2016 Addendum

Computing requirements for the Canadian subatomic physics community

Institute of Particle Physics of Canada and the Canadian Institute of Nuclear Physics
February 25 2016

This document describes the computing requirements for the Canadian subatomic physics (SAP) community for 2016-2021. In 2014, the SAP community submitted a white paper to Compute Canada; this document gives an update to the previous submission.

The main changes are found in the computing and manpower requirements of the ATLAS experiment and new input from the IceCube experiment (we provide a short description of IceCube at the end of this section). Minor changes were made to the resources requirements of the Belle II and SNOLAB experiments. In addition, there are updates and new additions from the SAP theory community.

The previous document focuses on the computing and storage requirements of our community as requested by Compute Canada. For many projects, this is sufficient for meeting their demands. However, manpower is also a key element in our computing requirements for the projects with international commitments.

ATLAS, for example, requires manpower from Compute Canada in order to integrate their resources into globally distributed and complex computing framework. International agreements require us to deploy and support services and portals on a 24/7 basis. Such is the situation for the ATLAS Tier-1 centre at TRIUMF, which will be migrated to Compute Canada in 2016-2017. The ATLAS Tier-2 centres have less stringent, but nevertheless, important manpower requirements. Other international experiments, such as T2K and BelleII, utilize the WLCG infrastructure at the Compute Canada sites.

Currently in Canada, a dozen FTE's are assigned to the ATLAS project for efficient operations of the Tier-2 centres and the dedicated Tier-1 (supported by Compute Canada and TRIUMF). It is essential that the Compute Canada personnel are fully engaged with our community and that adequate budgeting be put in place to ensure proper personnel skills development. Equally important, the personnel should attend periodic national and international meetings and workshops in order to keep up to date and develop synergies with their peers.

Compute Canada and the SAP community have established a world-class expertise in deploying large-scale storage systems, databases, complex clusters and network topologies, as well as Grid and cloud computing technologies. We strongly encourage Compute Canada to continue investing in such personnel so that the Canadian SAP community and others can play a leading role in discovery science.

IceCube Neutrino Observatory

The IceCube Neutrino Observatory is the world's largest neutrino detector conceived to detect high-energy neutrinos, up to the EeV-scale, originating in astrophysical sources. The detector consists of a three dimensional array of "strings" each supporting 60 Digital Optical Modules (DOMs), deployed over a cubic-km in the deep Antarctic ice near South Pole Station. Each DOM contains one 10-inch photomultiplier tube (PMT) to detect the Cherenkov light produced by the products of the neutrino interactions. Completed in December 2010, the observatory augmented its physics potential with the addition of the DeepCore subarray. DeepCore, consisting of eight additional densely spaced strings, lowers the IceCube neutrino detection energy threshold from a few 100 GeV to approximately 10 GeV. In addition, IceCube incorporates a surface array, IceTop, to measure atmospheric cosmic-ray induced air showers.

The computing needs for Canadian IceCube activities will grow over the next several years as the number of researchers increases, in particular due to the development of future extensions of the South Pole facility. Based on current analyses, disk storage requirements are estimated at ~25 TB per user, reaching ~1 PB for Canadian researchers by 2021. Compute Canada's CPU and GPU resources are utilized as Tier-2 facilities for IceCube simulation production and Tier-1 facilities for simulation of the proposed detector upgrades. The current IceCube activities use on the order of 2500 CPU core-years and ~70 GPU-years. The CPU requirements are expected to reach 4000 core-years by 2021. The GPU requirements must increase significantly now in order to reach the goal of real-time production, estimated to be 500 GPU-years. It is expected that this level of GPU production will be made possible by the addition of the proposed GPU cluster illume (Kopper CFI JELF).

Theory Projects

There are a number of updates and new projects fro the theory members of the SAP community. Typically they use interconnected clusters for large-scale calculations. We list the new contributions:

1. **Ab initio electroweak reactions with nuclei (S. Bacca, TRIUMF)**
   The calculation of electroweak reactions on nuclei from first principles heavily relies on high performance computing. We have recently developed a new method to tackle electroweak break-up reactions in nuclei of light and medium-mass number. The idea is based on the combination of the Lorentz integral transform method and the coupled-cluster theory for nuclei, which are both well established many-body methods in nuclear physics. We called this new method the LIT-CC approach. An equation of motion needs to be solved in coupled-cluster theory, which demands considerable computational resources even for sd-shell nuclei and is expected to substantially increase in the future, when we will address heavier nuclei.

   Numerical codes are written in Fortran90 and utilize MPI and OpenMP parallelization. At present, production calculations are performed at Oak Ridge National Laboratory on Titan, which we have access to via our collaborations with US scientists. For our last paper, appeared on Nature Physics, we utilized 15 million core hours.

   The Titan facility has 18,688 nodes with 16 cores each and 32 GB RAM (with a Gemini fast interconnect). We typically calculate with about 200 nodes. For calculating and storing all the matrix elements of the three-body force, we need instead very large memory/node (ideally 512 GB/node). At the moment we have access to a few of those nodes in the US, but we expect this not to be sufficient in the future.

   In terms of disk space, typically we need require 10 Tb per person. We have recently developed a new method and several high impact factor publications have appeared in the last couple of years from the LIT-CC method. We expect to be able to reach heavy nuclei such as 90Zn and 132Sn, which will be a breakthrough. We foresee our group to substantially increase, reaching up to 5 Canadian researchers (including students, postdocs and the PI) to run LIT-CC codes. This sets up our computing request for the future to be access to 1000 nodes (32 cores and 32 GB RAM), 50 nodes (32 cores and 512 GB RAM) and 100 TB of disk space. We expect our requirements to grow to 5000 core-years by 2020.

2. **Ab initio calculations of medium-mass nuclei (J.D. Holt, TRIUMF)**
   Generally production-level ab initio calculations of medium-mass nuclei require a moderate number (50-100) of nodes, with 16+ cores/node and a significant amount of memory per node.

   Currently we rely on international computing facilities to provide such capabilities, primarily the JURECA machine at the Jülich Supercomputing Center in Germany, where our 2014-2015 project was awarded one of only two John von Neumann Excellence Projects in 2014 and granted 4.5M CPU hours.

As our computing needs would be met ideally through Canadian computing resources, access to large-scale machine(s) with specifications similar to the above would adequately serve this research program for the coming 3-5 years. Beyond this, we anticipate an increasing need for high-memory nodes, driven by extending the many-body codes to heavier nuclei, which also requires having nuclear forces available in increasingly larger basis spaces (three-nucleon forces in particular present tremendous memory hurdles, and can now only be performed using certain truncation schemes).
We expect our requirements to be 3000 core-years by 2020.

3. **Monte Carlo Simulation of LHC Collisions (D. Morrissey, TRIUMF)**
   Theories of new physics beyond the Standard Model can be tested against data from the LHC through Monte Carlo simulations of collision events. We use a suite of collider MC codes to simulate each event: MadGraph is used to model the initial tow-to-two collisions of elementary constituents, Pythia to simulate the radiation of additional particles and collect them into a set of hadrons, and DELPHES to model the response of the LHC detectors.

   This simulation chain is highly CPU-intensive but only requires modest RAM (1GB/core). It is also very read-write intensive, and an additional requirement is fast short-term storage and a much larger long-term storage of the simulation output. A typical simulation run consists of 100k events with an output of about 1GB, and multiple runs are needed to fill out distributions of kinematic observables and to study multiple sets of theory parameters. Longer-term storage requirements are roughly 10 Tb per person. We estimate that we will need 200 core-years by 2020.

4. **Ab initio nuclear structure and reactions of light nuclei (P. Navratil, TRIUMF)**
   This project involves large-scale ab initio nuclear structure and nuclear reaction calculations, using as input modern two- and three-nucleon forces derived within chiral effective field theory. Using these forces, the quantum many-nucleon problem is solved for bound and unbound eigenstates. The method used is called the no-core shell model with continuum (NCSMC).

   For NCSMC calculations, it is important to use a large number of nodes with a large RAM memory; at least 16 GB per node but optimally substantially more. The storage requirements are modest (few TB). The codes are written in Fortran90 or in C and utilize MPI and OpenMP parallelizations. At present, these calculations are performed at parallel computers at Lawrence Livermore and Oak Ridge National Laboratories (USA), but the computations are expected to transition to Canadian facilities in the future. In fact, we already began calculations on MP2 of Calcul Quebec with a 2016 RAC allocation.

   In 2016, we were awarded 2500 core-years of CPU time on MP2 machine of Calcul Quebec. The first exploratory calculations have already been performed using about 100 nodes.

   Despite continuous formal and technical improvements of our codes, our computing needs will grow in the future, as we plan to tackle more complex problems, i.e., perform calculations for heavier nuclei (sd-shell and beyond). Further, we will study the alpha-clustering including the scattering and reactions of alpha-particles with nuclei. These problems will require a significant increase of computing power, i.e., by a factor of 10 or more. To meet our future computing requirements for this project, dedicated machines for running massively parallel MPI and MPI/OpenMP jobs with several thousand compute nodes with a fast interconnect and a large memory per node will be needed. A machine that would allow us to run parallel jobs on about 10,000 nodes with 32 cores and 128 GB per node would be ideal. We estimate that we will need 10,000 core-years in 2020 on Compute Canada.

5. Gravitational Collapse in Quantum Gravity Systems (Frey & Kunstatter, Winnipeg)

Studies of quantum gravity at the University of Winnipeg include analysis of gravitational collapse with HPC in two contexts. One research thread will investigate the effect of higher-derivative terms in the gravitational action (for example, those in Lovelock gravity theories and others designed to emulate the emission of Hawking radiation during collapse). The second major research program concerns gravitational collapse in asymptotically anti-de Sitter spacetime, which may occur at late times after matter reflection from the conformal boundary. Through the AdS/CFT correspondence, this process maps to thermalization of small amounts of energy in large N gauge theories similar to QCD.

As indicated, numerical solution of Einstein's equations for both research programs requires highly parallel computing (up to ~100 cores) with moderate memory needs per processor, which are currently supported on local computing infrastructure in Winnipeg and on WestGrid. The AdS collapse projects particularly require very high resolution to maintain numerical convergence over many time steps. Perturbative analysis of AdS collapse requires shared memory parallel processing (OpenMP), also performed locally and on WestGrid. In all, both projects are expected to require ~1000 core years between 2016 and 2022.

6. The numerical modeling of strongly interacting matter under extreme conditions

(C. Gale, S. Jeon, McGill University)

For the past 2 decades, the study of quark-gluon plasma (QGP) has been a major focus of activity in high-energy nuclear physics. In order to understand the properties of QGP manifested in the high-quality data being produced at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory, and at the Large Hadron Collider (LHC) at CERN, high-resolution simulation of ultra-relativistic heavy ion collisions is an indispensable tool. The nuclear theory group at McGill University has been at the forefront of the theory effort to investigate this deconfined matter of quarks and gluons under extreme conditions of temperature and density. The hydrodynamics approach, and the jet energy loss calculations developed by our group (MUSIC and MARTINI, respectively) currently represent the state-of-the-art in the dynamical modeling of relativistic heavy ion collisions.

The goal of simulations is to determine the properties of QGP by comparing the simulation results with the experimental data. The properties to extract include, but are not limited to, the specific shear viscosity, the specific bulk viscosity, and the location of the critical point and the bulk behavior of the fluid near critical points. To explore the parameter space and scan the collision energies and the system sizes, we need to generate more than 200,000 hydrodynamic simulations in 2016 alone. Each of these hydrodynamic events will then play the role of the background medium in which other observables, such as electromagnetic probes and jets, evolve. All combined, it will take about 1,000 core-years of computing time and 200 TB of storage in 2016 to carry out our studies on the Guillimin system, located at McGill and managed by Compute Canada.

In the next 5 years - up to and including 2020 - the LHC and RHIC will keep producing a large volume of data which will serve to characterize the QGP, and also explore other parts of the QCD phase diagram. Therefore, the need for HPC for our group can only grow. This is especially true, considering the fact that the continuing detector sophistication will permit to explore new levels of differential observables. We project that in the coming years, our projects will require at least 1,500 core-years of computing time each year and more than 300 TB of storage space.

Table 1

Processing requirements of the SAP experimental projects

The numbers are given in core-years for each year from 2014 to 2021 (a core is 14 HEPSpecInt). The upper third of the table shows the Tier-1 type (24x7 service levels) cores, the middle third shows the Tier-2 type (5x9 service levels) computing, and the lower third shows the total computing cores. The column labeled JLAB includes all the projects at Jefferson Laboratory except GlueX. The changes to the table from the previous document are highlighted in the yellow boxes.

Experimental projects CPU-core-year requirements

	Off-shore accelerator projects						***SNOLAB projects***			***TRIUMF projects***				
Year	**ATLAS**	**T2K**	**BelleII**	**GlueX**	**JLAB**	**IceCube**	**DEAP**	**SNO+**	**Other**	**GRIFFIN**	**TIGRESS**	**PIENu**	**Other**	**Total**
	Raw data (Tier 1) requirements													
2016	4,100	0	0	0	0	0	0	0	0	100	100	0	0	**4,300**
2017	5,200	0	0	0	0	0	0	0	0	100	100	0	0	**5,400**
2018	6,200	0	0	0	0	0	0	0	0	100	100	0	0	**6,400**
2019	7,500	0	500	0	0	0	0	0	0	100	100	0	0	**8,200**
2020	9,000	0	1,000	0	0	0	0	0	0	100	100	0	0	**10,200**
2021	10,800	0	1,500	0	0	0	0	0	0	100	100	0	0	**12,500**
	Analysis (Tier 2) requirements													
2016	4,040	730	400	1,500	200	2,500	300	300	100	0	0	30	100	**10,200**
2017	5,000	790	800	1,500	200	2,500	300	300	100	0	0	60	100	**11,650**
2018	6,000	910	1,000	1,500	200	3,000	300	300	100	0	0	60	100	**13,470**
2019	7,200	1,030	1,200	1,500	200	3,000	300	300	100	0	0	60	100	**14,990**
2020	8,700	1,150	1,400	1,500	200	4,000	300	300	100	0	0	90	100	**17,840**
2021	10,400	1,270	1,600	1,500	200	4,000	300	300	100	0	0	90	100	**19,860**
	Total requirements													
2016	8,140	730	400	1,500	200	2,500	300	300	100	100	100	30	100	**14,500**
2017	10,200	790	800	1,500	200	2,500	300	300	100	100	100	60	100	**17,050**
2018	12,200	910	1,000	1,500	200	3,000	450	450	400	100	100	60	100	**20,470**
2019	14,700	1,030	1,700	1,500	200	3,000	450	450	400	100	100	60	100	**23,790**
2020	17,700	1,150	2,400	1,500	200	4,000	450	450	400	100	100	90	100	**28,640**
2021	21,200	1,270	3,100	1,500	200	4,000	450	450	400	100	100	90	100	**32,960**

Table 2
Disk storage requirements of the SAP experimental projects.
The numbers are given in terabytes (TB) for each year from 2014 to 2021. The upper third of the table shows the Tier-1 type (24x7 service levels) cores, the middle third shows the Tier-2 type (5x9 service levels) computing, and the lower third shows the total disk storage requirements. The column labeled JLAB includes all the projects at Jefferson Laboratory except GlueX. The changes to the table from the previous document are highlighted in the yellow boxes.

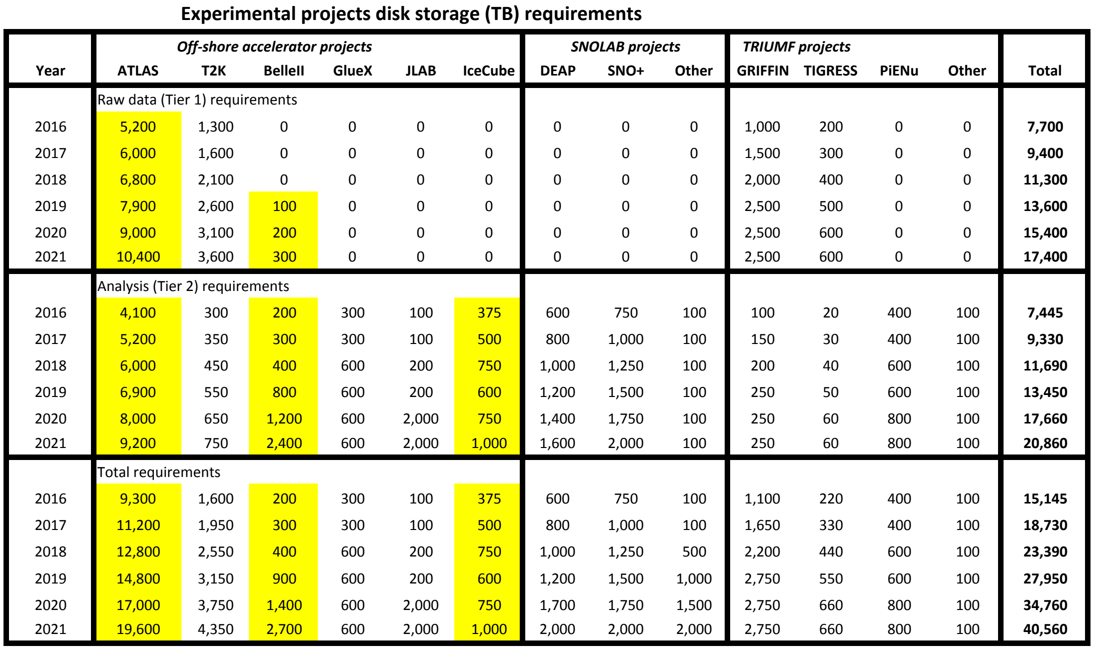

Experimental projects disk storage (TB) requirements

	Off-shore accelerator projects						***SNOLAB projects***			***TRIUMF projects***				
Year	**ATLAS**	**T2K**	**BelleII**	**GlueX**	**JLAB**	**IceCube**	**DEAP**	**SNO+**	**Other**	**GRIFFIN**	**TIGRESS**	**PiENu**	**Other**	**Total**
	Raw data (Tier 1) requirements													
2016	5,200	1,300	0	0	0	0	0	0	0	1,000	200	0	0	**7,700**
2017	6,000	1,600	0	0	0	0	0	0	0	1,500	300	0	0	**9,400**
2018	6,800	2,100	0	0	0	0	0	0	0	2,000	400	0	0	**11,300**
2019	7,900	2,600	100	0	0	0	0	0	0	2,500	500	0	0	**13,600**
2020	9,000	3,100	200	0	0	0	0	0	0	2,500	600	0	0	**15,400**
2021	10,400	3,600	300	0	0	0	0	0	0	2,500	600	0	0	**17,400**
	Analysis (Tier 2) requirements													
2016	4,100	300	200	300	100	375	600	750	100	100	20	400	100	**7,445**
2017	5,200	350	300	300	100	500	800	1,000	100	150	30	400	100	**9,330**
2018	6,000	450	400	600	200	750	1,000	1,250	100	200	40	600	100	**11,690**
2019	6,900	550	800	600	200	600	1,200	1,500	100	250	50	600	100	**13,450**
2020	8,000	650	1,200	600	2,000	750	1,400	1,750	100	250	60	800	100	**17,660**
2021	9,200	750	2,400	600	2,000	1,000	1,600	2,000	100	250	60	800	100	**20,860**
	Total requirements													
2016	9,300	1,600	200	300	100	375	600	750	100	1,100	220	400	100	**15,145**
2017	11,200	1,950	300	300	100	500	800	1,000	100	1,650	330	400	100	**18,730**
2018	12,800	2,550	400	600	200	750	1,000	1,250	500	2,200	440	600	100	**23,390**
2019	14,800	3,150	900	600	200	600	1,200	1,500	1,000	2,750	550	600	100	**27,950**
2020	17,000	3,750	1,400	600	2,000	750	1,700	1,750	1,500	2,750	660	800	100	**34,760**
2021	19,600	4,350	2,700	600	2,000	1,000	2,000	2,000	2,000	2,750	660	800	100	**40,560**

Table 3
Tape storage requirements of the SAP experimental projects

The numbers are given in terabytes for each year from 2014 to 2021. The upper third of the table shows the Tier-1 type (24x7 service levels) cores, the middle third shows the Tier-2 type (5x9 service levels) computing, and the lower third shows the total tape storage requirements. The column labeled JLAB includes all the projects at Jefferson Laboratory except GlueX. The changes to the table from the previous document are highlighted in the yellow boxes.

Experimental projects tape storage (TB) requirements

	Off-shore accelerator projects						***SNOLAB projects***				***TRIUMF projects***				
Year	**ATLAS**	**T2K**	**BelleII**	**GlueX**	**JLAB**	**IceCube**	**DEAP**	**SNO+**	**Other**		**GRIFFIN**	**TIGRESS**	**PiENu**	**Other**	**Total**
	Raw data (Tier 1) requirements														
2016	12,800	0	0	0	0	0	120	0	0		1,000	200	0	0	**14,120**
2017	20,400	0	0	0	0	0	320	0	0		1,500	300	0	0	**22,520**
2018	23,400	0	0	0	0	0	520	0	0		2,000	400	0	0	**26,320**
2019	26,900	0	1,000	0	0	0	720	0	0		2,500	500	0	0	**31,620**
2020	31,000	0	2,000	0	0	0	920	0	0		2,500	600	0	0	**37,020**
2021	35,600	0	4,000	0	0	0	1,120	0	0		2,500	600	0	0	**43,820**
	Analysis (Tier 2) requirements														
2016	0	0	0	0	0	0	0	0	0		0	0	0	0	**0**
2017	0	0	0	0	0	0	0	0	0		0	0	0	0	**0**
2018	0	0	0	0	0	0	0	0	0		0	0	0	0	**0**
2019	0	0	0	0	0	0	0	0	0		0	0	0	0	**0**
2020	0	0	0	0	0	0	0	0	0		0	0	0	0	**0**
2021	0	0	0	0	0	0	0	0	0		0	0	0	0	**0**
	Total requirements														
2016	12,800	0	0	0	0	0	120	0	0		1,000	200	0	0	**14,120**
2017	20,400	0	0	0	0	0	320	0	0		1,500	300	0	0	**22,520**
2018	23,400	0	0	0	0	0	520	0	0		2,000	400	0	0	**26,320**
2019	26,900	0	1,000	0	0	0	720	0	0		2,500	500	0	0	**31,620**
2020	31,000	0	2,000	0	0	0	920	0	0		2,500	600	0	0	**37,020**
2021	35,600	0	4,000	0	0	0	1,120	0	0		2,500	600	0	0	**43,820**

Table 4

Computing requirements of the SAP theory community.

The requirements in 2020 are given for the largest users. There is a wide range of needs: MPI-type jobs, large memory per core jobs and jobs requiring more typical (not interconnected) clusters (2-4 GB RAM/core).

Theory computing requirements

Group		**CPU ***	**Special requirements**
Lewis/Maltman (York)	Lattice QCD	5000	MPI (Infiniband), 100 TB storage
Jeon/Gale (McGill)	Relativistic quark-gluon plasma	2000	MPI (Infiniband), 250 TB storage
Navratil (TRIUMF)	ab-initio nuclear structure	200	MPI (infiniband), 32-64G RAM/core
Gezerlis (Guelph)	Neutron stars and nuclei	200	No special needs
Aleksejevs/Barkanova (Memorial and Acadia)	Electroweak physics	100	5G RAM/core, 600 TB storage
2016 additions:			
Frey & Kunstatter (Winnipeg)	Quantum Gravity Systems	1000	
Bacca (TRIUMF)	ab-initio nuclear structure	1000-5000	MPI (infiniband) and 512GB RAM/node
Navratil (TRIUMF)	ab-initio nuclear structure	10000-20000	
Morrissey (TRIUMF)	LHC physics simulation studies	200	
Holt (TRIUMF)	ab-initio nuclear structure	3000	
Gale & Jeon (McGill)	Numerical modelling of matter	1500	MPI (Infiniband), 250 TB storage

* CPU in core-years in 2020