

**Digital Research Infrastructure for Subatomic Physics in Canada**  
*Institute of Particle Physics (IPP)<sup>1</sup> and Canadian Institute of Nuclear Physics (CINP)<sup>2</sup>*  
December 12 2020

Innovation, Science and Economic Development Canada (ISED) has announced a restructuring of the digital research infrastructure in Canada. A new not-for-profit organization, New Digital Research Infrastructure Organization (NDRIO), has been established to advance the Digital Research Infrastructure (DRI) ecosystem for Canadian researchers. The purpose of this Position Paper is to provide NDRIO with input from the Canadian subatomic physics community with our requirements and vision for the future of our digital research infrastructure. The Position Paper was written by a task force<sup>3</sup> appointed by the Institute of Particle Physics (IPP) and the Canadian Institute of Nuclear Physics (CINP), and was circulated for feedback to the entire community.

Subatomic physics includes the fields of nuclear and particle physics with the common goal of understanding the fundamental nature of the universe. We estimate the size of our community, represented by IPP and CINP, to be over 750 faculty, staff, research associates and graduate students in universities across the country. We are involved in a wide range of international projects as well as projects within Canada at the TRIUMF Laboratory (Vancouver) and SNOLAB Underground Science Laboratory (Sudbury). Our community includes researchers in the Perimeter Institute and Arthur B. McDonald Astroparticle Physics Research Institute.

Computing is an integral part of our research and our use ranges from medium to large scale computing; storage of petabyte-scale data samples; and cloud and high-performance computing. High-speed networks are an essential part of the infrastructure as they enable us to transfer large volumes of data around the world. Today our community uses many thousands of compute cores and many tens of petabytes of data storage. Our software manages complex distributed heterogeneous systems of compute and storage resources. Further, the computing resources are a critical element for the training of highly qualified personnel (HQP). Our graduate students and research associates become skilled in a wide range of areas from data analysis, writing complex software, or using large high-performance computers or compute clouds. They find employment in research and in a wide range of industries from medical physics to financial modelling.

The international nature of our research requires that our resources integrate with the global infrastructure. Our connections to our institutions, Compute Canada sites and our international partners is provided by CANARIE, Canada's National Research and Education Network (NREN) organization. In addition to the national research network backbone across Canada, CANARIE has established two private networks for the subatomic physics community. One private network has a dedicated connection between the TRIUMF Laboratory in Vancouver and the CERN Laboratory in Geneva. The second network is a private routed network that connects the Compute Canada sites with our universities and laboratories to the global private routed network for particle physics. Our community is the largest user of the CANARIE network.

The subatomic physics community secures its computing through allocations of resources at shared facilities from Compute Canada and the regional computing consortia. In addition, the ATLAS experiment at the Large Hadron Collider (LHC) located at the CERN Laboratory in Geneva has a set of centres around the world that are designed to store and process the raw data in near-real time from the experiment. Canada has such a facility, called a Tier-1 Centre, that has been funded by CFI and is operated by TRIUMF and SFU. The ATLAS Tier-1 has been operating in a 24x7 mode for many years with an excellent track record. In addition, a similar facility for storing

---

<sup>1</sup> Institute of Particle Physics (Canada), Director J.M.Roney ([www.particlephysics.ca](http://www.particlephysics.ca))

<sup>2</sup> Canadian Institute of Nuclear Physics, Executive Director G.M. Huber ([www.cinp.ca](http://www.cinp.ca))

<sup>3</sup> This document was prepared by a joint IPP/CINP task force: Alexandros Gezerlis, Doug Gingrich, Garth Huber, Chris Jillings, Randy Lewis, J.Michael Roney, Pekka Sinervo, Randall Sobie (Chair) and Reda Tafirout.

and processing the raw data was recently funded by CFI for the Belle II experiment at the KEK in Tsukuba, Japan. The facility will be located at the University of Victoria and will start operation in 2021-2022 as a “contributed facility” to Compute Canada and NDRIO.

It is important to highlight the international nature of our projects. Our collaborators have access to the Compute Canada resources as Canadians have access to those of our partners. *It is critical that NDRIO continue to provide access to our international collaborators, irrespective of their country.*

Our requirements will continue to grow over this decade and beyond. Many of our existing experiments (such as ATLAS) are undergoing significant upgrades that will increase the volume of data by an order of magnitude. A number of smaller experiments (50-1000 Canadian and international researchers) are coming online now or later this decade. Experiments at TRIUMF and SNOLAB are increasing their data samples and their computing requirements are becoming significant. The theoretical researchers in our community use the large high-performance computing (MPI-type) systems for detailed calculations of hadronic or nuclear structures, lattice gauge theories and nuclear astrophysics.

The subatomic physics community is committed to the success of NDRIO as it is a critical infrastructure that will have a significant impact on our international commitments and reputation. We helped develop research computing in Canada, and participated in the formation of the regional consortia and Compute Canada. We have worked with CANARIE to establish a world class research network in Canada and links to the world. Further, we have worked with industry to study novel and innovative hardware and software technologies (including cloud computing). We believe the expertise in our community can be leveraged by NDRIO to build a national DRI infrastructure.

The NDRIO call for position papers requested our perspective on the current challenges accessing DRI tools, services and support; the future state of DRI in Canada; and how NDRIO could achieve such a state. We address each topic below.

### ***Current issues on DRI in Canada***

The subatomic physics community has been a large user of computing for many decades. We use the high-throughput (general purpose) and high-performance computing facilities, and the storage resources of Compute Canada. Our projects span years, and more often, decades. One of our concerns is that Compute Canada has been funded in an *ad hoc* manner by CFI (and provincial funding organizations). This makes it difficult to make long-term international commitments to our projects. A secure and continuing source of funding is critical to our success and a stable future of NDRIO.

Our community receives essential support from the Compute Canada personnel. Compute Canada has established a Subatomic Physics National Team (SPNT) that is responsible for the deployment, testing and support of systems, software and middleware. The SPNT personnel are supported by Compute Canada and local host universities. The SPNT meets regularly with the technical teams of the experimental projects and helps manage central services that are shared by the community. *It is essential that the SPNT be maintained for the ongoing projects, and the new projects under preparation.*

Our researchers have faced challenges using the Compute Canada resources. *Of utmost priority, is that NDRIO operate its resources on a 24x7 basis with well-established operating procedures.* No resource or centre should be offline for any length of time. This requires a technical design that ensures redundancy and the ability to take selected hardware components offline without impacting operations. A 24x7 service requires that system administration teams be available at any time. Currently, Compute Canada operates on a 9x5 basis and issues are addressed during working hours. Frequently sites are taken down with little warning or are down for extended

periods. Our research environment is very competitive and other computing facilities around the world operate on a 24x7 basis. *NDRIO needs to ensure that the facilities are designed and managed to be fully operational around the clock, and to ensure that scheduled downtimes do not result in loss of access to resources for any user.*

In addition, we find that resource allocations for projects are often unable to be fulfilled. We recognize that allocating resources is challenging but it appears that the resource allocation committees do not have a full understanding of existing and planned resources. As an example, we have received allocations on sites that do not allow the worker nodes (or processors) to connect to remote systems. This is a critical requirement as our resources are integrated into a global infrastructure. *NDRIO needs to improve the communication between the operations teams and the resource allocation committee to ensure that allocations can be effectively used.*

We have concerns with management of the Compute Canada systems and personnel. There are system administration and application support teams dispersed across the country. In most cases, the personnel answer to their university IT group. For example, a problem identified by the SPNT is often ignored (or put on low priority) by the system administration team of a host site. One of the challenges is that the site personnel are often employees of their local institution or university, and not Compute Canada, which leads to confusion about responsibilities. *It is important NDRIO technical personnel answer to the NDRIO management.*

Compute Canada previously employed a Scientific Officer, who was a well-respected researcher seconded from their responsibilities. The Scientific Officer was the only researcher on the Compute Canada staff and was able to understand the requirements of the research community, and convey the knowledge to the Compute Canada management. An NDRIO Scientific Officer could, for example, be the chair of the Researcher Council. *We strongly recommend that NDRIO maintain the position of Scientific Officer.*

### ***Future DRI state***

Our vision is that NDRIO be a national entity of compute, storage, software, networks and personnel resources providing 24x7 services to the research community.

Currently, Compute Canada sites operate as independent organizations with staff being employees of the local site. Each site has different rules and operating conditions. There is minimal interoperability and interchangeability between infrastructure elements of the sites. We believe NDRIO must unify its infrastructure and personnel to succeed, and this may require a significant reorganization so that NDRIO can avoid being a rebranded version of Compute Canada.

Our community is concerned that the proposed NDRIO model isolates computing, storage and software. Our computing systems must be tightly integrated to high-performance storage systems so that they can quickly deliver data to the compute nodes for processing. Our software finds the optimal location to run the application based on compute, storage and network metrics. Much of our software and storage systems are part of a global infrastructure, and we are constrained to use systems developed by the international collaborations.

We believe that NDRIO would be more effective and easier to manage if hardware systems and personnel were structured by computing requirements. For example, NDRIO could have sub-groups (each with dedicated hardware and personnel) for data curation, archive and preservation; high-performance (MPI-type) computing; high-throughput and high input/output data computing systems; and general-purpose systems. This would allow researchers to engage with NDRIO teams with specialized expertise. We believe this would result in systems that are more manageable, both technically and administratively, and will be more efficient and cost-effective to operate.

Compute Canada has a staff of system administrators and application specialists. The SPNT team is an example of a successful group of application specialists who integrate the complex systems of the particle and nuclear physics experiments into the Compute Canada infrastructure. In addition to the SPNT, we believe that NDRIO should also support software developers either by funding such personnel to work in the research groups or within NDRIO. Previously, CANARIE funded software developers though their focus had drifted from supporting world-class research to providing software systems that they felt might be of value to as many researchers as possible. *Research impact must have the highest priority in allocating personnel resources.*

Our community requires specialized computing facilities, such as the ATLAS Tier-1 Computing Centre at TRIUMF/SFU that stores the raw data from the ATLAS experiment in Geneva in near-real time. There are plans for similar specialized facilities in the coming years. Our understanding is that these specialized facilities will continue to be funded by CFI or other agencies, be considered as “contributed” resources and be housed in NDRIO (Compute Canada) data centres.

*Data archiving and preservation must be an integral part of NDRIO’s mandate.* This is important to almost every field of research and has not been part of Compute Canada’s mission.

It is critical that NDRIO adopt best practices when managing its resources and interacting with its user community. Establish lines of responsibilities between the site personnel, domain specific teams and the management with a working and transparent ticketing system that interacts effectively with the users. Ensure redundancy in expertise and skill-sets by cross-training existing staff to avoid any ‘single-point-of-failure’ associated with planned or unplanned leaves. Use a coherent set of rules and robust ticket procedures across all sites to address the issues raised by the research community.

NDRIO must demonstrate its value to the Canadian research community. NDRIO should provide the key performance indicators for each site that are reviewed by the Researcher Council and these should be made publicly accessible. For example, a quarterly report provided by the new organization should include the CPU utilization of a site, a list and reason for the downtime of each site, and percentage of resource allocation actually used by users compared to those allocated.

### ***How to bridge the gap***

NDRIO needs to improve the existing resource allocation process and ensure that the Resource Allocation Committee (RAC) members are knowledgeable about the type and availability of resources. Currently resources are often allocated on incompatible systems, or systems that were heavily oversubscribed, because of a lack of information provided to the RAC members. The RAC should also consider hardware decommissioning schedules, which is particularly important to maximize the stability through the lifespan of projects with a long life-cycle. Many users take advantage of the default allocation of resources provided by Compute Canada without having to apply for a resource allocation. This is an excellent way to provide modest levels of resources to many in our community and we recommend that NDRIO continue this service.

The NDRIO-RAC should work closely and collaboratively with CFI and NSERC. RAC awards should continue to be contingent on the proponents holding a peer-reviewed grant from the Tri-Councils, CFI or recognized funding organization.

We recommend NDRIO, at management and operational levels, consult with international organizations that run large computer systems such as National Energy Research Scientific Computing Center (NERSC) in the USA and CERN in Europe. These facilities have successfully deployed large systems with well-defined operational procedures. A review of the best practises at other sites will help NDRIO manage its resources and personnel. It

has been our observation that the management, operations and organizational aspects have been a weak point of Compute Canada and we encourage NDRIO to improve this area.

It is important that NDRIO investigate new technologies and build collaborations with experts in the research teams. Compute Canada, CANARIE and our community have participated in global network demonstration projects. Our community helped CANARIE test its new 100 gigabyte/second transatlantic link. We anticipate 1 terabit/second links will become the norm in this decade and new facilities should be designed accordingly.

NDRIO should explore the use of commercial cloud computing resources that could, for example, meet peak demands that exceed existing resources. The commercial cloud systems are extremely reliable and their technical teams are highly responsive to their customers' requirements. Members of our community have demonstrated that commercial clouds can be effectively utilized. The Canadian research community would benefit from improved ties with the commercial cloud computing industry.

NDRIO should support the development of new research software and data management tools. No organization can predict the software and services that will be adopted by the researchers, and NDRIO should invest in many pilot projects rather than trying to select a system or service for everyone. Funds in these areas should be focused on our highest priority projects that have significant international impact based on peer review.

### ***Summary***

NDRIO's mission is to enable Canadian science to be competitive on the world stage. The success of the Canadian particle and nuclear physics community is critically dependent on NDRIO. This document has described the impact of digital research infrastructure on our fields. We have highlighted the critical infrastructures and personnel requirements, and have made practical recommendations that will help NDRIO evolve into a world-class DRI resource for Canadian research.

### ***We summarize the key points of our position paper:***

*Of utmost priority, is that NDRIO operate its resources on a 24x7 basis with well-established operating procedures.*

*It is critical that NDRIO continue to provide access to our international collaborators, irrespective of their country.*

*It is essential that the Compute Canada Subatomic Physics National Team (SPNT) be maintained.*

*NDRIO needs to improve the communication between the operations teams and the resource allocation committee to ensure that allocations can be effectively used.*

*It is important NDRIO technical personnel answer to the NDRIO management.*

*We strongly recommend that NDRIO maintain the position of Scientific Officer.*

*Research impact must have the highest priority in allocating personnel resources.*

*Data archiving and preservation must be an integral part of NDRIO's mandate.*