compute canada

MSI Mid-Term Performance Report

For

Canada Foundation for Innovation

April 17, 2015

Compute Canada Submission to the Canada Foundation for Innovation

Date of Submission: April 17, 2015

Table of Contents

Section C: Financial Report and Proposal Justification

Appendices

A.1 | Cover Page

Date submitted: April 17, 2015 **Project No.:** 30198

Institution: Western University

Title of the Major Science Initiative (MSI): Compute / Calcul Canada

Designated contact person

Name: Mark Dietrich **Telephone:** 416-228-1234 x200

Title: President & CEO **Fax:**

Department: N/A **Email:** mjdietrich@computecanada.ca

Overall O&M budget

	2012-13 (actuals)	2013-14 (actuals)	2014-15 (forecast)	2015-16 (forecast)	2016-17 (forecast)	Total
Total eligible O&M costs	$21,153,044	$24,537,212	$34,109,036	$37,440,006	$40,441,216	**$157,680,514**
IOF or O&M from International fund	–	–	–	–	–	**–**
Funding requested from the MSI program	$7,718,536	$9,293,997	$14,907,186	$14,976,002	$16,176,486	**$63,072,207**

Keywords:
Advanced Research Computing (ARC), natural sciences, engineering, social sciences, humanities, medical, health

Research discipline/field codes:

Primary: 5000

Primary sub-discipline: 11800

Secondary:

Secondary sub-discipline:

Tertiary:

Tertiary sub-discipline:

Area of application codes:

Primary: 10

Secondary:

Signature of the applicant institution

It is agreed that the general conditions governing the partner contributions, and the use of CFI funds as outlined in the Institutional Agreement and the CFI Policy and Program Guide apply to the infrastructure project outlined in this application. These conditions are hereby accepted by the administering institution.

Name: Dan Sinai, Associate Vice President, Research, Western University

Signature: _____ **Date:** April 17, 2015

Note: Throughout this document, Compute Canada is abbreviated as "CC".

A.2 | Executive Summary

Compute Canada, in partnership with regional organizations ACENET, Calcul Québec, Compute Ontario and WestGrid, leads the acceleration of research innovation by deploying state-of-the-art advanced research computing (ARC) systems, storage and software solutions. Together we provide essential ARC services and infrastructure for Canadian researchers and their collaborators in all academic and industrial sectors. Our world-class team of more than 200 experts employed by 34 partner universities and research institutions across the country provide direct support to research teams. Compute Canada is a proud ambassador for Canadian excellence in advanced research computing nationally and internationally.

Compute Canada (CC) accelerates discovery and innovation by providing Canadian researchers, industry and international partners access to state-of-the-art advanced research computing (ARC) services to support a wide range of research needs. CC combines ARC resources, such as high performance computers, data storage, networking, software and services, with the talent and expertise of a cadre of skilled professionals working all across Canada. CC is dedicated to providing access to, and support for, internationally competitive ARC capabilities, regardless of the researcher's location in Canada.

Currently, more than 2,700 faculty-led research groups and more than 8,500 researchers and students across Canada, in a wide range of research disciplines, make use of CC's resources. Since the Canada Foundation for Innovation (CFI) awarded CC with funding under its Major Science Initiatives (MSI) in 2012, CC has more than doubled its user base. Access to CC facilities and expertise continues to attract internationally recognized researchers, including more than 300 Canada Research Chairs. CC hosts several large national research initiatives which share their data and compute resources through their own access mechanisms. In 2014, the largest of these initiatives supported data downloads from more than 250 Canadian users and more than 4000 unique international users.

More than 200 highly trained experts, working at more than 34 institutions across the country, directly support this diverse research community. These experts help researchers achieve more significant results more quickly, through a combination of direct one-on-one support, as well as extensive training in basic and advanced techniques in such topics as parallel programming, code optimization, bioinformatics, big data and the use of accelerators and visualization tools. Since 2012, CC has documented delivering more than 54,000 hours of training to more 11,000 researchers through approximately 573 events.

CC ensures that investments in this essential research infrastructure are managed efficiently and cost-effectively, while providing accountability and transparency to all of its stakeholders. CC also offers trusted guidance and expertise in this important area to decision-makers and funding bodies. CC works collaboratively with its regional partners ACENET, Calcul Quebec, Compute Ontario and WestGrid, using a federated model that provides national scope for the management of provincial and federal investments in ARC, as well as specialized regional and institutional support for local needs. Building on CC's five-year strategic plan (developed through a national consultation process and adopted in 2014), CC has established a detailed management and operations plan that will guide CC's efforts to achieve the strategic objectives agreed with stakeholders. In particular, CC has established an ongoing Sustainable Planning for Advance Research Computing (SPARC) initiative to support its continuous long-term planning process.

"If Compute Canada didn't exist, I'd have to get the computing power somewhere else. And if it were impossible to get it here in Canada or if I could only use minor-league computing resources, I'd move. It's really critical for me to have access to large-scale computations."

Dr. Josef Zwanziger, Professor of Chemistry at Dalhousie University and Tier 1 Canada Research Chair in NMR Studies of Materials

This robust structure has allowed several new major initiatives and partnerships to be led nationally and implemented across the regions. Activities include:

- Launching Globus: Enterprise quality data management services
- Implementing structures to ensure an international best-practice approach to national platform consolidation and renewal, governance and operations
- Creating a strong national management and governance structure and team
- Creating CC Cloud Service: Engaging new non-traditional users and expanding services
- Developing national data-visualization support
- Enhancing digital humanities engagement
- Becoming a recognized partner with Software Carpentry to deliver entry-level training in computation and data management
- Creating a multi-year allocation process for platforms and portals.

This document illustrates how CC has supported and accelerated research in priority areas such as the environment and agriculture, health and life sciences, natural resources and energy, information and communications technology and advanced manufacturing. CC has contributed to a number of notable benefits to Canada, enabling key research breakthroughs and augmenting industrial competitiveness. CC has trained HQP whose careers have taken them to jobs in industry, and has supported new ventures and spin-offs whose entrepreneurs have built their businesses around advanced research computing. Industrial partners span a wide range of industries including ICT, aviation, manufacturing, automotive and life sciences.

CC has maximized the ARC capabilities of Canada with the resources available to it. This has enabled researchers to undertake world-class research, much of which has contributed to significant social, economic and environmental benefits for Canada.

A.3 | Research Program and Researchers

The impact of advanced research computing in Canada, and the people who support it, is best illustrated by the innovative research programs supported by CC, and by the high calibre of those researchers. Note that, due to space constraints, only a small fraction of CC's more than 2,700 active faculty users (as of 1/1/2015) are featured in this report.

CC achieves its mission by enabling Canadian researchers to perform world-class research using Advanced Research Computing (ARC) strategies. The organization serves researchers who need ARC at all scales — from individual researchers to some of the largest research collaborations in the world. CC also serves the full spectrum of disciplines, including fields as diverse as computational criminology, chemistry, aerospace engineering, medicine and analysis of great works of fiction. This approach fosters both new users of advanced research computing and a culture of discovery and innovation, which has led to more than 3,700 peer-reviewed publications, 40 patents and 23 inventions between 2012 and 2014.

Many of the best researchers in the country rely on CC. For example, 32 per cent of NSERC Canada Research Chairs and more than 25% of Canada's highly cited researchers are CC users. As another example, consider the winners of E.W.R. Steacie Fellowships, awarded each year to five or six "highly promising university faculty members who are earning a strong international reputation for original research." These fellowships, awarded by NSERC, are generally considered the most prestigious awards for young faculty members in Canada. In 2014, five of six fellowship winners had CC accounts, and, in each of the last three years, one fellowship winner has had a major CC resource allocation to support his or her work. The awards span all fields of study in natural sciences and engineering and the prominence of the role of CC reflects the fact that computation and data-intensive research touches a broad spectrum of research areas.

To illustrate the impact of CC on the research accomplishments of Canadian faculty researchers, a selection of impact statements, submitted by the investigators themselves, are presented below, organized into seven research themes. These examples represent science of the highest quality, and reflect the critical role MSI-funded CC staff and computational and storage resources play.

Theme 1: Materials Science, Condensed Matter and Nanoscience

CC provides the primary computational support to Canadian researchers who are determining the basic physics and chemistry of condensed matter systems and those working closely with industry to explore exotic properties of novel materials. A few examples of the impacts enabled by CC resources include:

Dr. André-Marie Tremblay (Université de Sherbrooke, Physics) is designing new algorithms to solve the challenging problem of a quantum impurity in a bath of electrons [*Nature*, Scientific Reports 2(547) (2012)]. This Quantum Monte Carlo algorithm has become the workhorse of electronic structure calculations that combine traditional methods of density functional theory with modern treatments of electronic correlations.

Impact: Professor Tremblay's innovations allow the Quantum Monte Carlo algorithm to achieve unprecedented 500-fold speed-ups. The kind of electronic structure calculation that benefits from these algorithmic improvements is important to understand and predict from first-principle

properties of materials such as iron, copper, and uranium. Applications extend from magnets to high-temperature superconductors, chemistry and molecular biology.

Dr. Hong Guo (McGill University, Physics), in collaboration with Dr. Zetian Mi of Electrical and Computer Engineering at McGill University, has launched new experimental/computational research on the chemical transformation of carbon dioxide via solar-powered artificial photosynthesis. This project is related to the capture and chemical transformation of CO_2, thereby reducing its concentration in the atmosphere. Guo's group is using CC's facility to model and help understand the III-nitride nanowires (as catalysis) produced in Mi's lab for the chemical reactions. Prior to this project, a successful collaboration on the hydrogen evolution reaction was carried out [*Nature Comm*. 5:3825 (2014)].

Impact: Based on this preliminary work, their proposal to the Alberta-based Climate Change and Emissions Management Corporation (CCEMC) was selected in April 2014 as one of 24 winners for the round one of the $35-million international Grand Challenge: Innovative Carbon Uses. The winning group was selected from 344 submissions from 37 countries on six continents.

Dr. Roger Melko (Waterloo and The Perimeter Institute, Physics) studies the properties of $YBa_2Cu_3O_6$+x, one of the few materials that is superconducting at higher temperatures (eg. -179°C). Melko and collaborators found $YBa_2Cu_3O_6$+x oscillates between two quantum states in the pseudogap, one of which involves charge-density wave fluctuations. These periodic fluctuations in the distribution of the electrical charges are what destabilize the superconducting state above the critical temperature. Once the material is cooled below the critical temperature, the strength of these fluctuations falls and the superconducting state takes over.

Impact: Working with researchers at Harvard, Melko recently devised an explanation of the transition phase to superconductivity, or "pseudogap" phase, which is one of the last obstacles to developing the next generation of superconductors [*Science*, Volume: 343, Issue: 6177, pp 1336-1339]. This study was based on experimental data that came out only months before. CC facilities allowed the group to complete the critical Monte Carlo calculations in a very short time frame.

Theme 2: Chemistry, Biochemistry and Biophysics

Access to substantial ARC resources is critical to international competitiveness in chemistry, biochemistry and biophysics. Understanding the behaviour of molecules through complex simulations pushes the limits of high-performance computation with a wide range of applications including pharmaceutical, healthcare and oil sands-based resource extraction. A few selected examples include:

Dr. Régis Pomès (The University of Toronto and The Hospital for Sick Children, Biochemistry) specializes in the development and application of computer simulation techniques for the study of biological systems. He has recently been able to broaden the scope of his studies significantly to expand into investigating biological processes of increasing complexity. He is in a unique position to address challenging biological questions pertaining to the pathological effects of protein self-aggregation and to characterize rare events essential for the proper physiological function of important membrane proteins called ion channels. Together, these studies are helping to bridge the gap between microscopic and macroscopic scales of biological processes essential to life, providing meaningful insight into the molecular basis of human health and disease (eg. *Biophys. J.* 106:L29-31 (2014) selected for journal cover).

Impact: Ultimately, these advances will lead to a better understanding of human susceptibility to highly prevalent vascular diseases such as hypertension, atherosclerosis and aneurysms, and will facilitate the development of useful materials, such as artificial skin and vascular grafts, as well as new therapeutic approaches for the treatment

of numerous pathological ailments including pain, epilepsy, bacterial infections and neurodegenerative diseases such as Alzheimer's and Huntington's diseases.

Dr. Dennis Salahub (University of Calgary, Chemistry) works on accurate and efficient computation and simulation of important chemical reactions that take place in real-world environments. Enzymatic catalysis, biochemical electron-transfer reactions and nano-catalysis for upgrading the oil sands are examples. His group employs a multi-scale approach that involves quantum chemical methods for the heart of the chemical reactions and embeds these in a less-expensive framework of a classical mechanical forcefield (such quantum mechanical/molecular mechanical methods were the subject of the 2013 Nobel Prize in Chemistry).

Impact: The results of these demanding calculations are contributing to the better understanding of catalytic mechanisms and they are helping experimentalists to design new catalysts. In addition, the new methodologies have been expressed in a powerful set of computer codes known as deMon2k. The code is used by researchers in hundreds of labs around the world.

Theme 3: Bioinformatics and Medicine

Bioinformatics and medicine represent two of the fastest-growing sectors in usage of ARC resources. There are challenging requirements in these sectors around privacy of personal health information, data movement and the pace of technological change. Some examples of recent advances using CC resources are included on the next page.

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

"Every time I want to test a density functional theory, I have to run a thousand calculations. And, of course, the faster I see the results, the faster I know whether the theory has worked or not, and the faster I can work. Thanks to ACENET, a CC site, I can test a theory in an afternoon, and that would have taken me a month 15 years ago. [...] That was unimaginable to me 15 years ago, that I could test two or three density functional theories in a single day. But that's what's possible with the resources that CC has provided."

Dr. Axel Becke, Killam Chair in Computational Science at Dalhousie University, and the Harry Shirreff Professor of Chemical Research

Dr. Becke received the 2015 Gerhard Herzberg Canada Gold Medal for Science and Engineering.

extraction of 3D surfaces of the brain cortex and can compare it in normal versus diseased states across thousands of people. His lab can also create 3D brain reconstitution at extremely high resolution. 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-ever high-resolution 3D human brain reconstruction in the top journal Science. This breakthrough was named one of the top 10 breakthrough technologies by the *MIT Technology Review*. In 2014, Thomson-Reuters named Evans one of the most cited researchers in "Neuroscience and Behaviour".

Dr. Guillaume Bourque (McGill University and Genome Quebec Innovation Centre, Genetics) exploits new high-throughput technologies, and in particular next-generation sequencing (NGS), which are revolutionizing genomic sciences and enable the characterization of molecular processes of the cell with unprecedented resolution. The sheer volume of data associated with high-throughput technologies implies that one of the major challenges for these analyses is the development of an efficient framework for the management, processing and interpretation of sequencing data.

Impact: This group has been developing a Genetics and Genomics Analysis Platform (GenAP) that allows for the optimization of genomics software on CC resources. It has also deployed the IHEC Data Portal, which hosts and allows users to navigate through all epigenomic datasets that have been generated as part of the International Human Epigenome Consortium (IHEC). One example application of this work is the analysis of whole-genome datasets associated with the European Cancer Genomics of the Kidney (CAGEKID) consortium. Their study, published in *Nature Communications*, found that some of the tumours coming from a specific European population had a unique mutational profile that might be a consequence of dangerous environmental exposure.

Theme 4: Earth, Ocean and Atmospheric Sciences

CC plays a critical role in supporting research in Earth, ocean and atmospheric sciences in Canada. In 2012, NSERC launched a special program in Climate Change and Atmospheric Research (CCAR). Its mandate was to "support a limited number of large-scale research projects that are considered to be of high priority by both the Canadian academic research community and federal government departments." Of the seven research networks that were awarded funding, CC has granted resource allocations to directly support three of these networks, while the other four all have co-investigators with CC accounts.

Dr. Randall Martin (Dalhousie University, Physics and Atmospheric Science) uses CC resources in his effort to apply satellite observations and global chemical transport models to advance the understanding of atmospheric composition and its effects on air quality, climate and biogeochemical cycling.

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 Organization 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.

Dr. Paul Myers (University of Alberta, Earth and Atmospheric Sciences) uses CC facilities to construct high-resolution ocean/sea-ice general circulation models to examine scientific questions in the waters of the three oceans that surround Canada. Applications range from understanding how climatic changes might affect take-up of gases such as CO2 in the Labrador Sea, to studies of sea-ice changes in the Northwest Passage of the Canadian Arctic Archipelago. Longer-terms goals include coupling with biogeochemical models to examine questions of ecosystem evolution.

Impact: The group has become part of the European Drakkar Consortium for high-resolution modelling. Within Canada, their models are being used in two large NSERC Climate Change and Atmospheric Change networks: VITALS, which focuses on the Labrador Sea, and GEOTRACES, which focuses on the Canadian Arctic. The model development done as part of these projects will eventually be transferred to Environment Canada and Fisheries and Oceans Canada, for use in operational forecast models.

Theme 5: Subatomic Physics and Astronomy

CC provides a key national resource that allows Canadian scientists to participate in large international communities of scientists, to access and analyze the data and to serve results to the world. CC is currently supporting all major experiments at domestic subatomic physics labs (TRIUMF, SNOLAB) and experiments at offshore subatomic physics labs (ALPHA, ATLAS, Belle-2, IceCube, T2K). CC also hosts copies of all Canadian astronomy data (in partnership with NRC/CADC/CANFAR). A few examples include:

Dr. Robert McPherson (University of Victoria, Institute of Particle Physics) leads the ATLAS-Canada collaboration, consisting of 39 Canadian faculty members and more than 150 researchers at 10 Canadian institutions. The ATLAS experiment at the CERN Large Hadron Collider is designed to study matter at the highest energies and smallest distance scales yet achieved in the laboratory. The goal of this research is to understand the fundamental building blocks of matter and the interactions (or forces) that dictate their behaviour.

Impact: The observation of the Higgs boson gives insight into the origin of mass of elementary particles and is generally considered to be the most important experimental result in particle physics over the last 30 years. It was named scientific breakthrough of the year by Science (2012), and resulted in the ATLAS and CMS collaborations receiving the European Physical Society Prize. The theorists responsible for proposing the Higgs mechanism received the 2013 Nobel prize in physics. CC provides the required Canadian computing and storage contribution for the analysis and simulation required for Canadian researchers to be part of this endeavour. Further, CC contributed extra resources specifically to the Higgs boson analysis.

Dr. Christopher Pritchet (University of Victoria, Physics and Astronomy) leads the Canadian Advanced Network for Astronomical Research (CANFAR). CANFAR is an operational research portal for the delivery, processing, storage, analysis and distribution of very large astronomical datasets. An innovative, but challenging, new feature of the research portal is the operation of services that channel the onslaught of telescope data through Canadian data networks to the computational grid and data grid infrastructure.

Impact: CANFAR is currently used by many Canadian astronomy projects. These projects are using data generated by peer-reviewed allocations of observation time using three of Canada's telescopes: the Canada-France-Hawaii Telescope, the James Clerk Maxwell Telescope and the Herschel Space Observatory, as well as data from other tools, such as the Hubble Space Telescope. The projects include the Next Generation Virgo Survey, the SCUBA-2 Cosmology Legacy Survey and SCUBA-2 All Sky Survey, the Pan-Andromeda Archaeological Survey, the Herschel Debris Disk Survey and many others. The portal is operated and maintained by the Canadian Astronomy Data Centre (CADC) in collaboration with CC.

Dr. Darren Grant (University of Alberta, Physics) leads the Canadian effort in the IceCube experiment. This research focuses on the study of high-energy astrophysical neutrinos, atmospheric neutrinos and the search for dark matter with the IceCube Neutrino Observatory located near South Pole Station, Antarctica. Canada (through CC) is one of the largest contributors of computational resources for this experiment.

"Part of the reason for the existence of Memorial's master's degree program in Scientific Computing is ACENET and CC, that we know that there's that resource out there. Many of the faculty at Memorial use ACENET, and having students come through our program allows them to have their students trained in HPC techniques. ACENET has been a very important part of this program, and we train our students and they go off to do work in industry."

Martin Plumer, Professor in the Department of Physics at Memorial University

Impact: In 2013, the experiment provided a first measurement of the highest energy neutrinos ever detected as well as the world's first measurement of high-energy astrophysical neutrinos, a measurement of neutrino oscillations in a new energy regime, and new, world-leading spin-dependent dark-matter sensitivity. These results have been featured on the cover of *Physical Review Letters* and a special extended research article in Science. The observation of high-energy cosmic neutrinos at IceCube was named the physics breakthrough of 2013 by *Physics World* magazine.

Theme 6: Computer and Information Sciences

CC provides a core resource for researchers across the country in computer and information science. Some examples include:

Dr. Yoshua Bengio (Université de Montréal, Computer Science) has been a leader in the area of deep learning, which has been immensely successful in the last two-and-a-half 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 performance levels of predecessor systems in just a few months of research [arXiv:1409.0473]. 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 an auto encoder-based generative model.

Dr. Michael Bowling (University of Alberta, Computer Science) has a research program focused on machine learning, games and robotics. Games have a long history of being a proving ground for artificial intelligence.

Impact: In 2014, Bowling's group was granted a special CC allocation to attempt to solve heads-up limit Texas hold 'em poker. The new techniques developed in this effort are ground-breaking, in that the computer must be able to make strategic decisions in a human-scale setting, while handling the wide variety of uncertainty that exists in the real world. The resulting work was published in Science (Vol. 347 no. 6218 pp. 145-149) and has garnered world-wide attention. The impacts of this technology range from keeping our airports secure to robust decision support in evidenced-based medical treatments.

Theme 7: Social Sciences and Humanities

Traditionally, research in social sciences and humanities has not been as ARC-intensive as in natural sciences and engineering. However, there is a long history of activity in several areas of social sciences and humanities, and there is strong and growing activity and interest in some disciplines in recent years. Some examples include:

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 test bed 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. 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. Stéfan Sinclair (McGill University, Language, Literature and Culture) leads and contributes to several international projects that focus on the design, development, use and theory of text analysis and visualization tools that are specifically designed for humanities-based inquiry and interpretation. An example is Voyant Tools (www.voyant-tools.org), a web-based environment that combines user-friendliness with flexibility and performance to allow scholars to easily create a custom corpus and perform a range of analytic and interpretive operations.

Impact: Voyant Tools is widely used in the digital humanities community for research and pedagogy — it averages about 40,000 visits and 500,000 tool invocations per month. Voyant Tools has been hosted by CC since 2010.

A.4 | Highly Qualified Personnel

MSI funded CC to build on an already existing regional framework, in order to create the national system for training and development for Canadian researchers and the CC experts that support their work. CC helps Canada attract and retain researchers by providing training and skills development for highly qualified personnel (HQP), ranging from undergraduates to professors.

These researchers rely on a network of ARC experts that provide training and work directly with researchers to help them analyze data algorithmically, in order to accelerate knowledge and discovery. CC resources greatly increase capacity to spot trends, tackle interesting challenges on a larger scale, and make connections previously unattainable. Compute Canada experts provide more than 100,000 staff hours annually, in direct support of researchers. In addition, since 2012, CC has documented delivering more than 54,000 hours of training to more 11,000 researchers at approximately 573 training events.

A.4.1 | Training and Skills Development: Scientific Support Activities

CC provides a network of experts to train and support researchers. It also employs more than 80 full-time-equivalent (FTE) scientific support personnel working at more than 34 locations across Canada. These staff members spend the bulk of their time providing one-on-one support to CC users. This kind of activity can provide tremendous efficiency gains through expert review of researcher code and through one-on-one teaching of optimization techniques. This improved research efficiency benefits the researcher who receives the support and also optimizes the use of the shared resource by freeing infrastructure resources for the use of others. Support staff deliver more than 100,000 hours of one-on-one consultations every year.

CC has succeeded in increasing its complement of scientific support staff by more than 45 per cent since the start of MSI funding, exceeding the 76 FTEs targeted by the original proposal to approximately 80 FTEs. Over the same period, usage has grown by more than 100 per cent by almost any measure; for example, the number of active faculty groups has risen from 1,250 to more than 2,700 — a 116 per cent increase.

"Belaid Moa helped me very quickly and effectively to set-up and prepare my simulation code on Hermes cluster. At that time I had started my PhD for few months and I was not familiar enough with compilation of large software. I was using a finite element library called rheolef which didn't have a large user group and the compilation procedure was very tricky. I emailed for help and Belaid got everything set up and ready in just a day. He even managed to compile the library with Intel compiler for better performance which was not officially supported by the developers. Actually, later the developers asked me to tell how I was able to compile with Intel because they had problems with it! I asked for instructions from Belaid and sent it to the developers and the next version of library was released with support for Intel compiler. So Belaid helped me and the developers as well!"

Ali Roustaei, University of British Columbia

Continued growth is expected, so CC plans to grow scientific support accordingly, particularly in the areas of bioinformatics and computational biology, by leveraging partnerships with the Canada's Michael Smith Genome Sciences Centre (Vancouver) and the HPC4Health (Toronto). However, aside from this type of partnership, CC does not project significant additional hiring in university support during the current MSI period, since its consolidation of systems and data centres is expected to transform job assignments for many current staff and allow increased support levels with a relatively stable headcount.

Since 2013, all of these efforts have been guided by a national working group. The Research Support Committee, in consultation with the regions, and the Chief Science Officer, make or recommend decisions concerning the provisioning of all forms of researcher support atop the national platform, to ensure that all researchers in Canada can access the support needed for their research, regardless of their location or of the mix of resources (computational and personnel) they need.

A.4.2 | Traditional Training Activities: Acquiring High-Level Skills for Research and Other Careers

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 courses on 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 CC resources.

"I first met Dan Mazur at an excellent CUDA tutorial he had given at McGill. Since then, I have been in touch with him several times to help me port and modify my GPU single node code to the K20 nodes on Guillimin. He is extremely knowledgeable and really went above and beyond the call of duty to ensure that my GPU code was working on the Guillimin K20 nodes. At one point, I exchanged 32 emails with him in order to pinpoint a really subtle and specific bug. He helped modify several of my GPU codes, as well as wrote several example scripts for me that helped me port my codes."

Mohammed Harb, McGill University

CC combines national, regional and local events across the country to create a full suite of learning opportunities, reaching more than 4800 researchers in calendar year 2014. The table below provides a detailed breakdown of events run by CC from 2012-2014 and the total number of attendees.

Event Category	2012 (7 months)	2013	2014	Total
Accelerator	4	10	18	32
Application	0	7	0	7
Chemistry	0	4	1	5
Cloud	0	0	1	1
Data	0	0	22	22
Database	0	1	0	1
Development	7	6	14	27
Globus	0	2	2	4
Hpc	0	5	5	10
Hpcs	1	1	1	3
Industry	0	0	3	3
Introduction	32	87	88	207
Matlab	1	2	2	5
Mpi	10	11	10	31
Openmp	3	7	8	18
Outreach	10	8	6	24
Parallel	11	8	6	25
Python	2	5	11	18
Scheduling	2	5	8	15
Science	14	9	18	41
Summer Schools	1	3	1	5
Systems	0	0	8	8
Teachers	0	0	3	3
User	9	18	7	34
Visualization	6	4	13	23
Total Events	**113**	**204**	**256**	**573**
Total Attendees	**3,285**	**3,569**	**4,857**	**11,711**

New courses, developed by CC, include:

- Accelerators (GPU, Phi, CUDA)
- Applications (Gaussian, Gromacs, FEniCS, Matlab...)
- Big data (Map-reduce, Hadoop)
- Development (programming, databases, revision control, parallel debugging, optimization, profiling)
- Globus (high-speed data transfer)
- Parallel Programming (general, MPI, OpenMP, ...)
- Scheduling (MOAB, Torque, optimal job submission strategies, ...)
- Visualization (Paraview)

Occasionally very specialized courses for particular discipline-specific software packages or techniques are offered. For instance:

- Calcul Québec Presentation at the School of Planning, Architecture, Art and Design and in the faculty of literature.
- ACENET/HPCS 2014 Big Data analytics
- WestGrid Sensor Technologies for Managing Disease and Disability
- SHARCNET EEG signal processing

CC's annual flagship event is the High Performance Computing Symposium (HPCS, www.hpcs.ca), the only national conference in Canada devoted to advanced research computing. Each symposium includes two days of tutorials, which provide advanced training to technical staff and users. Approximately 300 attendees participate each year. This conference is now in its 29th year.

With MSI funding, CC expanded its training and partnership initiatives to participate and lead international training efforts. Today, Canada partners with Europe, Japan and the US to deliver world-class learning opportunities. 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 (Toronto, June 21-26, 2015). Ten Canadian students and several Canadian instructors will participate this year. The event is 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 also became a partner organization to Software Carpentry. This organization runs workshops all over the world, training researchers the basic software skills they need in order to effectively and efficiently use ARC resources. The software carpentry organization develops open access training materials and provides instructor training classes which are being offered to CC staff. As a partner, CC has committed to organizing at least 20 software carpentry workshops across Canada per year and has effectively become the coordinating body for these workshops in Canada. This kind of partnership has been made possible by MSI funding and CC's new corporate structure.

Since 2013, all of these efforts have been guided by a national working group on outreach, training and education. This working group establishes the national training agenda and works to enhance the biannual staff training workshops. This year, additional support was provided for technical and non-technical training. Sessions were held on Globus technology and a full day project-management training session is planned for June 2015.

A.4.3 | Developing Compute Canada's Cadre of Skilled Professional Staff

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 is essential to CC's success and growth in the future. The charts on the next page illustrate the expertise of staff members based on their level of degree attainment.

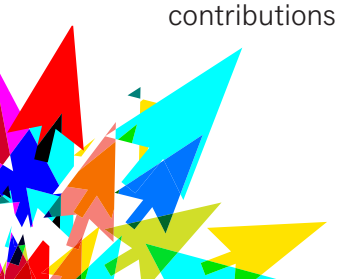

Compute Canada Staff by Education

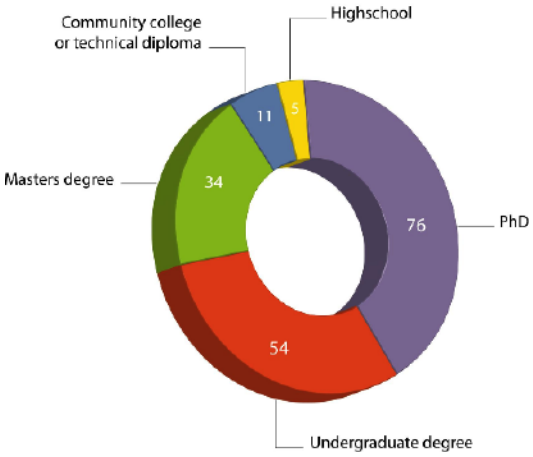

Compute Canada Staff by Discipline

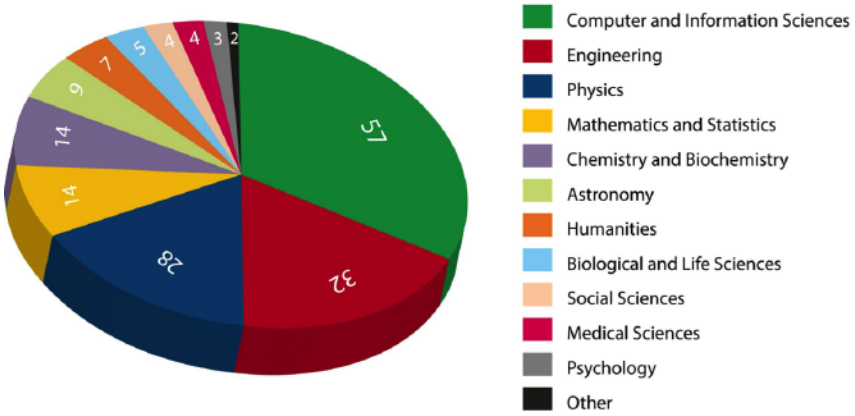

Activities targeting the advancement of this strategic objective are discussed further in *Appendix 6: Management and Operation Plan*. Specific professional and career development activities include:

- travel funds for external workshops, conferences and other events;
- workshops and seminars from expert staff;
- workshops and seminars from invited researchers;
- technical workshops and seminars from vendors;
- informal, internal working groups and collaboration venues;
- access to institutional training and education programs;
- access to compute resources for self-training.

In addition to these personally focused activities, CC undertakes several collaborative and team-building initiatives:

- The annual Supercomputing conference, held jointly by the IEEE and the ACM, is a unique opportunity for training and personal development, and a key opportunity for team-building across the CC organization.
- Every six months, CC holds a two-day workshop (TECC Summit) for staff, with one of these scheduled adjacent to the annual HPCS event. Historically, the focus has been technical topics. However the recent event in January 2015 expanded the scope to include work by our Education, Outreach, Training and Research Support Committees. These sessions promote sharing of best practices and national co-ordination of services. CC experts design and deliver workshops and presentations based on current gaps and challenges.

CC has also begun to offer staff-specific professional development opportunities through partnerships with other organizations. These include sending five staff members in 2014 and two staff members in 2015 as students to the Digital Humanities Summer Institute and sending 20 staff members in 2015 to a dedicated “train the trainers” course in partnership with Software Carpentry. These sorts of initiatives are only possible via MSI funding.

To illustrate the role played by CC staff in enabling excellent research, some quotes from researchers are provided on the following page.

"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 always only a Skype call away and will always tend to 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

"As a veterinarian trying to branch into research involving bioinformatics, the logistics of getting certain programs to work can sometimes feel daunting. Doug Phillips has been an incredible resource for helping me work through these challenges, making what can seem like an insurmountable obstacle look more like another regular road bump on the path towards completing my PhD."

Taya Forde, University of Calgary

"JF is amazing at what he does. His dedication to his user knows no bound. He answers questions quickly and clearly at every time of the day or even sometimes at night and even during the weekend. He can take the time needed to explain different concepts to use a supercomputer to any kind of users with any background. More than one time he helped us figure out why our code was slow. He also directed us to the right resources when he could not help us directly. Finally he was always really quick to help fix the experimental GPU cluster. All this helped us reach our deadline and let us focus on the science!"

Mathieu Germain, Université de Sherbrooke

"Dr. Pawel Pomorski in his day-to-day support of my research group has gone above and beyond what was expected of him to enable access to non-standard software that my research group uses. He has enabled the initiation of an industrial collaboration with BlackBerry through providing expertise in areas outside of my research group's strengths but required as part of the industrial research project."

Nasser M Abukhdeir, University of Waterloo

A.5 | Competitiveness of the Infrastructure

CC is the primary provider of ARC services to the Canadian research community and the only national organization dedicated to this purpose in Canada. CC's goals are different from many comparable international facilities because we provide ARC resources to all Canadian researchers, in any discipline, at all scales above desktop computing.

By contrast, most international facilities focus on one or two tiers of the Branscomb Pyramid. 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.

A.5.1 | Increasing and Improving Our Technology Platform

CC was established in early 2012, by which time nearly all of the most recent significant investments in CC's technology had already been made. CC was funded by a 2006 CFI National Platforms Fund (NPF) award of $60 million. Those investments contributed to a significant increase in computational resources in Canada, an increase that has maintained Canada's rank in terms of computing power per researcher. In 2008, Canada ranked 16th in the world on this basis, while in 2013, Canada ranked 15th in the world. In 2012, CC reached 198,000 cores of compute capacity across 50 systems (an aggregate of 1.85 Petaflops), as well as 20 Petabytes of storage. However, these investments, and the maintenance of Canada's technological competitiveness, are not due to MSI funding, since this funding only covers CC's operating costs.

Nevertheless, in the MSI era infrastructure investments continue to be made, using the last of the NPF award or pooling individual project funds for equipment purchases, as well as the skills of MSI-funded staff to guide the procurement and installation process:

- GPU systems
  - McGill (Guillimin): 100 nVidia K20, 100 Intel Phi
  - SHARCNET (Monk): 108 nVidia Tesla M2070
  - Laval (GPGPU): 120 nVidia K20, 48 K80 on the way; NPF funded 31 per cent of this system, with the balance from 5 research grants from 2 universities.
- 2 Openstack cloud systems ("CC-Cloud")
  - East (installed at Université de Sherbrooke): 576 cores, 128GB/node, 100TB disk
  - West (installed at University of Victoria): 640 cores, 256 or 512GB/node, 200 TB disk

On a stand-alone basis, these investments have not increased our competitiveness significantly, but when combined with expert people and enhanced services, they are powerful tools for Canada's researchers. For example, the investment in GPUs has been complemented by outreach, training and support for the effective use of this technology (as described in the HQP section above). Similarly, the CC-Cloud service is still in beta mode, to ensure users have all the tools and support they need to make using this new service worthwhile.

CC regularly makes small investments in its fleet of systems to improve performance. The following is a compressed list of activities since MSI funding began:

- Memory upgrades on some compute nodes (Jasper, Bugaboo, GPC)
- Upgraded interconnect (Toronto)
- Disk upgrades (Sherbrooke, Laval, SHARCNet, Toronto, Queen's, U de Montréal)
- Establishing archival tape systems (Sherbrooke, U de Montreal, Queen's)
- Investing, in collaboration with CANARIE and HEPNet, in a nationwide perfSONAR monitoring platform to monitor inter-centre network performance
- Network connectivity improvements (U de Montréal, Laval, Toronto, Queen's)
- National maintenance contract for DDN storage systems
- Improved reliability for auxiliary systems (Toronto)
- Creation of special purpose compute nodes to meet user needs (large memory Gaussian node at Queen's, GenAP fat memory nodes at Sherbrooke)

All of these actions were undertaken to improve the productivity and effectiveness of CC's users, as we seek to continuously improve CC's competitiveness.

A.5.2 | Increasing Our Investment in Skills and Expertise

Despite limited capital investment, CC has more than 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 aging equipment.
- System administrators ensure 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 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. To date, CC has achieved its commitments and increased its capabilities as shown in the table on the next page.

All of these staff members are highly skilled, with specialized training, experience and qualifications. Many have advanced academic degrees (more than 70 have PhDs). The quality of this team is competitive with the quality of staff at any comparable facility internationally.

The original proposal to increase research support staff was based on the expectation of increased demand from the research community. Within the overall staffing envelope detailed above, CC also proposed a number of specific staffing commitments in its original submission to the CFI:

- Three FTEs in biological and medical sciences: CC hired David Morais, Jean-François Lucier, Jean-François St-Pierre and others, and, through our partnerships with Canada's Michael Smith Genome Sciences Centre, the Toronto Hospital for Sick Children and the University Health Network, increase this number by an additional 13 specialists in bioinformatics and computational biology. These strong partnerships allow CC to serve this rapidly growing field (see Appendix 11).
- Six FTEs in the "path to exascale": No additional staff members were hired with this focus. CC will develop an appropriate strategy for this development as we execute our own infrastructure renewal plans, consistent with available funding, and continue to work with international partners. It is likely that some assignments with this focus will result from

Staff Category	Starting Staff (FTEs)	Proposed Staff (FTEs)	Staff as of Jan. 1, 2015 (FTEs in Fiscal 2014/2015)	Growth
Research Support	55.8	76.5	80.9	25.1 (45%)
Technical / System Administrators	63.6	69.5	73.0	9.4 (15%)
Administrative	17.5	20.25	25.4	7.9 (45%)

job evolution for existing technical and system administrative staff.

- Data management: CC hired Sean Cavanaugh, who is an expert in big data, data management and data management tools. He is well versed in big data management and analysis, including Hadoop, nosql, and other emerging technologies.

- One FTE for digital humanities: To invigorate its commitment to the digital humanities, CC hired John Simpson, PhD, in December 2014 to lead its digital humanities team, which is made up of 9 additional staff members. Dr. Simpson has a diverse background in philosophy and computing and, as a member-at-large of the Canadian Society for Digital Humanities, is an active contributor to the digital humanities. His focus is on building programs for outreach, training and support to the digital humanities community. He also works to understand and respond to the needs of this community through his position as chair of CC's digital humanities working group.

- Six FTEs for visualization: CC designated Alexei Razoumov, as its visualization coordinator. The visualization team is comprised of 15 additional existing staff.

- Three business development officers: Compute Canada hired an Executive Director of External Affairs (Kelly Nolan) who is developing strategic partnerships with like-minded organizations. CC's CSO and CEO are working with government laboratories and departments and several industry partners. In Atlantic Canada, Michele Fash has been hired by ACENET as Marketing and Business Development Manager. National staffing for business development will take place in coming months to maximize these opportunities.

Since the start of MSI funding, demands on CC have grown by more than 100 per cent by almost any measure, so the 45-per-cent increase achieved to date suggests that CC is just keeping up with this growing demand and that further increases are warranted. However before significantly increasing the level of research support personnel, CC intends to strengthen the management of the research support activity to ensure that CC is providing the right kinds of support to researchers in each field across Canada. In addition, as CC executes its consolidation and renewal plans (described in Section B), we expect significant evolution in job assignments for many of the more than 70 existing technical staff and system administrators, which should allow CC to further increase its research support in areas requiring additional capabilities, as well as to increase specialization in key areas, without significant additional hiring.

A.5.3 | Deploying More and Better Services for Researchers

CC is evolving from a resource-based organization to a service-oriented organization, with a new focus on providing researchers across Canada with well-designed, easy-to-use enterprise services that make their work more productive. Simultaneously, CC is less project-based, and instead acting as a national-level corporation. A few examples of these services, all enabled by MSI funds, are:

- **"Cloud" services.** CC has operated cloud environments for specific research platforms (e.g. Nimbus for CANFAR, DAIR at two sites for CANARIE) since the outset of MSI funding.

As described above, CC has used NPF capital funds to install two OpenStack cloud systems, named CC-Cloud, which will be made available as a generally accessible cloud infrastructure for Canadian researchers. In addition to the capital purchase, this has involved the formation of a national cloud working group and considerable investment of staff time to support this new ARC paradigm. This will serve researcher demand for virtual systems and virtual clusters, often with a specific software load. The virtual systems may be persistent, versus typical batch-oriented supercomputer jobs that only run for a limited time. The recent expressions of interest from CFI's Challenge 1 Stage 1 served to reinforce the need for such cloud-based services.

- **Data storage and transfer.** CC has seen growing demand for active and long-term storage solutions from the Canadian research community.

  CC's first service introduction (early 2014), was OwnCloud, a service similar to DropBox that exposes CC storage in a user-friendly way. OwnCloud is designed for relatively small transfers, so it does not address the needs of every user, but it has already transformed the pattern of work of many researchers, especially those who had never before used CC storage in their research.

  At a more "enterprise-grade" scale, in December 2014, CC deployed managed Globus endpoints at 24 major storage facilities across the CC network. The Globus service makes it easy for Canadian researchers to transfer large datasets to and from their own equipment and instruments as well as between CC storage facilities. As the de facto global data-transfer standard, Globus also allows researchers to share data with collaborators and institutions across Canada and internationally. CC's adoption of Globus helps bring CC on par with these international research facilities. As of March 2014, CC researchers have used Globus to transfer nearly 700TB of data in nearly 60 million files.

- **Platforms and portals.** In fall 2014, CC created a research-platforms and portals competition. The Canadian research community had been largely unaware of the ability of CC to support research platforms and portals. This competition formalizes and advertises this capability: CC selected 13 for support in 2015.

- **Continuous service improvement.** CC has executed a number of smaller initiatives aimed at improving service to researchers, including:

  - Negotiating improved licensing arrangements for key software applications (e.g. Allinea DDT and MAP, Fluent)
  - Improving the user experience through enhanced online services and outreach, online status indicators, email notifications, better documentation and wikis, consolidated ticketing systems, as well as implementing a formal "Tier 1" helpdesk service in Western Canada, which will be a model for deployment of national-level helpdesk services.
  - Forming working groups to improve support for key activities (CC-Cloud, CANFAR, perfSONAR)

A.5.4 | Establishing National and International Leadership

Developing and implementing harmonizing software and services: Using MSI funds, CC has partnered with Globus and XSEDE to build a common software stack across the CC network. This leverages work being done on software architectures in XSEDE and incorporates fundamental services such as identity management and data management. This work began in spring 2015 and will enable greater harmonization across CC sites and easier movement between those sites by CC researchers. This will allow researchers to more easily exploit a range of CC facilities, as well as preparing them for the planned consolidation and renewal of our resources.

Working with partners in the digital infrastructure: CCis working collaboratively with CANARIE and other partners in the eco-system on several national projects and collaborating on the CC technical upgrades. CC is working with CANARIE and the CFI to develop a coordinated long-term approach to the funding and provision of high-speed networking, high performance computing

and software tools. CC and CANARIE are working together to develop scalable solutions to the problem of research data management in Canada. With CC's recent identification of the location of CC's next round of system installations, and current development of a national data infrastructure plan, CC and CANARIE are coordinating plans for the high-speed networking that will be required to support those investments.

Breaking down regional and disciplinary silos: GenAP is a CANARIE-funded bioinformatics project that directly involves several CC staff. The software distribution mechanism used by the project is the CERN Virtual Machine Filesystem (CVMFS). This was possible due to existing expertise among CC staff that support the ATLAS particle physics project. The technology developed for particle physics and supported by CC staff was re-deployed in a genomics context and is currently being disseminated across several CC sites. CVMFS is now being considered for use in other disciplinary projects that need this capability.

Leading projects to ensure informed policy and strategic purchasing decisions: Two such projects are a secure computing pilot and a pilot in research-data management for the "long tail" of research:

The **secure computing pilot** involves next-generation network monitoring and firewall appliances (purchased from NPF funds) and the creation of a dedicated compute cluster that can host sensitive private data. The cluster is sharable, but also partitionable such that data in one partition is isolated from data in the other. This system is currently being used for population data studies (with Population Data BC) and for a health-related industry pilot (with a start-up, FusionGenomics). It has allowed CC to develop technological and policy solutions for sharable private data infrastructure; infrastructure that is critically needed by Canada and CC.

The **research data management pilot** has been undertaken with partners, including Research Data Canada, the Canadian Association of Research Libraries (CARL) and CANARIE, to address the infrastructure associated with a possible future requirement that all publicly funded data creation must be covered by data management plans. This pilot involves the deployment of Canadian-developed research data software stacks (e.g., from the library community) on CC systems and with extensive testing of their scalability. This pilot will involve ingesting research data from researchers in a variety of disciplines, and the design of future hardware and software systems to handle these data at scale.

Taking leadership roles on international projects: Since 2012, Compute Canada has played a leading role in several international projects.

The first formal project of this type involves the annual **International HPC Summer School** (see HQP Section).

CC is currently working with comparable international ARC facilities to explore the possibility of evolving the Globus data transfer and sharing service into a federated international model. Currently, Globus is built in part on the Amazon Web Services platform, but this raises security and privacy concerns for many potential users of the Globus service. In the federated model, partner countries would run their own Globus system and ensure local privacy and security regimes are respected, local language support is provided, etc. CC hosted the first organizing meeting for this project at SC14 in New Orleans and plans on participating in a follow-up meeting at ISC15 in Frankfurt.

International metrics: The expert review of CC's original proposal recommended developing appropriate metrics that would allow useful comparison with international facilities, while at the same time ensuring that CC aligns its efforts with the objectives of the strategic plan. CC has already starting working with its sister organizations across the world (XSEDE in particular) to identify the most useful and comparable metrics to support benchmarking, best practices sharing and effective operational planning. To date, CC has used the metrics collected through its CCDB database to measure and track resource usage — these efforts will continue with the adoption of more sophisticated tools (XDMOD and Splunk are currently being assessed) that will enable both benchmarking and the creation of operational dashboards for the entire CC portfolio of resources.

A.6 | Access to the Infrastructure

CC has always had a policy allowing all eligible researchers access to its facilities. However, the majority of CC compute resources (approximately 80 per cent) on each system are allocated through a competitive peer-review process led by the CC's Resource Allocation Committee (RAC). This committee review ensures that excellent science guides the usage of CC facilities. The remaining 20 per cent of available compute resources are left to the "default allocation."

There are two main reasons for maintaining a 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 their research. While an 80 per cent competitive allocation is the global target, it is not a hard rule applied to every system. Some CC systems have been allocated considerably beyond 80 per cent, while some are allocated below the 80 per cent level. These choices are assessed and monitored each year.

The CC RAC process has evolved considerably since the awarding of MSI funds in 2012. The original CC RAC process involved regional (local) LRAC committees handling most requests with a national committee (NRAC) in place to handle large region-crossing requests. The regional and national committees each had the power to assign a pre-determined fraction of each available resource. In the MSI era, this process has changed from a regionally focused process to a science-focused process. Currently, 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. 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 the request and the optimal site on which to implement it. This move to a national science review process guided by local system expertise better fits the character of the new national CC corporation and the goals of the MSI program.

In addition to funding the basic maintenance and operations of these facilities, MSI funding supports the salaries of expert staff, which enable researcher access by providing training and personalized support to them.

A.6.1 | Impact of MSI Funding on the User Base

The CC user base, as measured by the number of research groups, has grown substantially over time. It has increased from 1,250 registered groups at the time of the original MSI proposal to more than 2,700 faculty research groups today (including more than 300 Canada Research Chairs). CC also helps upwards of 8,500 researchers/students across Canada, in a wide range of disciplines, who use CC resources as an essential tool.

CC serves researchers from all parts of Canada. The figure below shows the breakdown of faculty research groups by the region of the principal investigator. CC also serves researchers from all disciplines, as reflected in the Figures on the following pages.

"We're increasingly reliant and dependent on CC computers to do our work. We simply couldn't do it without the resources that CC has to offer. You know, being able to access these state-of-the-art computing resources is crucial to making sure we stay competitive internationally in the research that we do, in order to stay on the leading edge of the research internationally."

Dr. Katja Fennel, Canada Research Chair and Associate Professor at Dalhousie University

Active faculty by region (January 1, 2015)

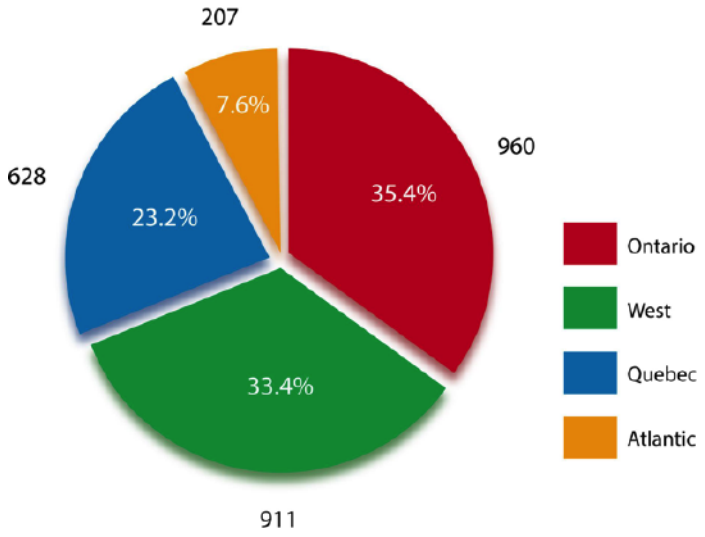

Active faculty by research area (January 1, 2015)

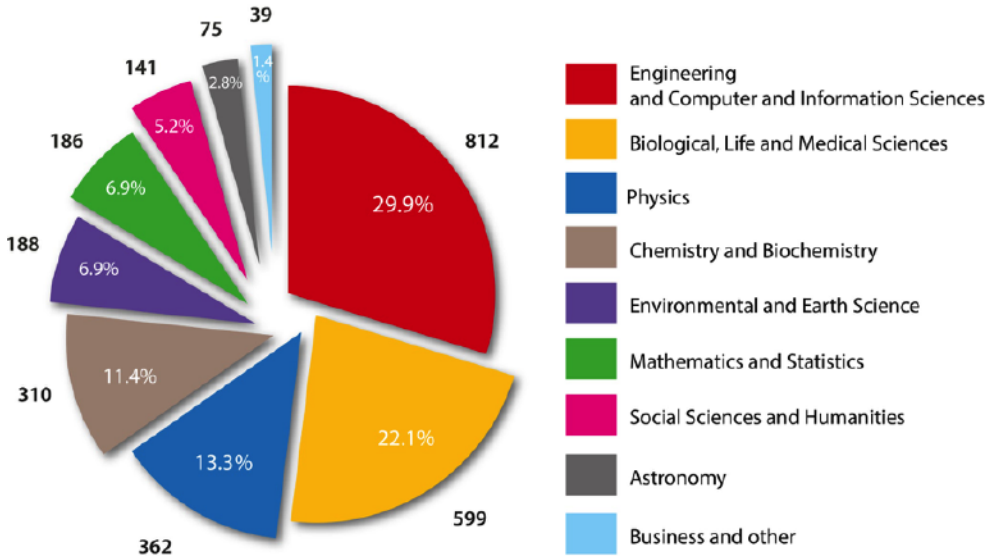

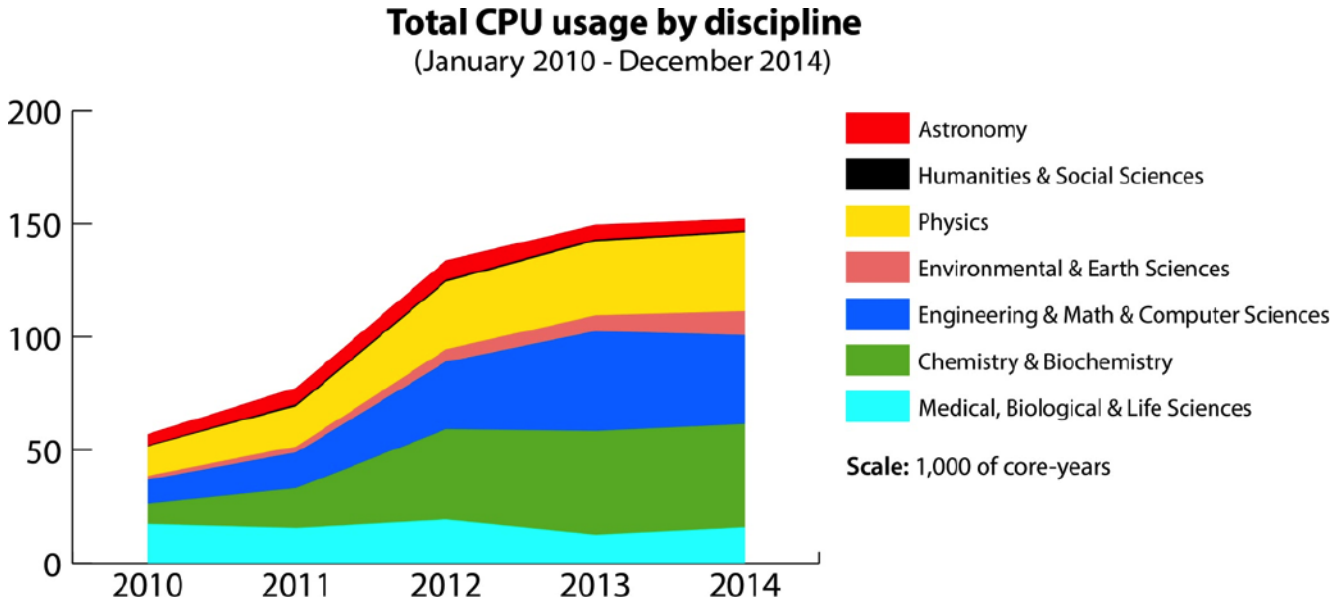

CC allows international users access to its computing infrastructure through sponsorship by a Canadian researcher. This is a critical service for researchers involved in international research collaborations. In 2014 there were 389 registered international users.. This international access is an integral component of a global system of reciprocity that enables Canadian research teams to gain access as needed to international resources as well, in the context of important international initiatives that depend upon the integration of data from already large databases residing at multiple laboratories located around the world.

CC also hosts several large national platforms. These platforms share their data or their allocated compute resources through their own access mechanisms (e.g. through a portal or grid tools). The table below illustrates the diverse user base of four such platforms in different areas of study. The numbers represent unique users in the 2013/14 fiscal year.

Platform	Canadian Users	International Users
ATLAS (particle physics)	184	3661
CANFAR (astronomy)	253	4181
CBRAIN (neuroscience)	172	82
Ocean Networks Canada	87	68

A.7 | Benefits To Canada

The extent to which Canada is able to exploit big data and high performance computing contributes directly to our economic success as a nation and our role as participants and leaders in the knowledge economy. CC helps the research and innovation community accelerate discovery and this is playing a central role in solving some major challenges facing Canadian society. Managing these resources efficiently on a national scale, Compute Canada is maximizing existing capacity and ensuring research that meets the threshold of excellence is supported.

The latest federal report, *Seizing Canada's Moment: Moving Forward in Science, Technology and Innovation 2014*, states:

> ***Supporting Cutting-Edge Infrastructure:*** *Our Government will enhance Canada's research capacity through investments in transformative infrastructure projects that underpin world-class research and enrich Canada's research landscape. This will be achieved by initiatives that include: Working with partners and government to develop a digital research infrastructure strategy to create a world-leading research and innovation ecosystem in Canada. This will also serve to advance Digital Canada 150 by positioning Canada as a global leader in "big data." We will include new policies on research data management and storage, and a co-ordinated long-term approach to the funding and provision of high-speed networking, high performance computing and software tools, to be developed by the Canada Foundation for Innovation, CANARIE and Compute Canada.*
>
> **Section 6.2.3 – www.ic.gc.ca/eic/site/icgc.nsf/eng/h_07472.html**

As the national advanced research computing organization, CC increases the capacity for research across the country by leveraging storage, computational resources, expertise and services. CC has enabled Canadian researchers to remain globally competitive, enhancing Canada's ability to attract and retain talent with the brightest and best minds. These cohort's research success and global research collaborations are inextricably linked to their access to advanced research computing infrastructure.

Canadian researchers increase the number and complexity of their experiments through simulation, visualization and modelling, and generate faster results at lower cost in terms of capital investment and human and financial resources. CC positions and maintains Canadian researchers on the forefront of global research opportunities, innovation and discovery. Specific additions to the national services and roles since the MSI award include Globus, Cloud, Visualization, Security and Digital Humanities.

A.7.1 | Achievements in the Transfer of Research Results to End Users Made Possible by the Award Over the Funding Period

CC takes a "multi-modal" approach to serving the research community and enabling international and industry partnerships, dissemination of best practices, consulting services, contract research, new product and process development support. This leads to patents and new companies, and the development of new standards, policies, regulations and codes of practice.

We play a central role in contributing to the success of many major Canadian science and innovation investments focused on end user and private sector engagement. Examples include: Networks of Centres of Excellence (including: BioFuelNet, MEOPAR, ArcticNet), Canadian Light Source, SNOLAB, TRIUMF, ATLAS, Ocean Networks Canada, and CIHR institutes.

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.

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. They also train highly qualified personnel for employment in extremely competitive sectors, such as aerospace engineering and film/entertainment. CC infrastructure and expert staff support a broad range of research projects across all tri-council disciplines. These 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. *(see full list in Appendix 3: Knowledge Transfer/Tech Transfer).*

"A significant amount of Bombardier's joint research with academia relies on numerical modeling and simulation to test new aerodynamic methods, new composite materials and other systems technologies before testing them in a factory setting. Compute Canada provides access to its infrastructure and expertise to the university researchers that make these computationally intensive tasks possible. We fund a variety of different projects and increasingly these are collaborative projects where engineers in our companies and engineers in universities will have to work together. That requires many projects to use the same computing architecture, if not the same computing platform, and for the work we support, that will be Compute Canada. It's in our interest to see that platform optimized."

Fassi Kafyeke, Director Advanced Design and Strategic Technology, Bombardier Aerospace, Montreal

"I can't live without Globus. The files that I process are huge and I move them often between Silo, Jasper, and Breezy. I also find Globus way faster and also more reliable than any other services. For someone who does big data analysis, it is a huge plus to have Globus."

Charles Martineau, Sauder Business School, University of British Columbia

Research Area: Empirical Asset Pricing, Microstructure and Multifractals

Industry partnerships are currently active in the aviation, automotive, ICT, film and gaming, medical devices and diagnostics, oil and gas, fisheries, forestry and optical-glass sectors *(see full list in Appendix 3).*

Datasets are valuable national assets that have been acquired through decades of significant investment from Canadian taxpayers. The international datasets for major projects such as CANFAR, ATLAS, Ocean Networks Canada and several genomics projects are housed and managed using CC resources. This infrastructure directly contributes to our ability to participate in large global research challenges.

For example, data used in the CANFAR project is gathered from major astronomical facilities, including the Hubble Space Telescope, the Canada France Hawaii Telescope, the James Clerk Maxwell Telescope and the twin 8-m telescopes of the Gemini Observatory. CC manages full copies of the data from this project and supports the tools for its manipulation, storage and accessibility.

Researchers working with CC resources reported partnerships jointly using CC infrastructure and support from 65 industry partners *(see Appendix 3)*. The development of national services, such as Globus, make it easy to move, synchronize and share big data. The CC-Cloud computing environment, currently in beta testing with some pioneer users, are especially enabling to groups in the digital humanities sectors. The new visualization working group addresses visualization of complex models and output data, seeking trends for vast amounts of data and helping to accelerate results. These and other services further enhance the robust shareable advanced research computing infrastructure and facilitate knowledge sharing and international collaboration.

Through partnership with Compute Canada, world-class institutions and their international and local partners are directly applying research efforts to real-world challenges in industry, health and policy areas.

Some examples include:

- Canada's Michael Smith Genome Science Centre: A Compute Canada project with an agreement to support researchers across the country in bioinformatics and for Compute Canada technical experts to support their ARC requirements.

- The International Human Epigenome Consortium (IHEC): IHEC is a global consortium with the primary goal of providing free access to high-resolution reference human epigenome maps for normal and disease cell types to the research community, supported by GenAp. The consortium uses sophisticated software developed by Compute Canada experts that cuts through the complexity of genomic data and makes it accessible to researchers and practitioners.

- Bombardier is partnering with Compute Canada and Calcul Quebec for a large-scale multi-institutional R&D research project.

A.7.2 | Most Important Benefits to Canada Realized to Date

Accelerating discovery and innovation in Canada is the overarching benefit to Canadians. Compute Canada experts and resources provide the environment to expand and explore more complex and interesting challenges which lead to better innovations and benefits. A small sample of some key benefits realized to date appears on the following pages *(see full list in Appendix 3).*

Information and Communications Technologies

Programming by Optimization for Resource Allocation, Scheduling and Verification
Dr. Kevin Leyton-Brown (UBC) and Dr. Holger Hoos, University of British Columbia

A team led by Steacie Award winner Dr. Kevin Leyton-Brown and Dr. Holger Hoos (UBC) uses CC resources to develop methods to be applied to improve state-of-the-art algorithms for problems such as scheduling, time-tabling, hardware and software verification, industrial process planning and optimization and supervised machine learning. This project is facilitating better solutions of hard computational problems important to everyday life, such as the optimal allocation of resources (e.g., energy, natural resources, health services, education), scheduling, and verification of hardware and software. The results of this research have led this team to work with IBM Canada and Actenum.

Department of Physics and Physical Oceanography, Memorial University
Dr. John Whitehead and Dr. Martin Plumer

Drs. Whitehead and Plumer use theoretical and numerical models to study fundamental physics with very real-world applications. Their research into the magnetic properties of thin ferromagnetic films — those at the core of many data-storage devices — leads to faster, denser and more energy-efficient devices for memory and computation. Western Digital sought help from Memorial University professors to develop a new generation of hard-drive technology to store more data at less cost, to better meet consumer expectations. Access to CC's resources has allowed Whitehead and Plumer to enhance and extend their work on similar problems facing the magnetic recording industry. The Memorial University team was able to use CC resources to do on a computer what would have otherwise required making and testing 10 prototype devices and testing theories on 10 different devices, a process that would have taken months and cost millions.

Advanced Manufacturing

Advanced Design and Strategic Technology, Bombardier Aerospace, Montreal
Dr. Fassi Kafyeke, Director

CC, in collaboration with Dr. Fassi Kafyeke, is supporting a joint research project with three Quebec universities to study numerical modelling and simulation that validates new aerodynamic methods, new composite materials and other systems technologies before testing them in a factory setting. CC provides university researchers with access to its infrastructure and expertise to make these computationally intensive tasks possible. The aerospace manufacturing industry is one of Canada's most R&D- and export-intensive industries. However, one of the defining challenges for the 21st century is understanding and reducing this global industry's environmental impacts. Together, Bombardier, CC experts and Canadian researchers are finding ways to manufacture planes that are quieter, lighter and more fuel-efficient.

Institute for Aerospace Studies at University of Toronto
Dr. David Zingg

Dr. Zingg has become a recognized world leader in the field of aerodynamic shape optimization. In addition to developing a world-class methodology for aerodynamic shape optimization, the group has applied the methodology to investigate several unconventional aircraft configurations and concepts, including a major collaborative project with Bombardier Aerospace on environmentally-focused aircraft. Outputs of the research include an extensive software package known as Jetstream capable of aerodynamic analysis and aerodynamic shape optimization and applicable to the development and investigation of novel aircraft configurations. The impact of this research comes through its influence on future algorithms for aerodynamic shape optimization as well as its influence on the design of future aircraft, particularly at Bombardier Aerospace.

McGill University
Dr. Nikolas Provatas

Advanced materials are a crucial part of Canada's economy. The prosperity of sectors such as aerospace, automotive, microelectronics, energy, pulp and paper, to name a few, is critically dependent on the development of advanced materials with improved properties. Dr. Provatas uses advanced research computing and dynamic adaptive mesh-refinement techniques to investigate the evolution of solidification microstructures in metals, a process underlying most industrial alloy casting processes. Using CC resources, Provatas' research has attracted the interest of Novelis Global Technologies and generated two NSERC strategic network grants with support from Dofasco, Evraz Steel, Novelis, GM Canada and, most recently, IBM Canada.

A.7.3 | Training Highly Qualified Personnel for Industry

CC has provided internationally competitive research, education and training opportunities to student and technical staff at Canadian universities, opening up new and long-term career paths in this field in Canada.

There have been 2,021 HQP-positive outcomes reported since 2012 out of a total of 2,444 user role changes. Role changes occur when a user transitions from one defined role type (such as "undergraduate") to a new role type (such as "graduate student"). Many such changes are reflective of career advancement or progression. At least 83 per cent of the role changes have resulted in users advancing their careers, either by moving to industry or the public sector or by advancing academically. A breakdown of these role changes is provided in the table on the following page.

Destination	Percentage
Faculty positions	3.7%
Post-doctoral positions	20.5%
Industry	25.4%
Public sector positions	8.4%
Further education	42%

The expertise developed in computational science by CC users has been extremely valuable. Examples include Dr. Katja Fennel, Canada Research Chair in Oceanography at Dalhousie University, who reported that her students' work in ocean modelling resulted in finding employment in digital film production in Montreal, and Dr. Nikolas Provatas of McGill University, who reported that more than half of his students are hired by industry and continue their work in advanced materials in commercial settings.

There is a continual influx of new users, usually at the student and post-doctoral level. These users undertake basic training either through CC or through their graduate courses and institutional training 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, some are becoming expert computational scientists. *Please see the HQP section for more details.*

A.7.4 | Enhancing the Health and Wellbeing of Canadians

Partnerships with Canada's Michael Smith Genome Science Centre, The University Health Network and The Hospital for Sick Children in Toronto are examples of centres delivering primary care and conducting research. CC accelerates the ability to deliver cutting-edge healthcare by providing the computational and data-management services necessary to deliver personalized medicine and conduct health research that is data-intensive in an accessible and secure environment.

Fusion Genomics

A Canadian company that was able to transform its discovery research into a commercial interest. Access to powerful computing resources, supported by CC and WestGrid, enabled a Simon Fraser University-based research lab to commercialize products that are available in the marketplace today. Fusion Genomics has a number of different research assays for use in cancer and infectious diseases. Its line of innovative, highly sensitive and accurate assays are designed around its Quantum Probe™ technology and can be rapidly developed and configured for a number of applications. A team led by Dr. Poul Sorensen, Senior Scientist at the BC Cancer Agency and Professor of Pathology and Laboratory Medicine at the University of British Columbia, developed the first product: ChildSeq-RNA. This is the first diagnostic tool to employ next-generation sequencing (NGS) technology to produce highly sensitive results (at the DNA level) in a fraction the time. Childhood sarcomas are aggressive cancers that affect the bones and soft tissues of children and adolescents.

Drug design for chemotherapy

Jack Tuszynski, University of Alberta

The major thrust of the Tuszynski lab is *in silico* drug design and *in vitro* drug testing for cancer chemotherapy applications. Without the resources supported by CC, years of additional effort would have been required to achieve the same results. One of the computationally designed anti-cancer agents, CR42-24, has now entered preclinical trials with the hope of reaching the clinic in one to two years. Furthermore, this team has initiated large-scale computational work in the area of antiviral agents and their first results produced potential compounds for treating hepatitis C. Other projects with near-term promising results include testing new vaccines, liver disease RNA-based drug development, a better understanding of Alzheimer's disease and control and modelling of infectious disease outbreaks.

A.8 | Governance

CC's governance strategy is designed to maximize trust, communication, transparency and accountability with users and stakeholders, as the organization consolidates and renews its infrastructure, and leads the transformation of its professional staff. In particular, CC recognizes several kinds of stakeholder institution: institutions operating facilities on behalf of CC, institutions providing staff and expertise to the CC enterprise, and institutions whose researchers use CC resources. In addition, CC makes it possible for institutions of any kind to participate directly in the governance of CC by becoming "members" of the Compute Canada Corporation.

CC's governance model is unique, combining a federation of institutional consortia with a corporate national organization, which reports to an independent board. CC seeks to follow best governance practices from the business and not-for-profit sectors.

When CFI awarded the CC project with MSI funding in March 2012, the CC enterprise was not well defined, and the CC Corporation did not exist. CC's transformation from a loosely co-ordinated "consortium of academic groups and regions" to a clearly defined national enterprise with effective governance, able to deploy resources across the enterprise and across Canada was in large part thanks to CFI's funding. MSI funding has already enabled CC to offer improved service, to optimize usage, to allocate capacity based on scientific merit, to maximize resources and to develop a clear value proposition to Canadian researchers and research institutions. CC is now launching a series of new national services and partnerships and is positioned for a new phase of improvements in service and competitiveness by the end of the project period.

A.8.1 | Governance Improvements Made Possible by MSI Funding

CC is the realization of an alliance of regional partners wishing to provide a national advanced research infrastructure and maximize expertise sharing across Canada. This evolution from regional consortia has offered many benefits. Most important, it has maintained positive local relationships with researchers and institutions, while focusing the national body on optimization of services, investments and science-based allocation. Many user groups rely fully on CC for their research, so ensuring these conduits remain accessible and promoting organizational transparency is essential. CC has adopted a robust governance model — guided by the recommendations of the review as well as the Advisory Committee on Governance and Management (convened in June and July 2012). To that end, CC has:

- Incorporated, rather than operating as a consortium or partnership of academic institutions.
- Chosen a membership structure for the CC corporation, which establishes direct accountability to members.

- Created an independent corporate board in which directors have the fiduciary interest of the corporation as their first priority and on which only a minority of directors may come from institutions or organizations served by CC.

- Hired a strong management team for the corporation, with clear authority, responsibility and accountability for the operation and performance of the corporation.

- Developed an inclusive, yet focussed, strategic plan to guide the direction of the organization.

- Continued to enhance and strengthen the role, governance and management of the regional organizations, communicated effectively with all partners and users, and worked to formalize the internal structure with appropriate agreements between CC, its regional partner organizations and the institutions we serve.

MSI funding explicitly made all of these actions possible. Documenting and formalizing these relationships will ensure the federated organizational model is robust and will operate effectively and efficiently for the benefit of all Canadians.

CC is now implementing additional steps in the area of governance:

- Developing effective inter-institutional agreements (IIAs) linking all participating institutions within each region, and each region with CC. IIAs will do more than enable the flow of funds; they will authorize transfers of funds and performance expectations among signatories, while respecting the rights and obligations of the institutions that employ CC staff, that actually own CC assets, and/or that have agreements with other agencies contributing funds to CC.

- Establishing additional governance mechanisms, such as a membership council, to provide a forum for consultation and consensus-building that will support CC and the regional organizations.

- Creating service-level agreements (SLAs) with institutions, defining respective rights and responsibility in connection with the services CC provides to that institution and its researchers. Key components of the SLA:

  - Universal service: availability of default resource allocations regardless of institution or discipline, regardless of the physical location of ARC systems
  - Acceptable use policy
  - Intellectual property, privacy, confidentiality
  - Terms of reference for, access to, resource allocation process
  - Financial arrangements — how CC aligns cost burdens with usage
  - Contribution of hardware in exchange for allocations

A.8.2 | Comparing Our Governance Model with International Best Practices

Best practices in governance require clear transmission of accountability from the funders, stakeholders and intended beneficiaries of an organization through all of its mechanisms of governance. CC's model as a not-for-profit corporation establishes a clear relationship between the members (the beneficiaries) and the board, management and employees of CC. CC and the four regional organizations extend that accountability relationship through to the individuals working in each regional organization and on each campus in order to fulfil its mission to serve the research community. This relationship between CC and

the regional organizations is currently informal; agreements between CC and each regional organization will formally extend accountability to regional and institutional staff.

The CC board, elected by the members, is a critical conduit for this accountability and, consistent with international best practice, has responsibility for the following functions:

- Legal and primary oversight
- Setting the strategic direction, guided by collective board experience in both industry and academia, as well as experience in advanced research computing
- Managing performance
- Financial control
- Managing risk
- Appointing, training and monitoring senior management
- Policies, procedures and compliance
- Reporting and communications

The nomination of directors is based on skill and experience, rather than representation, in accordance with international best practices.

Three committees whose members are drawn from the board, as well as by the Advisory Committee on Research (ACOR), advise the board. ACOR guides the board on issues relating to trends in Canadian science as well as informally providing a voice for CC's user community. The International Advisory Committee (IAC) provides the board with an international perspective on trends in science and best practices in the management of enterprises such as CC. Each of these bodies is governed by a terms-of-reference document, and the board chair, committee chairs and individual members are bound by code-of-ethics, conduct and confidentiality policies.

CC's board reviews its governance model and composition annually to ensure it adheres to best practices and adequately represents its stakeholders. This reflects CC's determination to build trust with the Canadian research community and with the many institutions that depend on it for access to advanced research computing resources.

Finally, CC has separated the board's responsibility for setting policy, overseeing management and planning strategy, from management's role, with executive authority delegated to the CEO, who has assembled and leads a team of competent functional experts. The management committee ensures that regional organizations and CC are aligned on all major initiatives.

The recently formed member council establishes an additional governance mechanism that increases CC's accountability to three specific classes of stakeholder: institutions owning and/or hosting CC systems ("host institutions"), institutions employing CC-associated personnel ("employers") and the research community at each institution. The Member Council has been directly involved in oversight of CC's Site Selection Process and will work to develop equitable funding models that can be implemented over the next few years.

A.8.3 | How We Addressed the Governance Conditions and Recommendations Expressed in the Initial Review Process

We have addressed all governance conditions and recommendations presented in the initial review process.

1. **Recommendation:** Incorporate CC and identify the appropriate parties to be "members" of the corporation.

   Response: CC was incorporated under the Canadian Not For Profit Corporation Act on Sept. 27, 2012, with potential members defined as institutions eligible to receive funding from the CFI.

2. **Recommendation:** Populate CC's board with independent directors who bring relevant skills to the board.

   Response: CC's board is composed of 11 directors with diverse skill sets and experience.

3. **Recommendation:** Form an International Advisory Panel. (This was also a commitment in CC's original proposal.)

   Response: An International Advisory Committee (IAC) was formed in late 2014. The IAC met for the first time on Feb. 21, 2015. Subsequent meetings are planned in conjunction with key international meetings (ISC, SC).

4. **Recommendation:** Revise CC's strategic plan to reflect the foundational role of ARC in Canada.

   Response: CC revised its strategic plan, supervised by the board. The document reflects extensive consultation with the research community across Canada. The board adopted this plan in July 2014. Although expressed at a high level, the strategic plan is a living document that must guide the organization through a time of rapid change in technology and the research environment in Canada and internationally. CC will revisit the plan at least every two years in light of this changing environment and the needs of CC's stakeholders

5. **Recommendation:** Strengthen the leadership of CC by hiring a chief executive officer (CEO).

   Response: The board has hired a president and chief executive officer, with authority to manage CC, including developing its budget, hiring and supervising key managers and implementing the strategic plan. A key component of the performance criteria for the CEO will be his success in achieving the objectives set out in the strategic plan.

A.9 | Management and Operations

CC's management and operations practices have advanced significantly since the outset of MSI funding in April 2012. Operational planning has evolved from regional (and even consortia within a region) to a more integrated process led nationally and executed collectively with the regional partners.

Areas of improvement include:

- Expanded coordination and planning mechanisms (national technical committees and working groups)
- Efficiency and organization of resource allocation processes (the "Resource Allocation Competition")
- Increasing integration and capability of performance management tools (CCDB, Perfsonar)
- Asset management tools (system databases, cost benefit analyses)
- Outreach and advocacy with federal and provincial stakeholders
- Development of demand forecasting processes and mechanisms (Sustainable Planning for Advanced Research Computing — SPARC)

Many of these activities have benefited from the simple ability of CC's national office to engage or hire specialists and contract for national licenses or hardware. CC's operating progress, as well as our plans for the final two years of the MSI project period, are more fully described in the attached Management and Operations Plan.

CC has accepted the recommendations from the original proposal review and taken steps that increasingly allow the organization to operate as a truly national facility. This was the intent of CC's original proposal, but adopting the review recommendations has yielded greater capability and expertise in the CC national office. This, in turn, has allowed CC to make more progress than might otherwise have been achieved in challenging areas, including developing a rational plan for reducing the number of CC systems, building consensus for that plan with stakeholders and developing a transparent and business-like process to select sites for a smaller number of new systems. The attached Management and Operations Plan documents this in great detail and will not be reproduced here. Instead, we summarize the direct response to initial review recommendations:

1. **Recommendation:** Revise CC's management structure to integrate board-led policy development with staff and line management within the organization, principally through the hiring, by the board, of a CEO with authority to manage the affairs of CC, the ability to articulate the vision and mission of the corporation and the ability to attract investment and create new opportunities for CC to meet its full potential.

   Response: CC's management structure has been revised as recommended, with the CEO hired by the board effective April 1, 2014. Since joining the organization, the CEO has made it a priority to clearly define the mission and vision of CC, to articulate these to diverse stakeholders and to enable new opportunities for the organization that are consistent with its strategy (e.g. collaboration with CANARIE, with federally funded research initiatives and with the private sector).

2. **Recommendation:** CEO to hire and directly supervise the Chief Science Officer (CSO), chief technology officer (CTO), chief financial officer (CFO) and communications officer, all reporting to the CEO and each with clearly defined roles and authorities.

   Response: A capable management team has been assembled, reporting to the CEO:

   a. CSO: Dr. Dugan O'Neil, seconded on a 50 per cent basis from Simon Fraser University, from January 2014 through March 2016. The CSO has responsibility for the scientific leadership and outputs, activities and vision for CC.

   b. CTO: Dr. Gregory Newby, hired starting January 2015. The CTO is responsible for all issues relating to acquisition, deployment, operation and decommissioning of systems.

   c. CFO: Paul Hart, CA, hired on a contract basis from May 2014 through March 2015. Going forward, CC will provide financial management capabilities through in-house personnel.

   d. Communications officer (Executive Director of External Affairs): Kelly Nolan, hired starting May 2014.

3. **Recommendation:** CEO should supervise the regional directors directly.

   Response: The regional organizations remain responsible for hiring and supervising the respective regional directors, who, in turn, manage regional staff. Regional directors report to the CEO informally, with the management committee providing a mechanism for integrated management. CC is working to establish inter-institutional agreements with all the regional organizations to improve accountability and transparency among these five organizations.

4. **Recommendation:** Translate strategic plan into management plan.

   Response: CC's initial management plan, attached, reflects the evolution of CC's role in the Canadian digital infrastructure ecosystem, the refinement of the relationship among CC, the regional organizations and the institutional staff, and the fact that our national efforts are setting precedents in Canada and in some cases internationally.

5. **Recommendation:** Conduct a system cost/benefit analysis and create a consolidation plan.

   Response: Both are complete. CC has identified 24 systems for defunding as part of the Stage 1 renewal (installation of new systems by early 2017), and a further 14 systems for defunding in Stage 2 (installation of new systems by late 2018). The 24 systems identified for defunding in Stage 1 represent the highest cost (per unit of usable capacity) systems, as well as the oldest systems, in the CC fleet; by 2017 they will be well past their useful lives. The Stage 2 renewal will enable the complete replacement of CC's entire current portfolio of sharable ARC systems.

6. **Recommendation:** Develop a cybersecurity plan, defining roles and responsibilities for the implementation and operation by the cybersecurity personnel for the life of the project.

   Response: A director of information security (Dr. Jonathan Ferland), reporting to the CTO, has been seconded by CC effective May 2014. Dr. Ferland has developed the CC cybersecurity strategy (see management plan) and is developing the associated policies and procedures. A key challenge with cybersecurity for CC is aligning security policies and procedures across a federated organization, whose staff are employed by 34 different institutions, each with slightly different, university-style security and IP regimes that generally promote and enable sharing rather than emphasizing the protection of "client"

information. At the same time, as part of CC's site selection process, CC has specified that hosts for new Stage 1 ARC systems must comply with the security and privacy policies now being developed; this compliance will enable enforcement of a higher degree of security and privacy protection. As CC's cybersecurity policies and procedures are more fully documented, CC will validate its approach through external review and consultation with experts.

7. **Recommendation:** Establish financial control and reporting processes, which should be modelled on, but not necessarily duplicate, the reporting and certification usually found in a public company.

   Response: Financial reporting and accountability have been significantly improved under the direction of CC's CFO, who brings public-company expertise to the organization. Since joining the organization in April 2014, the CFO has supervised closing the year-two financial reports and submitting them to CFI four months' earlier than in the prior year, finalized development and submission of year-three budgets two months earlier than in the prior year, and supervised a comprehensive CFI audit of both the MSI and National Platform Fund capital projects. CC financial management staff are developing best practice financial policies for the CC Corporation and the CC project. Establishing robust inter-institutional agreements among CC, the regional organizations and their respective partner institutions is critical to achieving the level of financial accountability desired for CC.

8. **Recommendation:** Nationally co-ordinated software acquisitions should be part of the role of the CTO.

   Response: The CTO is responsible for software licensing, acquisition and development. Given the distributed nature of the CC enterprise, as well as the wide range of disciplines it serves, CC has an extensive portfolio of commercial, open-source and researcher-developed analysis software, and a diverse range of systems software and middleware, installed on CC systems across the country. There are no nationwide licences in place — licences have been arranged on a local basis, and are often paid directly by research groups or institutions. The CTO will inventory existing software and licensing arrangements, examine opportunities to expand accessibility of licences while reducing their costs and establish mechanisms for licence consolidation and joint procurement. Detailed technical planning for the consolidation of systems in Stage 1 will address the "migration" of licenses from defunded systems to new systems, and this in turn will identify opportunities for efficiencies and/or opportunities for expanded availability of software resources.

9. **Recommendation:** Education, outreach and training must appear in the CC organizational structure.

   Response: Specific individuals within the CC management address education, outreach and training functions. National outreach is the responsibility of the executive director, external affairs; outreach initiatives are detailed in the communications section of the management plan. Education and training are supervised by the CSO, working through the Education and Training Committee, documented in *Appendix 6: Management and Operations Plan*, and discussed further in the management plan. CC's philosophy is that subject-matter experts are best qualified to deliver training on technical subjects, so CC's research support personnel allocate a portion of their time to this activity. The Education and Training Committee works to create a national "curriculum" of training topics and to align local efforts with that program.