Role of TRIUMF within the Digital Research Infrastructure Ecosystem[1]

December 14 2020

Digital research infrastructure in Canada in undergoing a national restructuring with the recent establishment of the New Digital Research Infrastructure Organization (NDRIO) by the Ministry of Innovation, Science and Economic Development Canada. The new organization is in the process of replacing Compute Canada with a new governance model, and going through a national consultation process to gather information from the research community about the state of digital research infrastructure (DRI) and assess needs.

TRIUMF is Canada's particle accelerator centre, and one of the world's leading laboratories for particle and nuclear physics and accelerator-based science. We are an international centre for discovery and innovation, advancing fundamental, applied, and interdisciplinary research for science, medicine, and business. At TRIUMF, we're passionate about accelerating discovery and innovation to improve lives and build a better world. Equity, diversity, and inclusion are integral to excellence and enhance our ability to create knowledge and opportunity for all. TRIUMF strategic plan is publicly available online[2].

The subatomic physics community, represented by the Institute of Particle Physics (IPP) and the Canadian Institute of Nuclear Physics (CINP), is making a White Paper submission to NDRIO. Therefore, the information provided in this document is to be considered complementary and expands further on the role of TRIUMF within the DRI ecosystem.

Traditionally, and at both the national and international levels, TRIUMF has been a key player in all facets of the DRI ecosystem, namely advanced research computing, data management, research software and networking. TRIUMF's international network not only allows data to flow, but also people and ideas, connecting Canada's DRI ecosystem to the broader world, building Canada's global reputation for innovation and excellence. TRIUMF has been at the forefront of technological advances in all areas, and a driver for innovation. Notable areas being advanced research networks and large-scale distributed computing. TRIUMF has been the largest user of CANARIE's bandwidth, and has hosted at one point the largest data storage facility in academia. TRIUMF researchers have also been closely involved in the establishment of a national DRI strategy prior to the inception of Compute Canada. Several solutions adopted for instance by Compute Canada today and which are to the benefits of all Canadian researchers from all fields, have initially been deployed at TRIUMF and for which TRIUMF has played a significant role.

To maximize its scientific output, the TRIUMF community has access to both dedicated resources and shared resources. Access to shared resources, outside of the purview of TRIUMF, have been through formal resources allocations from Compute Canada. The TRIUMF community has also used Compute Canada default allocations via opportunistic use which have been valuable to a number of researchers, including students and postdocs.

During the past several years, TRUMF has deployed several clusters and solutions to meet the needs of

1 On behalf of TRIUMF: Jonathan A. Bagger, Director; Reiner Kruecken, Deputy Director, Research; and Reda Tafirout, Head, Scientific Computing (corresponding contact via tafirout@triumf.ca).

2 https://fiveyearplan.triumf.ca/about-triumf/about-our-five-year-plan

both large international projects such as ATLAS and T2K experiments with large-scale and dedicated resources funded by CFI, but also for the on-site experimental program, as well as the local theory community. The ATLAS experiment at CERN's Large Hadron Collider has led to the 2012 Nobel Prize in physics for the Higgs boson discovery, and the 2013 Breakthrough Prize in Fundamental Physics. The T2K neutrino experiment in Japan has led to the 2016 Breakthrough Prize in Fundamental Physics. TRIUMF has also been involved in research software solutions that are benefiting not only the TRIUMF community, but also the entire Canadian research communities that are using Compute Canada resources. More recently, new initiatives with respect to machine learning and quantum information systems applications have gained significant traction, and a number of scientific projects are already benefiting from these advances.

A) Advanced Research Computing & Data Management

ATLAS Tier-1 centre

For more than a decade, TRIUMF has been hosting and operating the Canadian Tier-1 centre, one of ten centres worldwide as part of the Worldwide LHC Computing Grid (WLCG). The centre provides large-scale and dedicated resources necessary for the storage of the raw and secondary data sets, as well as computing capacity for data processing, simulation, and physics group activities. The Tier-1 centre is fully integrated within the overall ATLAS distributed computing operations worldwide and provides essential services to the entire collaboration around the clock. The Tier-1 systems and associated services are maintained in a highly secure environment and designed for high resiliency and performance. More recently, since 2018, the centre has been relocated to SFU and co-located with Cedar in response to a national policy to consolidate high performance computing in a limited number of sites for efficiency and economy of scale. A solid base of expertise has been assembled while deploying large-scale storage systems, databases, complex clusters and network topology, and grid computing technologies. The Tier-1 centre's dedicated personnel (close to 10 FTEs) have been key players in the initial prototyping of baseline WLCG services, as well as their deployment and evolution over many years. The personnel also play a key role within ATLAS distributed computing workflow and data management systems. The non-ATLAS Canadian research community has benefited from and leveraged the extensive knowledge and expertise from the Tier-1 centre while adopting some of its solutions. Currently deployed resources consist of 12,000 cores, 11 PB usable disk and 30 PB of tape capacity. Over the years, the average annual investment has been close to $1.5M for the equipment alone. In order to meet future needs of the LHC program, the Tier-1 centre capacity has to grow at a projected rate of 15-20% per year.

T2K Tier-1 centre

TRIUMF has been hosting a Tier-1 centre for the T2K experiment; a medium-scale storage-only facility with 700 TB capacity. There are only two Tier-1 centres for the experiment, the other site being hosted at RAL in the U.K. Computing needs for simulation and analysis have been provided by Compute Canada. The Tier-1 centre at TRIUMF is now being de-commissioned and the main T2K computing efforts are now using both storage and cpu resources at Cedar via Compute Canada. TRIUMF provided significant resources, including code repositories, various Plone websites, database servers, and electronics logbooks. Most of the services were deployed on a virtualization infrastructure. The local TRIUMF group also has access to modest cpu and storage resources (32 cores and 50 TB) used primarily by students for data reduction and analysis. Access to a large-memory footprint machine has proven valuable while performing complex statistical analysis.

Theory cluster

TRIUMF theory group have access to a dedicated medium-scale cluster (320 cores) that is hosted at the UBC ARC data centre since 2017. The nodes have the same architecture as the ones deployed at Cedar with large memory and fast interconnect. The cluster is used by the entire theory department, including students and postdocs. Overall, the cluster is more heavily used for nuclear physics projects, although there is particle physics and quantum computing work done as well. Typically, codes are parallelized on the nodes. However, MPI calculations are also used for multiple nodes occasionally. TRIUMF is contemplating an expansion to double it in size within the next five years. For large-scale massively parallel computation (>1000 cores), the theory group is using Compute Canada resources as well as resources in the US such as OLCF at ORNL through an INCITE award.

ATLAS Tier-3 cluster

Since the beginning of the ATLAS research program data taking, TRIUMF deployed a small-scale cluster (90 cores and 60 TB). The cluster is primarily used for analysis code development, along with a user interface for job submission to large-scale grid resources and to perform final stage statistical analysis. Derived data sets from various physics groups are replicated locally using distributed computing tools. The cluster has been used as training ground for students.

On-site experiments

TRIUMF also provides resources and support to several experiments that are located on-site as part of the Isotope Separator and Accelerator ((ISAC) facility; the great majority require modest storage resources. Overall, they require a few dozen TBs for each year of operation. The experiments also receive critical support from TRIUMF with respect to data acquisition infrastructure and data model designs.

FLUKA simulations

The Advanced Rare Isotope Laboratory (ARIEL) is a major endeavour at TRIUMF and requires extensive cpu resources to conduct major simulation campaigns using the FLUKA package. The simulations are needed to perform dose calculations and validate shielding. Currently, a medium-sized cluster is being used at TRIUMF. The cluster is also used for the proton and neutron irradiation facilities (PIF/NIF) studies which are used for materials characterization. The TRIUMF Life Sciences program also makes uses of these resources for FLASH studies as it recently joined the Optimization of Medical Accelerators (OMA) consortium. The overall simulation efforts are currently using 360 cores.

Machine Learning platform

TRIUMF is also operating a dedicated machine learning platform (with 32 cores and 8 GPUs) that is currently hosted at UBC ARC data centre and was recently moved from TRIUMF. The platform has been successfully used by several projects and strong support is provided by a recently hired data scientist who also provides intellectual contributions to the projects, while working closely with students and postdocs. The goal is to enhance the scientific output of projects by implementing innovative techniques.

B) Research software

TRIUMF has been involved in various software projects that are benefiting the wider community. For many years, TRIUMF has been collaborating on international software projects such as GEANT4, which is an essential tool and widely used in designing particle detectors and to do simulations needed for data analysis and extraction of physics results. GEANT4 is also used in medical physics applications such as for PET detectors. TRIUMF has also developed software for simulations of

accelerator beams dynamics (ACCSIM) that has been used by international laboratories around the world. TRIUMF also developed the PHYSICA package which is a high level, interactive programming environment with user friendly graphics and sophisticated mathematical analysis capabilities.

TRIUMF has also been a key developer and collaborator of the Maximum Integrated Data Acquisition System (MIDAS), which is extensively used by both the TRIUMF on-site experiments and also other medium-size experiments around the world. The data acquisition system is a key component in the design of an experiment and TRIUMF provides strong support, including initial design consideration of the data model to be adopted which is intimately related to the amount of computing resources that would be needed for data storage and processing during the lifetime of an experiment.

Regarding ATLAS distributed computing operations, TRIUMF personnel have developed extensive tools and a robust software framework to help thousands of users use and access efficiently both local and distributed computing resources worldwide. The framework has helped tremendously the overall output of the ATLAS research program and is also used to run data processing and simulations tasks on the grid. The framework is also able to integrate the use of containers technology.

The Canadian Tier-1 Data Centre is the only ATLAS Tier-1 site that has developed an in-house software to manage its tape storage system because of limited functionalities of commercial solutions. The system was developed from the beginning of the Tier-1 prototyping phase and perfected during the course of the last decade, while being used and stressed under various operating conditions. The Canadian Tier-1 system has one of the best efficiencies because tape families are used to store different data types onto different tapes, as well as file grouping and packing, i.e., files that belong to the same data set are written on the same tape because when files are read, the entire dataset is recalled. This greatly improves I/O and minimizes tape mounts, and the overall system attains very high performance.

C) Networking

TRIUMF has been a driver in the deployment and evolution of high-speed networks in Canada, and has been by far the largest user of CANARIE in Canada for several years, largely dominated by the ATLAS Tier-1 centre operations as part of WLCG. Excellent synergy and interplay exist with CANARIE, BCNET, HEPNET and Compute Canada in order to continue to drive network innovation in Canada with sate-of-the-art topology. In 2002, TRIUMF was involved in the first high-speed transatlantic data transfers tests at record speeds, using end-to-end “light path” connections and was one of the determining factors for hosting the Canadian Tier-1 centre.

TRIUMF was also an early adopter of the performance Service-Oriented Network monitoring A**r**chitecture (perfSONAR) suite; a network measurement toolkit designed to provide federated coverage of paths and help to establish end-to-end usage expectations. The PerfSONAR suite does regular latency and bandwidth measurements using thousands of mesh connections. There are ongoing projects to try to implement machine learning techniques to produce decisions that could be used in the ATLAS distributed computing workload and data management systems, which would ease operational burden and minimize failures.

D) Future outlook on DRI in Canada

There are extensive developments and concerted efforts within the LHC computing grid which are under constant evolution in order to meet stringent demands from the experiments and users, and

challenges are ever increasing. WLCG utilizes state-of-the-art workload and data management systems as well as extensive analytics tools. Some of the technologies adopted by the LHC computing grid have already been deployed at Compute Canada. For example, the CERN Virtual Machine Filesystem, used to dynamically distribute primarily experimental software and grid tools onto the WLCG infrastructure, is now used at all Compute Canada sites to distribute not only research software but also genomics datasets. Continuing to leverage these developments going forward would be of great benefit to NDRIO. A notable ongoing project concerns Data Organization, Management and Access (DOMA), with a focus on the medium and long term evolution. This includes the concept of a data lake, intended as a storage service geographically distributed across large data centers connected by fast network with low latency which would be applicable to NDRIO's infrastructure.

As part of its strategic plan, TRIUMF considers it imperative to exploit advances in data science and computing which are aligned with a national DRI strategy. Scientific computing is a critical component of much of the work that takes place at TRIUMF and our partner facilities around the world, particularly in high-energy physics experiments. In recent years, the amount of data produced has increased exponentially at research facilities and private businesses alike. Rapid advances in large-scale computing, big data, machine learning, and quantum computing are starting to have serious implications for how we do our work; it is imperative that we remain at the forefront of this fast-changing field. Using our experience in advanced high-performance computing and the application of machine learning algorithms gained through our work on various projects, we can play a critical role in advancing data science capabilities and techniques across disciplines, while at the same time teasing out new discoveries. Recent advances in theoretical methods and in computing power, in particular supercomputers, have enabled nuclear physicists to confront theory with data and calculate, from first principles properties of heavy nuclei once thought impossible to reach. This has enriched both theory and experiment at facilities such as ISAC and ARIEL.

The emergence of usable quantum computers has the potential to take these calculations to a new level with quantum computer-based nuclear science. This is but one example of how a deeper exposure to quantum computing could enrich our science program, given the extensive national and international networks we have established. TRIUMF has signed a cooperative agreement with Germany’s Helmholtz Association, along with Canadian companies D-Wave Systems Inc. and 1QBit, to jointly establish Canadian and German quantum computing and machine learning networks. Our aim is to facilitate national and international collaboration in the use of applied quantum computing and machine learning tools to enhance research across our fields of endeavour. The Canadian quantum computing network is one of the initiatives related to our involvement in Canada’s Digital Technology Supercluster. TRIUMF is considering, based on requests from the Canadian quantum computing community, developing a quantum computing user facility, based on ion trapping and would welcome collaboration with NDRIO regarding eventual access and development of a software services platform. Overall, our aim is to facilitate national and international collaboration in the use of applied quantum computing and machine learning tools to enhance research across our fields of endeavour. It is important for NDRIO to be part of that evolution.

TRIUMF is looking forward to collaborate and develop further synergies with NDRIO once firmly established and to continue to build on existing collaboration with Compute Canada. Leveraging both existing Compute Canada and TRIUMF extensive experience and expertise would add significant value to the national DRI strategy, and would be to the benefit of all researchers, beyond just the TRIUMF research community. A state-of-the art DRI infrastructure in Canada will increase competitiveness and international standing as well as attract and retain world-class talents.