturally closer to the Canadian average. Traditional ARC disciplines such as engineering, physics, computer science and chemistry report approximately 40-50% of their publications as enabled by CC while for social science and humanities researchers the number is 25-30%.

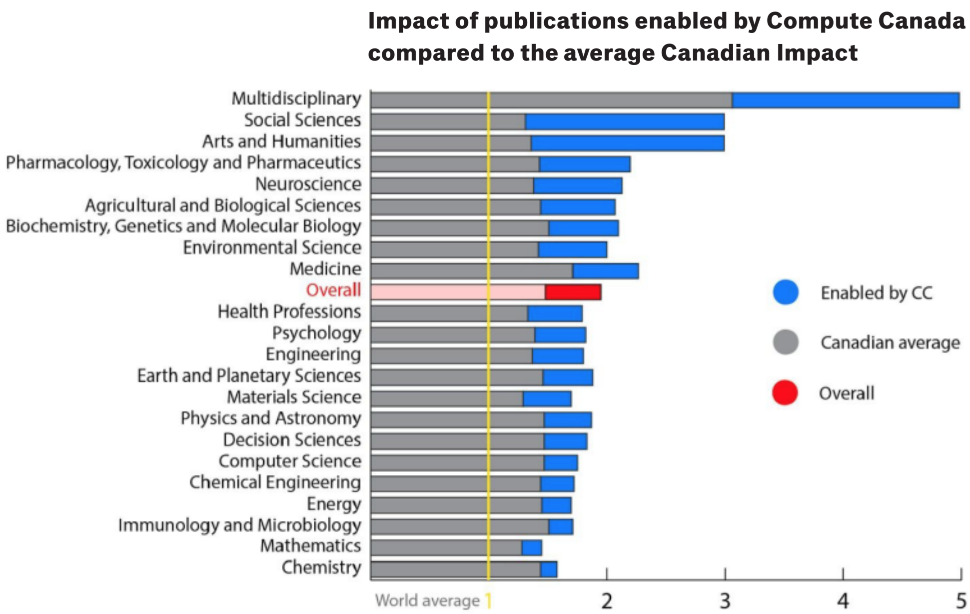

Figure 1.3: FWCI of CC-enabled publications

1.3.2 Key Accomplishments of CC Users

CC supports many of Canada's best researchers, and does so over a broad range of disciplines. This section provides a brief summary of recent research impact of the 5 selected principal users of the facility as well as a list of other recent highlights. The 5 principal users are leaders in the Canadian community and are chosen to represent the breadth and depth of impact enabled by CC.

Dr. Alan Evans (McGill University, Centre for Integrative Neuroscience) is focused on mapping the human brain. "Brain mapping" refers to the combination of brain imaging with sophisticated computational analysis to capture 3D maps of neuroanatomy and neurophysiology. His lab has developed an algorithm for the extraction of 3D surfaces of the brain cortex and can compare it in normal versus diseased states across thousands of people. CC's infrastructure is used to support the Canadian neuroimaging community with a large neuroimaging toolbox through the CBRAIN portal. Hundreds of researchers across the country process their research through this portal.

Impact: In 2013, this group published the first high-resolution 3D human brain reconstruction in the journal Science. This was named one of the top 10 breakthrough technologies by MIT Technology Review. In 2014, Thomson-Reuters named Evans one of the most cited researchers in "Neuroscience and Behaviour".

Dr. Randall Martin (Dalhousie University, Physics and Atmospheric Science), applies satellite observations and global chemical transport models to advance the understanding of atmospheric composition, and its effects on air quality, climate and biogeochemical cycling. He co-leads the development of the GEOS-Chem global chemical transport model that is used by about 100 research groups worldwide.

Impact: One of the achievements of the research team using CC resources is the first satellite-based estimate of long-term changes in global fine particulate matter [Environ. Sci. Technol, 10.1021/es502113p, 2014]. These data are being used by a variety of international organizations including the Organisation for Economic Co-operation and Development as well as the World Health Organization to assess global air quality and its implications for human well-being. Particulate matter has become the leading risk factor for environmentally-related premature mortality; improvements in outdoor air quality could reduce the 3M premature deaths annually associated with this risk factor.

Dr. Susan Brown (Guelph and University of Alberta, English and Film Studies) is the director of The Canadian Writing Research Collaboratory (CWRC), which is charged with providing an open web-based environment to foster the use of digital tools and resources for literary studies in and about Canada. The Collaboratory is developed around active research projects and comprises two major elements, a database and a toolkit, linked through a web-based, service-oriented architecture. The infrastructure facilitates digital literary research and provides a testbed for computational research, ranging from text analysis and visualization to the design and use of social networking tools.

Impact: CWRC extends the work of the Orlando Project, a mature humanities computing research project that continues to publish new findings and serves as an experimental dataset. Poised to launch a beta version in May 2016 and the full version in September, CWRC hosts 22 active projects, and has trained 115 scholars in research computing. Substantial digitization by one contributing project is already underway, and others are poised to start. Partner projects such as Canada's Early Women Writers and Editing Modernism in Canada have published significant new research findings, including, in the latter case, a smartphone app based on the writings of Sheila Watson.

Dr. Yoshua Bengio (Université de Montréal, Computer Science) is a leader in the area of deep learning, which has been immensely successful in the last three years. Deep learning has been taken up by major corporations such as Google, Facebook, Microsoft, Apple, Intel, Qualcomm, Samsung, and Baidu for speech, computer vision and natural language applications.

Impact: In 2013 Bengio's lab won the international competition for emotion recognition in the wild (from movie clips, including both audio and image), using deep learning [ICMI '13 pp. 543-550]. Over the past year, they have developed new algorithms for deep learning that aim to replace the state-of-the-art in machine translation, and they have reached [arXiv:1409.0473] and then beaten (WMT'2015) the performance levels of predecessor systems in just a few months of research. There have also been significant theoretical advances in unsupervised learning algorithms for deep networks in Dr. Bengio's lab over the past two years, especially using autoencoder-based and adversarial generative models. Theoretical advances on the optimization of neural networks were selected by the La Recherche magazine among their 10 chosen discoveries of the year 2015.

Dr. Robert Wolkow (University of Alberta, Physics), is focused on the principle limitation facing the computer electronics industry today, specifically, that microprocessors generate too much heat, preventing further improvements in speed and computing power. While retaining silicon as the circuit platform, he is doing away with the transistor and using instead new atom-sized quantum dots as the main circuit building block. The new quantum dots allow electrons to be stored and manipulated efficiently to achieve computations while consuming very little power.

Impact: In 2015 Robert Wolkow received Alberta's highest award that recognizes the transfer of scientific developments to the commercial realm, the ASTECH award. The new Alberta-based company Quantum Silicon Incorporated has acquired the patents Wolkow's team has generated at the National Institute for Nanotechnology and the University of Alberta and is building the world's first atom scale electronic devices. This technology could be transformative. Compute Canada's contributions make possible a close coupling of state of the art computational work with world-leading academic and industrial laboratories.

Other high profile recent contributions enabled by use of CC resources include:

- Edward Sargent (UofT) used CC resources to perform the Density Functional Theory (DFT) calculations needed to screen potential catalysts in the process to split water into its constituent parts (hydrogen and oxygen), before attempting to produce the catalyst in the laboratory. The result, published in Science in March 2016, was a tungsten-based catalyst that is 3x better than the previous world record-holder.
- The observation of gravitational waves by the LIGO experiment (February 2016) relied heavily on detailed simulations for interpretation of the results. Simulations performed by Harald Pfeiffer (UofT) using CC resources, including those shown at the public announcement, were key to finding the gravitational waves and to identifying the source of the waves as a pair of colliding black holes.
- The 2015 Nobel Prize in Physics was shared by Arthur MacDonald (Queens) for his leadership of the SNO experiment, which discovered neutrino oscillations. CC (and its regional partners) have supported SNO data analysis for more than 10 years. Ongoing SNOLAB experiments such as SNO+ rely on CC for their ARC needs.
- In 2015, Michael Bowling (Alberta) was the first to solve an imperfect information game (Texas Hold'em Poker). This artificial intelligence breakthrough was published in Science and was supported by a special CC resource allocation.
- In 2015, the New Horizons space probe broadcast to the world the first human encounter with Pluto. Data from the Canada-France-Hawaii Telescope, served from the CC-supported CANFAR platform, was key to allowing the probe to avoid collision hazards on its journey.
- The observation of high-energy cosmic neutrinos at IceCube was named the physics breakthrough of 2013 by Physics World magazine. Through the work of Darren Grant (Alberta), and support of CC resources, Canada was one of the largest contributors of computational resources to that discovery.
- In 2012, the ATLAS and CMS experiments at the CERN Large Hadron Collider announced discovery of the elusive Higgs Boson, responsible for giving mass to the fundamental particles. CC provides all so-called Tier-2 Canadian computing for ATLAS, and provided a special resource allocation in support of this specific result in the months leading up to the announcement. This discovery led to the 2013 Nobel Prize in Physics being awarded to Englert and Higgs.

1.4 Benefits to Canada

As described in the previous section, CC provides critical support for academic researchers across a broad range of disciplines. The single most significant benefit to Canadians realized to date is the support of more than 3,000 research programs across the country over the last 5 years. Collectively, the work presented in the 33,000 CC-enabled papers represent significant socio economic benefit to the country. The quantity and quality of those papers represents benefits similar to that of a major research university.

These federally funded research teams leverage their access to CC resources to attract industry partners and apply their research efforts towards generating solutions to complex industrial challenges. Medical researchers also use CC resources to directly address the health of Canadians. There are also benefits to Canadian culture through CC support of digital humanities research and to social policies through support of social science research. Through support of this research, CC takes part in training highly qualified personnel in fields as diverse as aerospace engineering and film/entertainment. The supported researchers are working closely with national and global industries across all sectors and a number of government agencies in Canada, Europe and the U.S.

CC's support of the digital humanities includes projects with a strong historical or cultural focus on Canada, including the Canadian Writers Research Collaboratory, FolkwaysAlive!, and the Ukrainian Folklore Archive.

Social Sciences researchers in Canada use CC resources to study many aspects of Canadian life which have a strong influence on Canadian social policy. As an example, Dr. Patricia Brantingham studies computational criminology, in particular the geography of crime in urban centres. This research has implications for urban planning and allocation of police resources. There are also ongoing discussions with the Canadian Research Data Centre Network, Statistics Canada, and CC to improve the access of thousands of researchers to significant computational resources to analyze Statistics Canada datasets.

In the last 5 years, CC investigators have reported 188 R&D collaborations with industry, 116 instances of technology transfer, 83 involvements in creation of a start-up company, 183 technology product or process developments, and 102 instances of consultation for industry. As a cohort, CC faculty users have reported more than 2000 patents. Furthermore, more than 148 registered user groups have reported receiving funding from NSERC's Collaborative Research and Development program which supports industry-relevant research by university researchers and their private-sector partners. Two examples of industrial collaboration are described in more detail:

Vancouver Coastal Health Research Institute, & UBC: Dr. Artem Cherkasov & Dr. Paul Rennie

A potential new treatment for prostate cancer was developed by researchers at the University of British Columbia (UBC) and the Vancouver Coastal Health Research Institute by utilizing advanced research computing. Developed by a research team led by Dr. Artem Cherkasov and Dr. Paul Rennie, the treatment is designed to outsmart cancer that has become resistant to current treatments. This new drug attacks a receptor that promotes tumour growth by binding itself to a specific piece of essential DNA. This means that mutations to escape binding would not be viable, offering hope that a drug designed to exploit it could be effective for a long time. "The success of our research program was effectively backed by the modern computing facilities provided by Compute Canada," said Artem Cherkasov. "Using computer simulations, we sometimes go through 50 million compounds to find a molecule that will seat in a precise and accurate way."

The breakthrough is now being licensed to the pharmaceutical company Roche. Under the terms of the agreement with Roche, UBC and VCHRI can expect to receive an upfront payment, and up to $141.7 million US in pre-clinical, clinical and sales milestone payments for the first product to reach the market, and royalties thereafter.

Mechanical Engineering, McGill & Polytechnique Montreal: Dr. Siva Nadarajah & Dr. Eric Laurendeau

The long-term objective of Dr. Siva Nadarajah is to facilitate the design of the next generation of commercial aircraft to meet future global aerodynamic performance, and accurately resolve the far-field airframe acoustic signatures within a single numerical framework, would be an invaluable computational tool. Compute / Calcul Canada infrastructure is used to perform three-dimensional Unsteady Reynolds-Averaged Navier-Stokes calculations around complex aircraft geometries for problems with up to 100 million unknowns. In 2016, Professors Nadarajah and Laurendeau were awarded a prestigious NSERC CRD grant worth $650K, to work closely with Bombardier Aerospace (BA), a prominent Canadian aerospace company and Cray Inc., a manufacturer of specialized supercomputers to develop the next generation of numerical tools for massively parallel high-fidelity computational fluid dynamics (CFD). This partnership will contribute and advance the academic field of CFD for massive heterogeneous computing. The research will provide the Canadian Aerospace and Computing industries with valuable technical achievements, new opportunities that would strengthen the Canadian Aerospace economy, and HQP who would be valuable assets to the industry. The investment into the proposed project demonstrates BA's strong interest to remain competitive through research and development initiatives. The benefits for Canada's aerospace industry, number three in the world behind the USA and Europe, will be significant.

As CC usage and the CC user base continue to grow, so will the associated benefits to Canada. The creation of a national data infrastructure, expansion of cloud infrastructure, construction of hybrid systems and general increase in capacity described in Part B, will enable benefits to be realized across a wide spectrum of research areas. While benefits will continue to be realized from existing infrastructure and services through 2018, this is expected to accelerate with the commissioning of these new systems, which begins in late 2016.

155 University Ave, Suite 302 Toronto, Ontario, Canada M5H 3B7
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

1.5 Producing a Highly Qualified Workforce

CC has provided internationally competitive research and training opportunities to student and technical staff at Canadian universities, opening up new and long-term career paths in their fields in Canada. There is a continual influx of new users, usually at the student and post-doctoral level. These users undertake training either through CC or through their graduate courses and institutional opportunities. Most are developing expertise in their chosen disciplines with a computational or data focus, through their graduate programs and research. As they work through the system, many become expert computational scientists.

As shown in Section 1.2, CC currently supports more than 3700 graduate students, 847 undergraduate students and more than 1000 postdoctoral fellows. Since 2013, 1889 CC users have transitioned from academia to industry, and 588 have transitioned to the public sector. Among CC users reporting transitions within the academic sector, 738 undergraduates have become graduate students, 525 masters students have become PhD students, 719 PhD students have become postdoctoral fellows and 321 postdoctoral fellows have become faculty members.

CC training is essential for success in many research disciplines. While advanced research computing may be an essential tool for research in the field, it is often left up to the student or the researcher to acquire the skills necessary to use this tool, and to use it effectively. CC's one-on-one support, combined with a diverse array of introductory and advanced topics, is often critical for the personal success of the individuals CC serves. These skills are useful not just in research, but also in the broader workforce, and CC regularly encounters companies and organizations wishing to use its resources, because they already employ highly qualified people previously trained by CC and familiar with our resources.

Compute Canada experts provide more than 100,000 staff hours annually, in direct support of researchers. In addition, between 2012-2014, CC has documented delivering more than 54,000 hours of training to more 11,000 researchers at approximately 573 training events.

New courses, developed by CC during the previous MSI, include: Accelerators (GPU, Phi, CUDA), Applications (Gaussian, Gromacs, FEniCS, Matlab, etc.), Big data (Map-reduce, Hadoop), Development (programming, databases, revision control, parallel debugging, optimization, profiling), Globus (high-speed data transfer), Parallel Programming (general, MPI, OpenMP, etc.), Scheduling (MOAB, Torque, optimal job submission strategies, etc.) and Visualization (Paraview).

With MSI funding, CC expanded its training and partnership initiatives to participate and lead international training efforts. Canadian trainers and students now participate yearly at the International HPC Summer School, which CC joined for the first time in 2013. In recognition of CC's contributions to the school, CC was invited to host this prestigious program in Canada for the first time in 2015. The event is sponsored by CC, the Extreme Science and Engineering Discovery Environment (XSEDE), the Partnership for Advanced Computing in Europe (PRACE) and the RIKEN Advanced Institute for Computational Science (RIKEN AICS) in Japan.

In 2015, CC became a Software Carpentry national partner. Software Carpentry runs workshops all over the world, training researchers in the basic software skills needed to effectively and efficiently use ARC resources. CC leverages open access instructor training materials to bring a high level of standardized introductory training to new researchers across the country. As a partner, CC has committed to organize at least 20 software carpentry workshops across Canada every year, and has effectively become the coordinating body for all Software Carpentry workshops in Canada. This kind of partnership has been made possible by MSI funding and enabled by CC's corporate structure.

CC has sponsored the Digital Humanities Summer Institute since 2014. This is one of the leading training events in the world for Humanities and Social Sciences researchers, with around 800 students participating in 40 courses at the University of Victoria, over a 2-week period each year and many others at satellite events at other campuses. In the first year of sponsorship, CC sent 5 staff to the event as students. In 2015, an intro to programming course was taught by a CC team member, and in 2016 CC will teach 3 courses: Introduction to Programming, Big Data, and Cloud.

2. International Competitiveness

2.1 The National Platform Provided by CC

Today, the CC platform includes approximately 200,000 CPU cores, several modest-scale GPU deployments, and approximately 20 PB of disk storage and significant tape resources. Access to these resources is supported by more than 200 people, most of whom are employed by research institutions spread across the country.

The CFI Cyberinfrastructure Initiative (Challenge 2 - stage 1) award, announced in July 2015, is being used to address pressing and urgent needs for a technological update of the hardware platform. Stage 1 includes investment in four shared national systems, deployment of more than 120,000 computational cores, substantial investment in accelerator capacity, and creation of a national data infrastructure. It also includes investment in common middleware solutions to address such topics as research data management, authentication and ID, data movement, monitoring and resource publication. The stage 1 technology deployment plan status is provided in Part B of this proposal. Part B also provides a detailed account of further proposed upgrades to the facility over the next 3 years using the stage 2 award.

The more than 200 experts who are part of the MSI project are critical to operation of the platform and support of the researchers who use it. Approximately half of this team have graduate degrees and so have a research background. The majority of team members have degrees in computer science, engineering or physics. There is a growing number of team members with backgrounds in life sciences, as well as targeted new hires in humanities. More details are provided in Section 3.2.

The hardware for the facility is owned by host institutions, not CC. Institutions also directly employ most CC team members funded through this proposal. The institutions therefore play a critical role in providing all facility services and local support of those services.

CC's national platform allows Canadian researchers to collaborate with researchers around the world. CC users have reported collaborating with researchers in more than 100 countries. Of the approximately 3,000 documented collaborations, the four highest frequency involve researchers in the US (848), France (292), Germany (187), and the UK (144).

An important case in point of how CC is enabling international collaboration and competitiveness is the Genetics and Genomics Analysis Platform (GenAP) project spearheaded by professor Guillaume Bourque. GenAP facilitates the distribution and analysis of high-throughput data for the Canadian life sciences research community by leveraging CANARIE high-speed network and CC HPC resources. Since its launch in February 2015, GenAP has had a steady growth with more than 120 users, 147 projects and 110 applications created such as Galaxy, a tool facilitating data analysis for non-bioinformaticians. The GenAP-Galaxy applications alone have launched over 10,000 jobs on CC resources. The platform also includes a framework for advanced bioinformatics analysis pipelines that have been used to run over 11,000 jobs. Finally, GenAP also includes web tools like the UCSC Genome Browser and the International Human Epigenome Consortium (IHEC) Portal with an average of 2000 and 500 page views per week respectively. This hub is used by the international community to share globally-generated reference epigenomics datasets. The analytical tools developed for GenAP are examples of how CC can play a crucial role in facilitating HPC-related technology creation for advanced medical research. Besides their importance in medical research, tools like GenAP are expected to be transferable to other disciplines, such as advanced materials manufacturing, where large volumes of emerging materials property data are being mined to optimize the design of functional materials.

2.2 Comparison to International Facilities

CC is the primary provider of ARC services to the Canadian research community and the only national organization dedicated to this purpose. CC's goals are different from many comparable international facilities because we provide ARC resources to Canadian researchers, in any discipline, at all scales beyond desktop computing. By contrast, most international facilities focus on one or two tiers of the "Branscomb Pyramid," a concept that illustrates the distribution of computing resources from desktop (at the bottom) to leadership class (at the top). Canada is in a position to deliver an integrated system, providing a clear advantage as researchers move from solving problems with desktop-scale solutions to problems with supercomputing-scale solutions. CC provides resource sharing to maximize resource use, and subsequent science impact. This science impact is the truest measure of competitiveness. CC strives to ensure we keep pace with the needs of science. The chart below shows the science impact (measured by FWCI - see section 1) for CC-enabled publications compared to the average in several comparator countries around the world for a wide range of disciplines. These results show that CC-enabled publications do very well in comparison to papers from other countries.

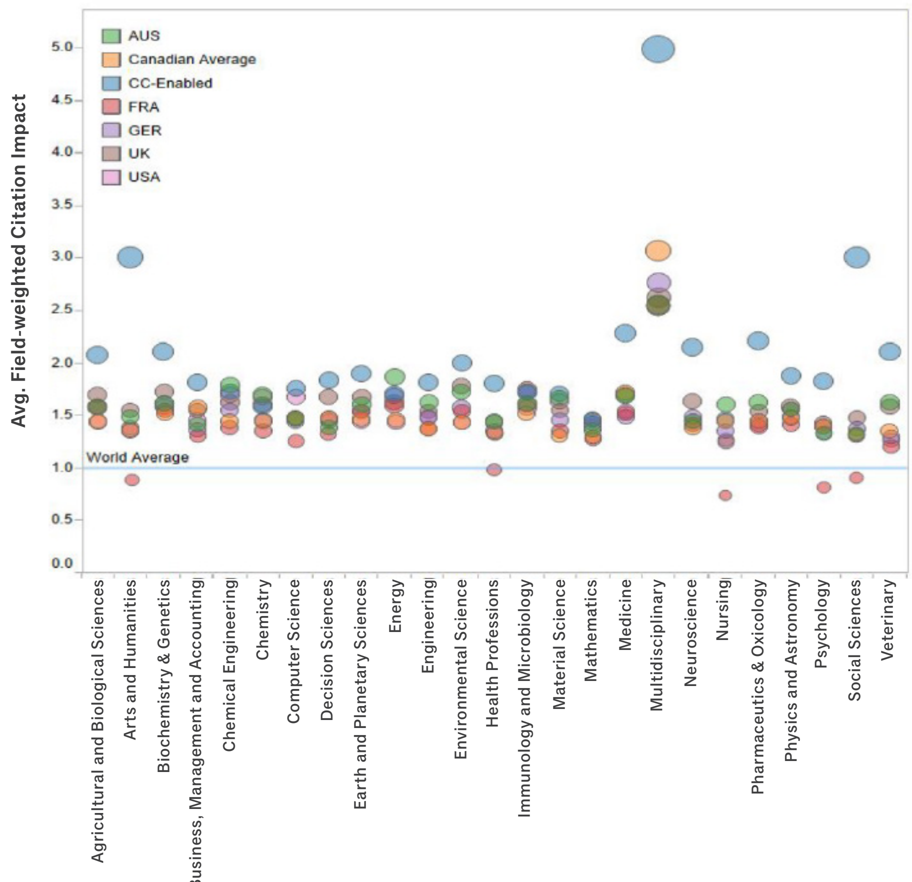

Figure 2.1: FWCI comparison between CC-enabled publications and the average from comparator countries over a wide range of disciplines.

Outside Canada, there are a few example of ARC organizations which have a similar role. XSEDE (www.xsede.org) serves a broad range of US researchers with facilities located at multiple service provider's centres across the country. PRACE (www.prace-ri.eu) and EGI (www.egi.eu) serve a broad range of European researchers, with facilities located across Europe. PRACE's facilities are traditional supercomputers, while EGI's facilities focus on high-throughput or capacity computing. CC provides both of these capabilities. However, in contrast to Canada, researchers in the US and Europe can also access significant other domestic facilities, so CC's universal service model appears to be unique.

The competitiveness of the infrastructure provided by CC to the research community has slipped in recent years. A 2013 assessment of available computational power per academic researcher ranked Canada 15th among nations. There has been no substantial deployment of ARC resources in Canada since 2013, while other countries have brought major new systems online. The capital investment enabled by CFI Cyberinfrastructure Challenge 2, Stage 1+2 awards addresses the existing technical debt.

2.3 The CC Platform is Built to Serve Canadian Researchers

CC has consulted extensively with the community over the last 2 years, informing all aspects of planning both capital and operational investments. The growing user base (2.5x over the current 5-year MSI period), community survey responses, and submitted community white papers all support the strong need in the research community for the specialized services that CC provides. (See Sections 3.1 and 3.2.)

2.4 Data Sharing and Access in Compute Canada

CC's practice is to provide the user with the tools required to share data according to their needs. For some researchers, this means providing physical and network security layers which provide protection against inadvertent sharing of private data. For others, it means providing a platform which gives researchers enough control to allow sharing only for collaborators within their university, only among their collaborators, or with the world.

In 2014, CC deployed the Globus File Transfer Service (FTS) across its entire network (24 storage sites). This platform allows researchers to move data in and out of CC storage through a web interface or through code written to leverage the Globus FTS API. Researchers have full control over data access and can choose to share their datasets with any collaborator anywhere in the world. Since the launch of the CC Globus FTS, researchers have used it to move more than 2 PB of data.

In 2014, CC also introduced its Research Platforms and Portals (RPP) competition for the first time. This allowed researchers or groups of researchers to apply for support of portals or platforms which serve shared datasets to a research community. Data sharing platforms for many fields including astronomy, neuroscience, genomics, and digital humanities are currently supported on CC resources as a result. In 2016, 30 RPP awards were made to projects in a wide variety of disciplines.

The first major purchase using the Cyberinfrastructure Challenge 2, Stage 1 award is the new national data infrastructure (described in Part B). This infrastructure meets needs of diverse communities for data replication, availability (through a common namespace and integration of distributed storage systems) and durability (through policy-based backup and capacity management). In addition, this infrastructure will establish a nationwide service layer that will support both metadata and persistent identifiers, key requirements of more ambitious efforts to manage and exploit growing volumes of research data.

CC and the Canadian Association of Research Libraries (CARL) have entered into an agreement to create a production-scale system for research data management in Canada. The solution will leverage existing platforms, and custom software to provide a solution for data ingestion, curation, discovery, transfer and geo-replication. The system is designed to be federated (bring your own storage) and scalable to a national level. Working in collaboration with Globus, and the CARL Portage project, a team of CC developers have been working on this solution since January 2016 and will provide first demos at CC's annual HPCS conference in June 2016.

3. Need for CFI Funding

3.1 Assessment of User Demand

CC has undertaken extensive consultations in order to engage the community and assess user demand. This consultation has included:

- A needs survey in the autumn of 2013 (more than 200 faculty responses)
- More than 20 in-person consultations at various Canadian campuses in the winter of 2013-14.
- A call for white papers issued in summer 2014.
- In-person consultations at 6 locations across Canada in January, 2015.
- A needs survey in January 2016 (more than 100 faculty responses).
- A call for white papers issued in January 2016.
- Interviews with some of Canada's largest HPC users in 2016.
- In-person consultations at 15 locations across Canada in January-February 2016.
- Several online consultation sessions in 2014 and 2016.

In addition, user data from the Compute Canada Database (CCDB) was mined for the 2010-2015 period to search for usage trends.

3.1.1 Growth in the User Community and Resource Requests

The number of CC active faculty users has grown from approximately 1,250 at the time of submission of the last CC operations grant (2011) to nearly 3,200 as of January 1, 2016. This 2.5x growth in the user base over 5 years is driven by the essential nature of ARC in a broadening list of disciplines. From January 1, 2015 to January 1, 2016, CC added more than 260 new faculty members in medical, biological and life sciences, representing 27% year-over-year growth. For the first time, there are now more CC faculty accounts in medical, biological and life sciences than in all of physics and chemistry combined. There is no evidence of saturation in the growth of the CC user base. CC must therefore plan for both continued growth and continued diversification of the CC user base over the coming years.

Faculty growth (unique users)
(until Jan 1, 2016)

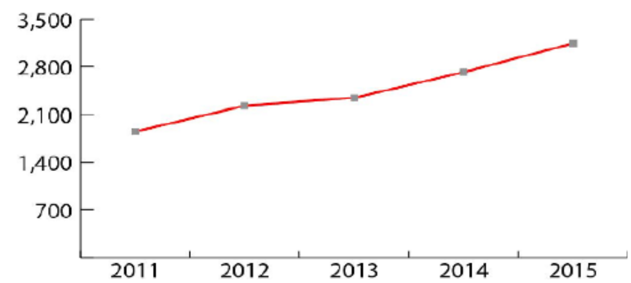

Growth in usage is strongly constrained by available resources. The allocated CPU was 85% of the request in 2012 and only 54% in 2016. In each case, these are the numbers after technical validation and review and so do not include exaggerated or inaccurate requests. It should be noted that, unlike in other jurisdictions, virtually all applicants are federal grant-funded researchers who have already been subject to a merit-based peer-review process. They generally have already had their program funded and have capital and personnel funding invested in the project. In the current Canadian funding system, they essentially have no alternative but to meet their computational needs through CC resources (i.e. there is no alternative grant agency program for local or regional computational infrastructure). In this context, a 54% allocation rate is very low, leading to stranded capital and operating investments in those funded research programmes, significant delays in student graduation, and reduced international competitiveness. The rate of GPU allocation was essentially 100% in 2012. As community adoption of GPUs has grown, CC capacity has been unable to keep pace, leading to a 20% allocation rate in 2016. The situation for storage is better due to deploying stage 1 storage in mid-2016. Without the new stage 1 storage, the 93% allocation rate for 2016 would have been 65%.

155 University Ave, Suite 302 Toronto, Ontario, Canada M5H 3B7
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

3.1.2 User Needs Survey (2016)

The most recent user needs survey was distributed to the Compute Canada user community early in 2016. The 189 submissions included responses from 109 faculty members from a wide variety of disciplines. Natural Sciences and Engineering dominated the responses, with 76% of respondents identifying NSERC as a source of funding, 19% identifying CIHR (health research) and 7% identifying SSHRC (social sciences and humanities). Each respondent was asked to identify their current usage of both compute and storage and to project the relative increase in that number over the coming 5 years.

The demand projections over 5 years are

- An expected 5-year compute demand growth of 12x
- An expected 5-year storage demand growth of 19x

A number of requests for specific resources were also identified from the survey data. These are described in Part B of this proposal (i.e. Cyberinfrastructure Stage-2).

3.1.3 White Papers (2016)

The Canadian ARC community submitted 17 white papers to CC in early 2016. Representative bodies that submitted papers include: The Institute of Particle Physics plus Canadian Institute of Nuclear Physics, the Canadian Astronomical Society, the Canadian Association of Theoretical Chemists, and the Canadian Society for Digital Humanities. Several other multi-institutional groups of researchers submitted papers related to genomics (4 largest Canadian centres), neuroinformatics (CBRAIN platform), ocean modelling (25 faculty from 14 institutions), materials research (28 faculty from 12 institutions) and light microscopy (3 researchers from 3 institutions). In addition, several individual research institutions or research projects submitted papers.

Projected growth in compute over the next 5 years vary by discipline from 3x in subatomic physics to 12x in theoretical chemistry. Weighting by current usage in each discipline, this leads to an average expected increase of 7x over 5 years. It should be noted that, in some cases, the range of responses within a discipline may include researchers who need 100x over the next 5 years. In these cases, the community is already providing an average which may not accommodate the peak need in their discipline. Further, the age and size of current CC systems has inhibited the growth of the high-end compute community in Canada and leads to a downward bias in the responses. The 7x projection could therefore be an underestimate of true demand.

Projected growth in storage needs over the next 5 years also varies by discipline, from 3x in subatomic physics to 30x in astronomy. Very strong growth rates are seen in genomics (15x) and neuroscience (18x). Based on these submissions, we estimate an aggregate growth of 15x over 5 years. It is noteworthy that while growth in compute consumption is expected to be dominated by the needs of the existing CC user base, the growth in storage need is directly tied to the emergence of new data-intensive disciplines, new instruments and new CC users. The 15x projection could therefore be an underestimate of true demand.

Full supply and demand projections are provided in Part B, along with white paper comments which directly inform infrastructure planning.

155 University Ave, Suite 302 Toronto, Ontario, Canada M5H 3B7
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

3.2 Use of MSI Funding to Enable Exploitation of the Facility

MSI funding makes it possible to operate the CC facility through provision of funds for utilities and services and through salary support for the cadre of 200 expert team members across the country who ensure that the facility's technical capabilities are fully exploited. To date, CC has achieved its mandate to enable excellent research by using MSI funding to invest in people to support CC's users and in enhanced services for those researchers.

Skilled and experienced people are a critical component of CC's competitiveness:

- Technical staff maximize uptime and availability of equipment.
- System administrators ensure that scarce resources are efficiently scheduled and maximize the productivity of research users.
- Research support specialists make sure that users are employing CC resources productively and efficiently.

CC's original MSI submission to the CFI proposed to increase research support personnel to better meet growing research needs, and to increase technical and system administration staff to support the expanding complexity of our distributed portfolio of ARC systems. This has been accomplished, with an overall growth in staffing from 137 FTEs in 2012-13 to 212 FTEs today. The current CC team complement is expected to be relatively stable throughout the next MSI period, with 226 FTEs projected for 2022. Roles within this complement will shift in time as site consolidation proceeds and researcher-facing service offerings expand.

CC team members are highly skilled, with specialized training, experience and qualifications. Many have advanced academic degrees (more than 70 have PhDs). In recent years, staff with specialized skills in biological and medical sciences, data management, digital humanities and visualization have been added in response to shifting community needs. The quality of this team is competitive with the quality of staff at any comparable facility internationally.

"In the world of large-scale computing, the speed with which barriers are overcome and problems are solved is absolutely critical. It is often the case that our research team hits a computational barrier that we cannot overcome. These barriers become stopping points for us: we cannot move any further with the project and we sit waiting for a response. I have dealt with support staff at numerous institutions in the past, including the Max Planck Institute in Germany, Cornell University and the Stanford School of Medicine. It was routine at these institutions to wait for days for a response from the staff who manage the computational infrastructure. Ross Dickson, however, assured me during my first week of my faculty position at Dalhousie University that he is always only a Skype call away and will always tend to solve problems as soon as he possibly can. I am convinced that our research lab remains competitive in computational and statistical genomics because of the dedication Ross has to timely answers to our queries. We consider him a critical part of our research team."

Sean Myles
Dalhousie University

Without this highly qualified staff, the facility would not run, the HQP training described in section 1.3 would not occur, and researchers would not get the support they need to efficiently exploit ARC resources. CC staff represent a critical asset of the organization, the development of which is a strategic objective. Retaining and increasing the advanced skills, detailed knowledge and long experience of our personnel are essential to CC's success and growth in the future.

CC provides a structure through which local experts from campuses across the country can share knowledge and work to solve common problems. Several national technical teams have been formed to work on cross-cutting topics such as storage, networking, monitoring, help desk, national platforms and cloud operations. Recently, CC has also formed several national disciplinary support teams in digital humanities, biomolecular simulation, subatomic physics and bioinformatics. These disciplinary teams are involved in community outreach, documentation and creation of common technical solutions for various common disciplinary use-cases.

MSI funding is also being used to deliver software and middleware services to the research community to optimize their use of the facility. As an example, the Globus File Transfer Service is enabled at 24 different CC storage locations. This allows users to move files to and from CC systems in a straightforward and efficient way. Stage 1 capital funding is currently being used to build several national infrastructure services in support of research data management, and platforms such as those funded through CFI Cyberinfrastructure Challenge 1.

3.3 Securing and Diversifying Funding Sources

CC is actively working to develop new sources and approaches to its funding. Compute Canada has worked directly with key government departments to advocate for coordinating funding for advanced research computing to leverage investments for the benefit of the greater ARC community. Through extensive and targeted government relations Compute Canada has demonstrated the value of pooling regional development funding and government research ARC funding to maximize capacity. CC's "Event on the Hill" in December 2014 highlighted to Members of Parliament CC's key role in Canada's digital research infrastructure; a similar event is planned for later in 2016. These efforts have been successful: CC was named as a key player in the government's Science Technology and Innovation Strategy in December 2014, $100M was earmarked for Digital Research Infrastructure in April 2015, and CC was encouraged to submit a proposal for a new funding model to Finance Canada for consideration for Budget 2016. Consultations continue as Canada develops a national innovation strategy in time for Budget 2017.

These efforts at the federal level are mirrored by similar efforts provincially, typically led by the respective regional organization. The regional organizations have been actively communicating with provincial governments across Canada about both the important role of ARC in the province and the value of the partnership between CC and the regional organization.

Several institutions and regions have provided services on CC systems to the private sector on a fee-for-service basis. This has the potential to be a source of significant revenue, and the CC organization is examining policies around such fee-for-service arrangements, however this is currently limited by the lack of capacity to deploy to the industrial research sector and is currently focused primarily on consultations. The regional partners, along with Compute Canada, receive regular requests for services. As capacity grows this will be an important source of external non-governmental funds.

3.4 Impact of Failure to Fund

CC supports more than 10,000 researchers across the country. These researchers have reported more than 33,000 publications enabled by use of CC resources since 2010. Failing to fund CC would have a serious negative impact on Canadian research productivity across a wide range of disciplines.

The recent $30M Stage 1 award ($75M including match) is currently being used to purchase major new infrastructure at four national sites. The operations and maintenance plans for these facilities rely on MSI funding for basic operations and maintenance charges and for funding the personnel required to provide services to researchers. Failing to fund this MSI project would require already-selected Stage 1 infrastructure hosts to place procurement on hold until any alternative funding for operations could be secured. Similarly the Stage 2 infrastructure investment also being proposed by CC relies on this MSI funding proposal for the sustainability of operations.

155 University Ave, Suite 302 Toronto, Ontario, Canada M5H 3B7
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

CC operations have been funded through the MSI program since 2012. The staffing levels set by the first MSI award are being continued in this second MSI funding proposal. Since there is no flexibility in expenses like power costs, any cut to the operations budget would lead to staff reductions and reduction in services to researchers. While the CC user base is projected to grow by approximately 60% by 2022, this proposal only proposes an increase in staffing levels by 6-7% over the same period. A reduction in staffing levels from the proposed levels would therefore be felt immediately by the research community.

4. Excellence in Governance

CC's governance model is unique, combining a federation of regional consortia with a corporate national organization reporting to an independent board of directors. CC is the realization of an alliance of institutions wishing to provide a national advanced research infrastructure and maximize the sharing of expertise across Canada. While no one form of governance appears to represent best practice among other ARC facilities around the world, CC seeks to follow best governance practices from the business and not-for-profit sectors.

The Board's responsibility for setting policy, overseeing management and planning strategy, is separate from management's role, with executive authority delegated to the CEO, who has assembled and leads a team of competent functional experts. The CEO sits on the Management Committee with the heads of the four partner Regional Organizations; the Management Committee works to align the Regional Organizations and CC on major initiatives.

4.1 Governance Structure of the Facility

CC is a national not-for-profit corporation with an independent board of directors. The members of the corporation (equivalent to shareholders in a for-profit structure) are Canadian research institutions, either universities or research hospitals. The board of directors is accountable to the members, who elect the board. The board is comprised of directors who, together, hold expertise in a core areas including: corporate governance, finance, risk, research, and advanced research computing, and who are bound to act in the best interest of the members. This structure reflects international best practice in governance.

The CC board includes standing committees on Audit and Disclosure, Management Resources and Compensation, and Governance and Nominations. The CC board is advised by both an International Advisory Committee (IAC), comprised of experts from ARC organizations from around the world, and by an Advisory Council on Research (ACOR), comprised of Canadian experts representing a wide range of disciplinary expertise.

4.2 Accountability to Members and the Administrative Institution

The Member Council augments the role of the CC Board, enabling decision-making where member institutions cannot delegate their interests to the CC Board, primarily in connection with facility assets (owned by individual institutions) and facility personnel (employed by individual institutions). For example, CC's cybersecurity policies are being developed to reflect best practice, but formal adoption of each policy is strictly the responsibility of each member institution. Therefore, extensive consultation is required to create uniform policies that can be adopted, and then those policies will eventually be adopted by each institution using its own governance procedure; each institution may require one or more exemptions to the policy to reflect specific constraints. Regular consultation with members and with the Member Council is essential to maintain CC's accountability to members, whether they represent the owners of facility assets, the employers of facility personnel or the research community at each institution.

It is challenging to convene the representatives of roughly 35 institutions, much less achieve consensus with such a group, so CC also works with the Management Committee on plans, policies and procedures with the objective of developing consensus either on decisions that may be made by CC or on materials that still require consideration and approval by CC's members (particularly as asset owners and/or employers). Working closely with members, CC is committed to continual improvement in governance.

155 University Ave, Suite 302 Toronto, Ontario, Canada M5H 3B7
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

CC is accountable to its administrative institution, Western University, not only as a member but also in connection with the legal obligations placed on Western by CFI Award Agreements, particularly in connection with financial reporting and accountability. CC, Western and the four Regional Organizations/Leads intend to enter into an Inter-Regional Agreement that will specify the respective roles and responsibilities of each party, including the division of responsibilities for the project's financial management, principally between CC and Western.

4.3 Performance Measures

Compute Canada, in partnership with its regional partners, institutions, and experts has developed an initial performance measurement framework with over 260 metrics capturing important operational activities across the organization, mapping them to CC's strategic goals and objectives. Work is underway to identify the most significant metrics that will be selected for implementation. This work is also informed by continued evolution of our Management Plan -- making sure that there are metrics in place for all key management activities. For example, even where operational responsibility for an activity (such as operation of a specific system) rests with a region or institution, effective management of the facility requires that relevant performance metrics are captured to allow oversight of this type of activity across the organization, and to ensure consistent high performance.

CFI has mandated that CC provide a performance measurement framework by September 30, 2016. Significant progress has been made in this area, but the definitive framework is not available at this time. Initially-identified KPI groupings focus on the four goals from the 2014 Strategic Plan:

1. Provide services which enable excellent Canadian research across a broad range of disciplines.
2. Build capacity and connections nationally and internationally to tackle future research needs.
3. Maximize the positive effects of Canadian research on the lives of Canadians.
4. Ensure sustained organizational excellence.

KPIs will roll up to scorecards for each strategic goal, as well as a scorecard view (on the same KPI metrics) reflecting progress in areas such as Internal, Finance, Learning and Customers. Initial KPI groupings under consideration span a wide range, such as Employee Workforce Diversity, Issues & Risks, Research Outcomes, Training, Innovation, Cost, and Technology & Innovation. CC is working with Nick Berente of Georgia Tech, who helped XSEDE with their KPI process. CC will coordinate with CFI to ensure KPI metrics, groupings, and scorecards are appropriately reflective of MSI interests.

4.4 Consultation Process Which Informed the Strategic Plan

In December 2013 and January 2014, CC conducted nationwide consultation that included 24 townhall sessions at institutions across the country, two online consultations and consultations with CC team members. The plan was developed by CC management and key regional representatives. Drafts of the full plan were shared with the CC membership, all functional and advisory committees, and the full CC user community, with all inputs reflected in the plan before it was adopted by the CC Board in July 2014.

4.5 The Added Value of the Networked Model

The effectiveness of CC builds on the effectiveness of the institutions and regional organizations that make up the CC federation. The institutions and regional consortia that predate the regional organizations have successfully served the needs of their local research communities for many years, and bring expertise and unique experience through their personnel, trust, effective inter-institutional working relationships, and important provincial funding relationships through their organization. CC's federated model parallels Canada's federal structure politically and represents a natural and effective way to integrate distributed operations across a range of operating environments and circumstances. The federal model encourages co-funding from provincial sources that must be confident that their own funding objectives and targets will be met.

155 University Ave, Suite 302 Toronto, Ontario, Canada M5H 3B7
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

5. Excellence in Management and Operations

The CC facility is operated by a team that, while distributed and federated, has for many years already operated excellent local facilities with high efficiency and optimal service to their local communities. An important objective of creating CC has been to build on that excellence and experience, maintaining high levels of service even as many legacy systems are defunded and replaced by fewer, larger systems serving researchers across Canada, and to enable increased excellence through the sharing and adoption of best practices across the CC organization and continuous improvement in all aspects of CC's operations.

CC's management structure and plan are detailed in the separate Management Plan. CC is working to develop a new management framework that integrates national leadership with local experience and expertise, balanced in turn by the rights and obligations of the institutions that represent a critical part of the CC federation. CC encourages initiatives on the part of self-organizing internal teams, which are exploring the use of operating principles such as are found in agile development and holacracy.

Recognizing the different degrees of operational capability and accountability across the federation, CC follows several organizational patterns for ensuring effective management of different operational areas:

- National projects, such as resource allocation competitions, account renewals, user migration, specifying and selecting "national" scale technologies, developing and revising CC budgets, consulting the community on their needs.
- National procedures, including asset management decisions (extended warranty, life support), grant consultations and budget adjustments.
- Policy-based operations, such as technology operations and researcher support. Clearly these are major areas of activity for CC, as well as the primary areas in which the CC federation relies on its partners for success. Policies focus on adoption of best practices, interoperability, performance measurement and continuous improvement.
- Principle-based frameworks to coordinate and harmonize institutionally-based policies in order to achieve clear objectives across the CC federation. Examples include cybersecurity, human resources, risk management and financial control.

Regardless of the organizational pattern, all activities involve operational work or initial policy development by inclusive teams of experts rely on consensus across the organization, and sometimes involve consultation with/agreement by responsible partner institutions.

5.1 Description of the Current Management Structure

CC's management structure and practices marry international best practices, both for comparable research infrastructure and for effective organizations of any kind, with mechanisms that recognize both the federated nature of CC as well as the limits of CC's authority in areas where institutions have continued fiduciary responsibility:

- CC is managed by a capable team of executives, with clearly defined areas of responsibility. Each senior manager has been recruited to bring specific functional expertise to the organization.
- Individual executives have responsibility and accountability for their functional domains but must fulfill their objectives by working collaboratively across the CC organization, particularly through inclusive teams, both standing and ad hoc, comprised of experts from all the regions of the organization, as well as institutions and regions with responsibility for various operational functions (e.g. system operations, user support).

155 University Ave, Suite 302 Toronto, Ontario, Canada M5H 3B7
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

- The Management Committee works to ensure alignment of objectives across the organization, with specific responsibility for developing overall budgets, providing resources for all required activities, and ensuring the effective representation of regions on key teams. Input from this group is also sought on any policies or frameworks in advance of partner consultations or on policies that impact overall budgets.
- Key teams include the Technological Leadership Council (TLC), Science Leadership Council (SLC) and Security Council (SC), supported in turn by a variety of specialized teams (Storage, Research Support, Visualization, Finance, etc.). For example, the Research Support National Team (RSNT) was charged with planning national support activities, determined that selection of a national ticketing system was required, reached consensus on the best system to select, and is now rolling out implementation of that system across functional areas. The TLC supervised this effort and approved the recommendation to select the system.

Detailed organizational charts illustrating these relationships can be found in the Management Plan.

5.2 Team Member Expertise

CC works with over 200 experts, including 60 with PhDs, at 35 institutions. This team provides support to the broad community in Canada which includes 70 institutions. This is complemented by regional staff and the senior management team at Compute Canada.

CC Team Members represent a critical asset of the organization, the development of which is a strategic objective. Retaining and increasing the advanced skills, detailed knowledge and long experience of our personnel is essential to CC's success and growth in the future.

5.3 Tracking Outcomes Enabled by the Facility

The process by which CC tracks and measures outcomes has undergone significant change in the last year. CC continues to improve both the resource allocation and account renewal processes implemented by its CCDB management application, which now allows CC users to submit their Canadian Common CV (CCV) data to CC as part of these processes. Access to this data has resulted in a rich set of machine-readable information being submitted by more than 2300 faculty members during the annual account renewal process in early 2016. Detailed reporting is now available, including: publications enabled by CC support, industry partnerships, federal funding received, and HQP training performed. This resulted in the identification of nearly 10 times more CC-enabled papers relative to previous reporting efforts. CC is continuing to refine the analysis and use of the collected information.

Some of the outcomes enabled by the CC facility which have been collected through CCDB have already been analyzed and have been presented in previous sections including:

- Number of CC-enabled publications
- Impact of CC-enabled publications across a broad range of disciplines
- Positive career transitions among users
- Industrial contributions by CC users, including patents, start-ups and technology transfers

Other metrics which have been traditionally collected are also being used, including the number of active users and their distribution among disciplines and parts of the country.

In response to a 2015 CFI expert panel report, the organization is in the process of developing a comprehensive set of organization-wide KPIs, which will be delivered to CFI by September 30, 2016. The examples included above will be important measurable components of those final KPIs.

5.4 Key Risks for the Facility

The following are significant risks for the facility and threaten the facility's capacity to deliver ARC services required and needed by the Canadian research community to be competitive.

5.4.1 Lack of sustainable funding

The level of investment in ARC required to keep Canadian researchers competitive internationally is growing and represents a significant burden for hosting institutions. Reducing the number of hosting sites increases this financial pressure, and makes it difficult to establish a scalable and sustainable operating and financial model for ARC in Canada.

Mitigation: In the absence of long term predictable funding from federal and provincial governments, CC has a limited control over this risk. Nonetheless, CC is advocating to funding agencies, stakeholders and government on the importance for Canada of having a sustainable ARC platform. CC is also engaged with its members on a discussion about having a more sustainable and equitable model for the facility.

5.4.2 Insufficient supply

The required amount of resources needed to support the growing user community of the facility is larger every year. An insufficient increase in resource available to the facility users would impact their competitiveness and in some cases strongly affect their ability to continue their research program.

Mitigation: CC has little control over demand, and only impacts supply within the budget set by CFI. CC aims at having the highest efficiency and availability to ensure the resources available are maximized. CC allocates resources based on science review and merit in the limit of resource availability. Insufficient supply will lead to significant increase in institutional request to CFI for additional funding, adding to the complexity of the ARC picture in Canada and resulting in more disparity and lower value for the Canadian research community. CC is actively and continuously consulting with the community to ensure resources needed are well understood and taken into account in planning and providing services within available funding.

5.4.3 Insufficient Institutional engagement

Engagement of stakeholders is vital to the effectiveness of the facility operation. A lack of engagement would have several negative impacts, ranging from lack of understanding of the institutional needs, insufficient institutional financial contribution, increased risk with respect to cyber security and could also affect facility policies enforcement at the institutional level.

Mitigation: CC is consulting and engaging its member institutions on key issue to ensure alignment with their need. Compute Canada is creating more formalism around interactions with institutions to ensure clarity on roles and responsibilities, obligations and to ensure proper applicability and efficiency of Compute Canada policies at the institutional level. Compute Canada is adopting a more consultative approach with institutions to involve them in strategic evolution of the facility.

5.4.4 Accelerated decommissioning

Several services or platforms are scheduled for decommissioning in the coming 2 to 3 years. Many of these systems are already over their normal lifetime which might lead to sudden, early system decommissioning. An early retirement would lead to an abrupt disruption of resource that are critical to many research groups, and strongly affecting their ability to continue their research program.

Mitigation: Rigorous planning and management of deployment and decommissioning of systems will ensure a smoother transition, with lower risk of early retirement of systems. Compute Canada is exploring a diversity of funding to help stabilise the system portfolio and help to provide sufficient margin in case of unexpected delays for commissioning new systems.

155 University Ave, Suite 302 Toronto, Ontario, Canada M5H 3B7
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

5.4.5 HQP retention

The uncertainties associated with the lack of sustainable funding combined with the reduction in the number of operating sites creates an environment where many HQP are feeling insecure. Many other sectors of activities are in needs of the skills of these individuals, including similar international organizations. The knowledge and expertise lost upon departure could have a significant negative impact on the facility competitiveness.

Mitigation: CC has created National Teams to ensure all skills and talent have an opportunity to be involved in the future of the facility service delivery. Compute Canada is also engaging in new national and international projects that will stimulate our HQP and provide them an opportunity to contribute to novel service and technology deployment. Compute Canada works in partnership with the institutions to ensure continuity and stability in the context of the human resource affected to the facility.

5.5 Access and Allocation

The CC access policy allows access to academic, government and private sector researchers. Any grant-eligible faculty member at any Canadian university may obtain a CC account. The faculty member may sponsor any number of users under their account. These sponsored users may include students, research associates or external collaborators - including industry collaborators and academic collaborators outside of Canada. All usage by any sponsored user is accounted for under the allocation of the sponsoring faculty member. The vast majority of current CC users are sponsored in this way by Canadian academic faculty and access the facilities free of charge. Researchers who are not grant-eligible, for example, those from industry or government laboratories, may also create an account and sponsor others but must pay a fee to enable this access.

The number of access requests has been growing each year. The table below presents the number of active faculty accounts captured on January first each year for the last 6 years. In addition, it includes some relevant numbers for the annual allocation process, described below.

	2011	2012	2013	2014	2015	2016
Faculty accounts (beginning of year)	1,415	1,854	2,229	2,343	2,721	3,146
All accounts (beginning of year)	4,171	6,278	7,754	8,151	9,570	10,799
Number of RAC/RPP Requests	135	161	212	289	346	366
RAC Requests (1,000s of core years)	73	104	142	173	192	238
RAC Requests (storage in terabytes)	8,169	9,903	12,412	13,422	20,871	28,620

155 University Ave, Suite 302 Toronto, Ontario, Canada M5H 3B7
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada

The majority of CC compute resources (approximately 80%) on each system are allocated through a competitive peer-review process collectively referred to as CC's Resource Allocation Competitions (RAC). This committee review ensures that excellent science guides the usage of CC facilities. RAC includes two streams: "regular RAC" and "Research Platforms and Portals (RPP)". RPP was introduced for the first time in 2014 and includes many projects serving shared datasets to large collaborations or to the research community at large. The remaining 20% of available compute resources are left to the "default allocation". There are two main reasons for maintaining a significant default allocation: 1) it allows new users to build, test and debug applications on CC resources before requesting (and potentially wasting) a large allocation and 2) it allows researchers with modest needs to effectively "backfill" shared facilities, taking advantage of idle resources to perform high quality research. While the RAC allocates 80% of the resource, more than 80% of the community served by CC share the default allocation. Significant science impact is achieved through default usage, enabled by a low barrier to access to researchers with modest needs. While the number of papers and citations obtained by the largest RAC users is substantially higher than a typical default user, the impact per paper, as measured by FWCI averaged over similar sized cohorts, is roughly the same.

In the RAC process, all applications are categorized by discipline and assigned to one of seven national disciplinary expert committees for scientific review. The science review process produces a set of scores for each proposal, with a scoring system modelled on NSERC procedures. For the 2016 competition, 80 volunteer faculty members from across Canada served as reviewers for the RAC process. A national multi-disciplinary committee formed primarily of expert committee chairs then balances the "budget" between all proposals. In addition, CC staff with representation from all regions perform a detailed technical review to determine the reasonableness of each request and the optimal site on which to implement it.

The RAC model continues to evolve. The multidisciplinary committee of RAC chairs and the CC Advisory Council on Research have made a series of recommendations to improve the RAC process over the next 2 years. These include:

- The creation of 3 RAC streams with different required levels of documentation and different levels of scrutiny:
  - Default - no scientific review
  - Regular - technical review, at least 2 scientific reviews
  - Large - technical review, at least 2 scientific reviews plus at least one international review
- Implementation target: 2017.
- Adjusting the annual RAC schedule to provide more notice to researchers of pending changes in allocations. Implementation target: 2017.
- Allowing multi-year proposals to allow faculty members to plan research programs with some certainty surrounding resourcing levels. Implementation target: 2018.

These recommendations are a direct result of interactions with the community through the SPARC consultation process and through direct interaction between committee members and Canadian researchers.

155 University Ave, Suite 302 Toronto, Ontario, Canada M5H 3B7
www.computecanada.ca | www.calculcanada.ca | @ComputeCanada