



Compute Ontario:

A Blueprint for Action

2014-15 to 2019-20

## SECTION 1 – INTRODUCTION

Ontario is at a critical juncture in terms of its capacity to supply advanced computing at the level required for leading research and enabling industrial competitiveness in the global marketplace. Access to this fast changing infrastructure has become a key competitive differentiator for companies, research institutions (including health institutions), universities and governments. Advanced computing has become an essential tool on university campuses, akin to libraries for generations and the internet in more recent years. It is essential for research and more and more for product innovations for industries. Jurisdictions with which Ontario must compete have already left us far behind with respect to advanced computing capacity.

Compute Ontario is a newly formed not-for-profit corporation, designed to support advanced computing in Ontario. The organization is stakeholder driven. It arose from consultations with the research community and experts from industry, and the Board includes representatives from these backgrounds. Its early functions will include centralizing strategy and planning functions for advanced computing assets and investments in Ontario, including hardware, software, highly qualified personnel, data management, storage and protection, and connectivity. Universities, research institutions, companies and other stakeholders are coming together through Compute Ontario to address concerns about Ontario's ability to maintain its leadership position in advanced computing and to continue to be a player in the area globally.

Based on advice we have received, this Blueprint sets a course for action for Compute Ontario over the next 5 years to achieve the following 6 medium-term goals:

- To develop an investment strategy to secure sufficient and predictable funding
- To become a focal point for collaboration
- To develop and enable a strategy to support/expand Ontario's highly qualified personnel (HQPs)
- To facilitate a coordinated approach to storage and data access
- To expand the commercialization opportunities and to promote public awareness of the benefit of advanced computing, and
- To build our capacity in order to do all this.

The document contains high level action plans for realizing these goals. Our objective is to put Ontario in a position to utilize advanced computing to accelerate research and innovation to enhance its competitiveness.

### [What is Advanced Computing?](#)

*Advanced computing can be described as the use of high-end computing resources (computers, storage, data management, networking and visualization) to help solve highly complex problems in business, STEM (Science, Technology, Engineering and Mathematics), arts and humanities among others. Advanced computing typically involves computationally intense workloads (such as large and complex simulation models or extremely large data sets). Most importantly, it includes the highly trained scientific and technical people (HQPs) who support the systems that are critical for researchers.*

We have relied on significant prior work. The Ontario Ministries of Economic Development and Innovation and Training, Colleges and Universities established an Expert Advisory Panel, which released its report in 2012, to develop a vision for the next generation of research computing in Ontario. The panel included researchers, and industry experts. The development of a strategic framework to start up Compute Ontario followed in 2013, with the guidance of a Strategic Advisory Committee which brought forward the needs of stakeholders either supplying or using advanced computing. That document drew on expertise from within Ontario's universities, industry leaders, and researchers, all of whom provided valuable insights through consultations and focus group sessions.

Since its inception, Compute Ontario has been working with industry partners and research institutions considering the collective technology infrastructure needs of organizations wishing to invest in advanced computing assets, or wishing to be able to use these assets to further their business objectives. As we move forward, our approach will include working with Ontario entrepreneurs and businesses.

Support for advanced computing in Ontario needs to be multi-faceted. Advanced computing is a rapidly evolving ecosystem. In addition to investments in hardware, software, people and networks, advanced computing will require enhancements to storage and data management. As the move to open data accelerates, as health and research data grows exponentially, and as public and private "Internet of Everything" data explodes, so does the need to move quickly with sophisticated data management plans, curated to enable data access while protecting privacy and security.

Compute Ontario has already begun to establish partnerships and strategic alliances. Compute Ontario is seeking these arrangements in both the private and public sectors. At national tables, Compute Ontario is already providing regional representation for Ontario. Compute Ontario is working with Compute Canada on issues such as CFI's cyber-infrastructure initiative. It is also working with organizations preparing compute and data intensive FedDev proposals to identify shared resources already deployed and to work toward further collaboration.

Compute Ontario is also working with partners such as ORION and CANARIE with respect to their research networks and associated tools. Provincially, Compute Ontario will contribute to the Ontario government's work related to advanced computing, including its big data strategy.

#### Advanced Computing in Ontario today

The largest providers of advanced computing in Ontario are four consortia hosted by universities: SOSCIP (Western University and the University of Toronto, plus 9 more universities and growing, in partnership with IBM), SciNet (University of Toronto), SHARCNET (Western University plus many universities and colleges), and HPCVL (Queen's University plus a few universities and colleges). SOSCIP, the most recent, in 2013 invited Ontario businesses to submit research proposals. The latter three consortia have estimated that they support \$2.45 billion in leading edge research annually in Ontario, serving 21 universities, and a growing number of colleges and research hospitals. Additional advanced computing assets in Ontario are in the life sciences area and are hosted by the Ontario Institute for Cancer Research (OICR) and certain hospitals, with HPC for Health (which includes the University Health Network and the Hospital for Sick Children) having the largest capacity among these.

## SECTION 2: GOALS AND ACTION PLANS

Our vision is to drive advanced computing to accelerate research and enhance competitiveness in the global marketplace for a more prosperous Ontario. This Blueprint sets out our approach over the next five years.

Here are our guiding principles:

- Compute Ontario must be a trusted single voice for advanced computing in Ontario
- Our activities will be designed to promote the best interests of the advanced computing community (research institutions, industries, HQPs) to enable it to optimize services to researchers and businesses.
- We will work to deepen and broaden the use of advanced computing for research and industry applications including those in the health sector.
- We will encourage collaboration across the province to enable the best value for investment dollars, expertise and resources.
- We will undertake our work in a way that allows us to adapt quickly to changing circumstances and new opportunities.

Drawing from recent consultations as well as those used to develop the strategic framework. It is time to act on the advice provided. Over the next five years we plan to focus on achieving 6 goals by implementing their related action plans. One goal is to enable our organization to do its work. The others are intended to move advanced computing to the level needed to support leading research and to enable industrial competitiveness in the global marketplace. While moving ahead with our action plans, we will be refining our Blueprint as we develop our organization and further understand the needs of our stakeholders and alliances.

### *Goal 1- Develop and refine a balanced investment strategy for province-wide advanced computing infrastructure and secure predictable and sufficient funding*

Funding for advanced computing in Ontario is fragmented, provided through a number of sources including funding through competitive grants (CFI, FedDev, and ORF). Funding uncertainty discourages lifecycle planning and limits the ability to factor in the total cost of ownership. Ontario has held a leadership position in advanced computing historically but no longer – a significant funding injection is required to avoid adverse consequences to our international competitiveness.

**Organizations in the advanced computing area are not able to plan for the future, secure talent, execute multi-year projects or take advantage of early commercialization opportunities because of this uncertainty.** As well, a predictable funding stream is needed to address power utilization as electricity costs continue to rise while newer infrastructure becomes substantially more power efficient.

There has not been a systematic plan at a provincial level or ongoing funding stream, and until recently each consortium/institution planned and sought funding largely independent of one another. Planning at the provincial level will allow for the best value for investment dollars through use of shared resources. We need to look at opportunities to work in collaboration with private sector providers such as has taken place with IBM for SOSCIP.

As an initial step we have been working with the institutions applying for this round of FedDev funding to understand their requirements and to determine how the use of advanced computing components can be optimized.

However, much more needs to be done if Ontario is to have competitive infrastructure. For example, the systems used by our consortia providers are aging and we are falling generations behind. To illustrate, Ontario (and Canada for that matter) have yet to build a “petaFLOP class supercomputer”, and the current measure of leadership in supercomputing is now in the 10s-of-petaFLOPs. The US built its first petaFLOP class machine in 2008 and this machine was retired last year. There are currently over forty machines publically listed at this level of computing from 8 countries, including: the US (17), Japan (5), the UK (5), China (3), Germany (3) and France (3). Even if all of the current cyber-infrastructure “emergency” refresh spending is concentrated on two or three machines, it would take until 2016 to build a petaFLOP class system in Ontario. Although lists of machines like the above are most important to our “power users”, it is clear that Ontario researchers in general are constraining their research to the available computing capacity. As well Petaflop machines are increasingly important in industry. For example they are becoming critical in areas such as drug discovery and large scale device simulation (such as automobiles, energy systems including power grids, solar and wind).

Working with the advanced computing community, Compute Ontario needs to look at hardware refresh requirements from a total cost of ownership perspective. We need to minimize duplication, enhance networking, augment platform software and optimize inter-operability. We need to address storage requirements and data access. We need to see how we can leverage private sector investments. Perhaps hardest of all, we need to assess demand in a quickly expanding area where technology is rapidly changing.

It is important to note that not only are the needs of the current Ontario consortia’s users growing, but the user community is also expanding rapidly to include data-intensive research in all disciplines. The computing needs of this new data-intensive user base are expected to eclipse those of established users by many times. While historically the requirements have arisen from university based research, the universe has been enlarging to include health research institutions, business and potentially government. Therefore, the investment needs will inevitably be much higher than historical levels.

We need to consider how to maximize funding for advanced computing, as well as how to make the best use of available funds once in place. Compute Ontario will work with the advanced computing community to maximize access to federal and provincial funding. It will also explore other approaches to investment to supplement currently available funds such as employing partnerships and different funding models such as fee-for-service.

*Goal 2 - Establish a focal point for provincial collaboration across the advanced computing community that can also facilitate collaboration at the national and international levels.*

There is a need for collaboration: research institutions, researchers, HQPs who work in the area, governments and companies all have interests in advanced computing. While collaboration has now begun, generally each consortium or research institution has been planning primarily for its

own purposes, in part due to the way advanced computing is funded, on a project by project basis.

The environment is complex and fluid. Hardware platforms continue to evolve forcing changes in the research applications, programming models, and job mix management as well as the HQP skills required. The explosion of big data continues. The development of computing models, especially cloud, implies both technical change and new business models, such as changes in the management and governance of larger and larger shared services organizations.

Compute Ontario will only succeed if it is connected closely and credibly with its community. Stakeholder engagement will be a key activity over the next five years. The advanced computing community is diverse. Compute Ontario will engage stakeholders within its core group: the consortia leadership, and additional partner universities.

As well, Compute Ontario will work to engage other stakeholders such as research institutions and groups of clients of advanced computing as well as industry providers. Compute Ontario also recognizes the potential for advanced computing to benefit businesses through assisting start-ups and assisting established businesses so that they can meet their computation requirements. It will work with business incubators, accelerators, universities, community colleges and commercial firms in a number of areas. Compute Ontario also recognizes that new sectors need to be linked into the advanced computing community. Some future linkages are expected to include financial services, mining and information services, as well as a deepening of linkages with the health sector.

Compute Ontario also plans to move quickly to build critical partnerships and alliances, such as those with Compute Canada, the OCE, ORION, CANARIE and research institutions, beginning with leading health institutions.

It is important to note the value of turning vendors into partners. Research computing partnerships with IBM in Ontario have proven the benefits of these partnerships. Such benefits will be multiplied if we take what has been learned from our partnerships with IBM and apply it more broadly to the advanced computing ecosystem.

*Goal 3 – Develop and facilitate the implementation of a strategy to support the retention, development and expansion of the scientific and technical personnel (HQPs) supporting advanced computing.*

Currently the small number of HQPs in Ontario is a limiting factor to the utilization and growth of the advanced computing area. In part this is due to discontinuities in funding. However, there is no HQP strategy for training and development. Ontario has highly skilled and dedicated HQPs, many with long tenure, who provide programming support and design and maintain the systems. The capacity of this group is stretched—only 51 active HQPs were identified in Ontario’s academic advanced computing community. While the need for academic training for HQPs is recognized, and there are some specialized programs in place, they have limited capacity at this point. In addition, it is difficult for HQPs to access equipment for training purposes.

There are also highly qualified people who work with the researchers to assist them in making use of advanced computing resources—the “analysts”—usually with some form of computer

science related skills (e.g., programming, application development, bio-informatics). In addition, they have grounding in one or more of mathematics, physics, life sciences or computer sciences, often at the post-doctoral level. They are very valuable, can be hard to find and often move on after a few years. Businesses using advanced computing also require these HQPs.

For both types of HQPs, most are trained on-the-job, and it takes time to build up the required skill set. At present, HQPs are hired on short-term (1-3 year) contracts. They have little job security or opportunities for advancement. There are not enough personnel moving from university to industry. If Ontario develops HQPs, they often leave for other countries or giant corporations such as Google and Amazon where incomes are higher and job security is better. Many Ontario businesses do not seek to recruit for HQPs because they are not aware they will need them.

If Ontario wishes to continue to excel in this area, and if it wants to support the expansion of advanced computing in life sciences and other academic research areas and in industry, more HQPs are needed. As well, Ontario needs to find ways to retain the HQPs it already has and to make sure that it is making the most effective use of them. This is particularly important over the next few years until a human resources strategy can be put in place for recruitment, retention, training and development. An early activity for Compute Ontario will be to undertake a thorough assessment of the supply and utilization of HQPs and to make recommendations as to how the situation might be improved.

#### *Goal 4 – Facilitate a strategic and coordinated approach to storage and data access that incorporates curation, privacy and security*

Data management plans are now a requirement for research funding. Compute Ontario in partnership with Compute Canada and other stakeholders must develop an implementation plan to curate, preserve and provide access to research data to assist in broadening its use.

Attention to privacy and security is already recognized to be of paramount importance. We must coordinate data privacy and security provincially while implementing it institutionally and collaborating at a federal and international level so that unintended barriers to research are avoided.

Privacy is particularly important for health care and genomic data, where the individual can be identified. Complying with various governments' privacy legislation has created major challenges and costs for health research. New legislation is being introduced at the provincial level to protect privacy of personal health information for shared electronic records, and this will need to be taken into account in the development of new policies and standards.

Further, businesses undertaking research and developing innovation are critically concerned about controlling access to their intellectual property.

While protecting privacy and security, attention needs to be paid as well to how the data is best curated to facilitate access for research and innovation uses. This aspect is becoming increasingly important as expanding reliance is placed on big data to be shared across organizations.



From a security standpoint, business continuity and disaster recovery plans have concentrated to date on data protection but not alternative computing capacity. This could lead to major disruptions in research if a serious adverse event took place.

Compute Ontario is poised to provide leadership in privacy and security. It is demanding and time consuming to negotiate policies and standards for organizations, and for research projects to demonstrate how they will be adhered to, particularly when they entail more than one institution. We have heard from our community about the importance of obtaining funding to address issues of data sharing standards in particular. The situation will only become more complex as Ontario organizations become involved in international endeavours where such issues can take years to resolve.

*Goal 5 - Expand the commercialization opportunities available through advanced computing and promote industry and public awareness of the economic and social benefits from enhancing Ontario's advanced computing infrastructure.*

Many Ontario businesses do not understand how they can benefit from advanced computing; often they are not even aware of it. Developing bridges for businesses to meet and work with researchers in the advanced computing area could benefit all. The Ontario Centres of Excellence, universities, research institutions and business incubators are working to bring these communities together.

As advanced computing becomes a tool used by more companies across industries, Ontario companies will need to have this capacity to enable them to compete and lead within a global marketplace. To be a leader in their sector, most businesses must now engage in strategic development of advanced business processes using advanced computing capabilities.

Tools that were once only available to larger companies are now available to the individual/small business. Using advanced computing, small businesses can build rapid prototypes of their products and test them. There are new opportunities for entrepreneurs to take advantage of new technologies – some see it as a potential third industrial revolution— with personalized production on demand and web-based ordering. Advanced computing also enables sophisticated collaboration and innovation opportunities, in areas as wide-ranging as the health and screen industry sectors.

Where appropriate, Compute Ontario will work in partnership with OCE to explore utilizing promising new start-up technologies and software in advanced computing developed in Ontario. In particular, it will consider the needs and opportunities associated with creating software to serve research priorities.

Ontario's future prosperity is critically dependent on increased investment in and effective utilization of advanced computing technologies in both the public and private sectors. Advances in research and commercialization provide opportunities in employment, education, and business for those living in Ontario.

Ontario needs to share these stories across the province to enable enhanced public understanding of potential opportunities. Because of its technical complexity, advanced



computing can pose marketing challenges, particularly when it comes to attracting the sustained media interest necessary to educate the public about its benefits

Making the general public aware of successes to date and how advanced computing can be used to improve health outcomes and other aspects of quality of life will deepen the appreciation of opportunities for its use, and may encourage more people to support its activities. It also assists in setting the stage to enable government funding.

#### *Enablement Goal - Develop Compute Ontario Capacity*

In order to perform the Blueprint activities detailed above, Compute Ontario must be resourced. We are incorporated as an organization and have our Board in place. We are already making some progress through the use of loaned resources working part-time. But to address the issues people see as critical we need to rapidly accelerate our activities. To do so, we need to be able to hire key resources to undertake our work. This requires ongoing funding, initially from government. Through our activities, we will endeavor to maximize value for money through resource and knowledge sharing and through the use of fair and transparent processes.

### Section 3: Moving Forward

Over the coming months, we will begin to implement the Blueprint while consulting with our community as to how to make it even stronger. We will work closely with the providers and users of advanced computing in Ontario to make this plan a reality.

## ABOUT COMPUTE ONTARIO

Compute Ontario was incorporated as a not-for-profit corporation in the spring of 2014. Our vision is to:

*“Drive advanced computing to accelerate research and enhance competitiveness in the global marketplace for a more prosperous Ontario.”*

Our operations are guided by these core values:

*Trusted* – We will strive to be trusted by researchers, industry and institutions. We will strive to preserve integrity in research. We will be regarded as vendor neutral. We will provide credible advice and information that is accurate, timely and relevant. We will protect privacy. We will create a secure and predictable environment to foster research in Ontario and to support Ontario’s business.

*Collaborative* - Compute Ontario will work to bring partners and stakeholders together, to work towards common goals, enhancing the role of advanced computing for all partners. We exist to serve a number of stakeholder groups. We will strive to provide a place for research and industry to interact. We will serve as a resource for research and for industries in Ontario. We will work to form partnerships and alliances within the advanced computing community. Compute Ontario can only succeed with its partners.

*Accountable* - Compute Ontario will seek to provide value for money to Ontario residents. Our decisions will be transparent, and clear, and we will take responsibility for our actions.

*Innovative* – Compute Ontario will be able to respond to the rapidly changing nature of advanced computing. We will be agile, and able to adapt as needed to provide quality outputs in a streamlined fashion.

*Excellence* - We will promote excellence in research in Ontario and excellence in commercialization of products by Ontario’s companies through the use of advanced computing. We will strive for excellence in our own activities and will promote Ontario as a global leader in advanced computing.