

Computing requirements for the Canadian subatomic physics community

Institute of Particle Physics of Canada and the Canadian Institute of Nuclear Physics
July 31 2014

Executive Summary

This document describes the computing requirements for the Canadian subatomic physics (SAP) community for 2014-2021. Computing is an integral part of our research program from the analysis of large data sets from the ATLAS experiment at CERN to the theoretical calculations aimed at understanding the nature of the quark-gluon plasma.

Today the computing needs of the SAP community are dominated by the ATLAS research program. Canadians, as well as the international community, have access to the ATLAS Tier-1 centre at TRIUMF, which is currently managed and operated by ATLAS-Canada. In addition, ATLAS effectively uses the shared Tier-2 facilities that are operated by Compute Canada. The Tier-1 has dedicated resources due to its near-real-time operational and 24x7 service level requirements¹. Its primary function is the storage and processing of the raw data from ATLAS. Canada has international commitments for the operation of the Tier-1 centre and the provision of Tier-2 computing resources. The Canadian centres are integrated with similar centres around the world to form a unified computing grid infrastructure.

The SAP community is also involved in a wide-range of projects at the TRIUMF and SNOLAB Laboratories in Canada, and at a number of international laboratories. Many of these ongoing and future projects will require access to similar services for managing their raw data as provided by the ATLAS Tier-1 centre. Further, these projects require significant amounts of processing and storage capacities for the analysis of their physics data samples (Tier-2 type computing with 5x9 service level requirements).

The type of computing and the required infrastructure of the experimental SAP projects is very similar and the community would benefit, for example, from a Tier-1 centre for all its projects. In addition, the theoretical members have a more diverse set of requirements ranging from the use of highly interconnected clusters to the use of computers with very large amounts of memory.

The scale of our requirements will reach many tens of thousands of cores and tens of petabytes of storage by the end of the decade. Approximately half of those resources will need to be run in the current Tier-1 mode of operation with 24x7 service levels. The centres will require high-speed links to the research global networks with a plan for utilizing a terabit per second network by 2020.

We anticipate that these resources will be requested through either a thematic proposal directly to CFI or as part of the renewal process of Compute Canada (or a combination of both).

¹ We refer to 24x7 service levels to mean that the facility is managed 24 hours for 7 days a week and 5x9 service levels to mean a facility that is managed during working hours (5 days by 9 hours).

1. Overview

This document describes the computing requirements of the Canadian subatomic physics (SAP) community for the period 2016-2021. It has been written in response to the call for whitepapers by Compute Canada, as part of its Sustainable Planning for Advance Research Computing (SPARC) initiative that is mandated by CFI. The subatomic physics community is diverse and composed of researchers belonging to the Institute of Particle Physics² and the Canadian Institute of Nuclear Physics³.

The Institute of Particle Physics (IPP) represents the Canadian experimental and theoretical particle physics community, which conducts research into the properties of fundamental particles and their interactions. The current IPP projects are addressing fundamental questions about the nature of mass, the differences between matter and anti-matter and neutrino oscillations, as well as searching for evidence of new physics such as dark matter. The ATLAS experiment based at CERN (Switzerland) is the highest priority IPP project and it has the largest computing requirements of all projects within the SAP community. IPP researchers are also involved in the T2K experiment at JPARC/Kamiokande (Japan) and the Belle-II experiment at KEK (Japan). IPP projects at SNOLAB and TRIUMF include DEAP and SNO+, and PiENu, respectively. All these projects have significant computing requirements. There are also other IPP projects, either in a post-data-taking phase or under development; these projects tend to have modest computing requirements.

The Canadian Institute of Nuclear Physics (CINP) is a formal organization of the Canadian nuclear physics research community to promote excellence in nuclear physics research and education, and to advocate for the interests of this community both domestically and abroad. The scientific program of the CINP members includes experimental and theoretical studies of nuclear astrophysics, nuclear structure, fundamental symmetries and hadrons/QCD. The CINP projects with large computing requirements include experiments using the GRIFFIN and TIGRESS detectors at TRIUMF, the GlueX experiment at Jefferson Lab (USA), and theoretical calculations of the quark-gluon plasma, neutron stars and nuclei, lattice QCD and electroweak two-loop contributions.

The science program for the SAP community is summarized in the 2011 NSERC Long Range Plan (LRP) document⁴. As stated in the LRP, the scientific mission of subatomic physics is to identify the elementary constituents of matter and their physical properties, identify the fundamental forces through which they interact, and identify how these ingredients combine to produce the organization we see around us in nature. The LRP highlights the plans of the SAP community for the 2011-2016 period and discusses the long-term goals. The LRP discusses the impact of SAP on society and the training of highly qualified personnel. It also highlights the relationship of the SAP community with the funding agencies (NSERC and CFI), and the other organizations (Compute Canada and CANARIE) that provide critical support for our computing requirements.

Cyberinfrastructure is pervasive in all aspects of the SAP ecosystem. We use it to communicate and share information with our colleagues around the world. Our experimental detectors are highly sophisticated systems, controlled and monitored by computers integrated into all levels of the apparatus. Real-time computer systems reduce the vast amounts of collision data to more modestly sized physics analysis data samples. Computing facilities around the globe are used to process the raw

² <http://www.ipp.ca>

³ <http://www.cinp.ca>

⁴ http://www.ipp.ca/pdfs/LRP_English_final.pdf

data, generate simulated data samples, and analyze the physics samples. High-speed research networks span the globe connecting the centres into a unified system for the world community.

Canadian SAP researchers are leaders in experimental and theoretical research. An important contributor to our success is the access to world-class computing facilities and research networks. The ATLAS Tier-1 computing centre at the TRIUMF Laboratory, funded directly by CFI, is an example of such a facility. The Tier-1 centre is connected directly to CERN with a dedicated high-speed network connection provided by CANARIE that is used to stream the ATLAS raw data to TRIUMF for storage and processing. The significant resources provided by Compute Canada centres (also funded by CFI) provide Canadian SAP researchers with the critical resources to make key contributions to their projects.

This document describes the future computing requirements of the SAP community. There is a wide range of needs from the ATLAS experiment, the experiments at SNOLAB and TRIUMF, and the needs for specialized computing facilities for theoretical calculations. It is essential that the ATLAS Tier-1 centre and the Compute Canada centres be renewed to meet our current commitments. Further, the SAP community has a number of new projects that will require additional computing resources in this decade.

In the following two sections, we highlight the experimental and theoretical projects with large computing requirements. The fourth section describes the technical details and year-by-year estimates for processing and storage. The fifth section describes the current efforts in the area of software development (a new funding opportunity in the coming years).

2. Experimental Projects

A brief description of the experimental projects is presented with the focus being the projects with the largest computing requirements for the period 2016 to 2021. The ATLAS project dominates the requirements of the entire SAP community. The other projects based at international laboratories include T2K, Belle-II and GlueX, projects at TRIUMF and projects at SNOLAB.

i. **ATLAS**

The ATLAS experiment at the Large Hadron Collider (LHC) at the CERN laboratory in Geneva is studying proton-proton, proton-lead and lead-lead collisions at very high energy. The unprecedented energy densities created in these collisions allow us to study the structure of matter at very small dimensions. This makes it possible to extend investigations of the fundamental forces of Nature, to understand the origin of matter and to search for physics beyond the Standard Model (SM) of particle physics. ATLAS is collecting several petabytes of raw data each year during LHC operations and is producing petabytes of derived and simulated data sets on a continuous basis.

The vast amount of data collected by ATLAS is being analyzed on an international network of high-performance computing known as the Worldwide LHC Computing Grid (WLCG). There are ten Tier-1 centres around the world that are primarily responsible for storing and processing the raw data, and for producing various derived data sets to be distributed to the worldwide ATLAS community in a timely fashion. One of the Tier-1 centres is located at the TRIUMF Laboratory. Worldwide, ATLAS has over 150 petabytes of data stored on disk and tape storage, and on average there are 150,000 running jobs at any given moment.

There are about seventy Tier-2 centres around the world that are responsible for the generation of simulated data samples and user analysis. Four Tier-2 centres are located in Canada at Compute Canada shared facilities. In the East, we have SciNet (Toronto) and CLUMEQ (McGill), and in the West, we have Westgrid-II facilities that are located at SFU and Victoria.

One of the primary goals of the LHC was to search for the Standard Model Higgs boson, the particle thought to be central to the mechanism that gives subatomic particles mass, and the last missing ingredient of the very successful theory. In July of 2012, ATLAS and CMS experiments announced the discovery of a new particle consistent with a Standard Model Higgs boson. The discovery of the Higgs boson would not have been possible without the WLCG infrastructure. In particular, the Tier-1 and the Tier-2 facilities provided crucial extra computational resources and storage capacity that facilitated this discovery.

In 2013, the primary focus of the ATLAS experiment shifted into understanding the properties of Higgs boson. In addition, ATLAS continues to search for physics beyond the SM. The computing and storage resources are critical in extracting the science from the experiment and to full exploitation of the ATLAS data.

The ATLAS-Canada collaboration consists of 39 faculty members at nine universities, plus TRIUMF, Canada's National Laboratory for Particle and Nuclear Physics, as well as 25 postdoctoral fellows, and 69 graduate students. The group has made significant contributions to the detector and scientific output.

The Tier-1 centre at TRIUMF and the Compute Canada Tier-2 resources have been key contributions to the success of ATLAS. ATLAS-Canada has worked with Compute Canada since its inception to fully integrate its facilities into the WLCG. Compute Canada resources were used to help produce the first results from the LHC in a very short time after data started to be available at high energy in May 2010.

The computing resources requirements described herein form an integral part of the ATLAS-Canada computing system. These resources will allow us to meet our increasing commitments to ATLAS as more data are being recorded, and allow Canadians to remain competitive on the world stage. The Tier-1 centre at TRIUMF and the Compute Canada Tier-2 resources are crucial to our continued contribution to ATLAS. The computational needs for the upcoming 2015-2017 LHC collisions (Run II) will be significantly larger due to higher beam energy and intensities.

ATLAS Computing Model

ATLAS collects approximately 4 petabytes of raw data per year during LHC operations, with significant additional storage required for the secondary data sets and simulation. Analysis of particle physics data is done in stages, starting with calibration, and then moving to event reconstruction, event filtering, and finally to the extraction of physics results. Secondary data sets are produced at each stage of the analysis with each successive data set being selectively reduced from the previous event sample. The staged nature of the analysis makes it natural to use a tiered system of analysis centres. We briefly describe the issues related to the raw data, the simulated data and the physics analysis samples.

a) *RAW data handling*: the raw data from the collision events surviving the ATLAS trigger selection system are processed first on the Tier-0 system at CERN. The raw data and the results of preliminary first-pass processing are sent to the Tier-1 centres. Canada's Tier-1 centre is responsible for hosting a share of the raw and derived data. If there are times when the Tier-0 capacity is insufficient, then the Tier-1 centres perform the initial reconstruction. This was the case during heavy ion collisions running due to the complexity and size of the events. As better detector calibrations and reconstruction software become available, the Tier-1 centres reprocess the raw data to produce higher-quality derived data, which are distributed worldwide. The demands on the Tier-1 centre are significant. In addition to the processing and storage capabilities, the Tier-1 centre must continuously receive raw and derived data from CERN, as well as receive other derived data from other Tier-1 and Tier-2 centres. Because the Tier-1 centres play a critical, semi-online role in ATLAS, they must be dedicated facilities and must be operated on a 24x7 basis.

b) *Simulation data handling*: one of the primary functions of the Tier-2 centres is to produce the large amounts of simulated data needed for analyses. The simulation is CPU-intensive and the Tier-2 centres are designed for this purpose. As the simulated data at the Canadian Tier-2 centres are produced, they are copied to the Canadian Tier-1 centre for storage and further dissemination to the other WLCG sites. The simulated data is handled effectively in a similar manner to real data and derived data is produced accordingly.

c) *Physics analysis*: individual Canadian physicists typically use desktop systems for preparing and finalizing data analyses, but every analysis requires significant access to Tier-2, and in some cases Tier-1 capabilities. In all cases, analysis algorithms are developed at the Tier-2 centres on a subset of the data, before being run on the full data set at the Tier-1 or the Tier-2s. Canadian physicists also use CPU resources to create smaller simulation data sets to validate analysis strategies.

ATLAS specifies its resource requirements based on the following inputs:

- Number of seconds per year that ATLAS will collect data;
- Number of events recorded per second;
- Amount of storage required per event;
- Number of CPU-seconds required to reconstruct and analyze an event;
- Number of CPU-seconds required to simulate an event using Monte Carlo techniques and the number of simulated events required for the data analysis;
- Resources required for user analysis beyond the organized common analysis that constitutes a large part of the overall computing request.

ATLAS reviews its computing requirements on a monthly basis taking into account the performance and plans of the LHC together with the needs of the experiment. The ATLAS usage and requirements for the next 2-3 years are formally reviewed twice per year by the CERN Computing Resources Scrutiny Group, which reports its findings to the LHC Resource Review Board. In this document we use the official ATLAS resource estimates and then extrapolate the requirements through 2021.

Each country must formally pledge resources to the WLCG every August (as part of their MOU commitments to ATLAS) and those resources must be installed by April the following year. The Canadian MoU commitments are to provide 10% of the Tier-1 capacity needed

worldwide and 5% of the Tier-2 capacity; this percentage matches roughly the fraction of the total ATLAS collaboration made up by the Canadian group.

In addition to the pledged resources, ATLAS-Canada requires resources for user analysis beyond the limited amount included in the ATLAS model, which focuses mainly on the common analysis tasks that are performed by the collaboration as a whole. Experts on various aspects of the detector and event reconstruction algorithms must perform common tasks such as calibration, event reconstruction, event selection, and some of the first stages of physics analysis. This ensures that expertise is brought to bear on these crucial aspects of the analysis, and provides consistency across all analyses. Beyond these common tasks, significant resources are required to extract physics results from the secondary data sets resulting from the common stages of the analysis. We have included such resources, which are mostly outside of the ATLAS computing model, in our allocations requests to Compute Canada. The first three years of experience with real data has brought the ATLAS computing model to full maturity and the use cases are well understood. In addition to the 5% required for the Canadian contribution to ATLAS common computing, ATLAS-Canada requires additional resources for exclusive use by Canadian researchers; this corresponds approximately the same level as the international requirements. The Canadian-only portion is crucial because it allows Canadians to pursue avenues of investigation outside the standard ATLAS analysis program, providing us with a competitive advantage. The Canadian-only resources are being used in several high profile analyses, the best example being studies of the properties of the Higgs boson. The extra resources were essential in the discovery in 2012 and properties measurements carried out in 2013. For the Higgs boson discovery, large simulation samples were urgently needed to complete the analysis in time for the summer conferences and have a better understanding of the signal and background compositions in order to make a definite statement about the discovery; the Canadian Tier-1 and Tier-2 centres played a key role with their significant CPU and storage capacities and excellent performance and reliability.

ATLAS jobs are data-intensive serial jobs. It is critical that the CPU and storage resources are well connected in order to achieve the best performance. Distributing the Compute Canada resources for ATLAS across a few sites allows load balancing across sites; one site being down would have a lesser impact on the overall system. The use of multiple sites also leverages the knowledge accumulated at these sites over the past 5-7 years of operation.

ATLAS Access Model and Control

The ATLAS workload system has many components. First there are production tasks such as reconstruction, Monte Carlo simulations and common analysis tasks coordinated by different physics groups. Second, there is access to ATLAS data for the whole ATLAS collaboration (thousands of users). In effect, we are contributing resources to a worldwide system that Canadians use along with their international colleagues. Thirdly, there is access to dedicated resources reserved for Canadians that allow them to be competitive within ATLAS in the extraction of physics results. Finally, there is a need for certain Canadian users to have increased priority and access periodically to run over large data in order to complete an urgent analysis. This type of special access is called “power user access” and is implemented using grid tools. We have made all of this transparent to Compute Canada for the Tier-2 centres, once a fairshare allocation has been set. Currently, users or production jobs are handled in the following manner:

- ATLAS submits pilot jobs that once running on a node query a central database at CERN and download a “payload” job that will be executed. The pilot jobs are

submitted to the local batch system by a very limited number of users with production, analysis, and Canadian-only roles;

- The payload job runs and the output is copied to central storage and registered in the Grid catalog so it can be made available to other Grid users and for distribution and replication to other sites around the world;
- The compute node is then released to the batch system for the next job.

Priority between ATLAS users is controlled by a central service.

ATLAS Compute requirements

The grid middleware and ATLAS software are non-commercial and provided free of charge. The grid middleware for WLCG services are provided by the European Grid Infrastructure (EGI), while the ATLAS software suite is developed, controlled and distributed by the collaboration. Many Canadians are involved in ATLAS software development. Analysis codes are written by individual users or physics groups via a collective effort. The Compute Element (CE) is where ATLAS connects to the batch queues at the Compute Canada sites. Furthermore, the extensive ATLAS software suite is validated via the CE. The ATLAS software is compiled and validated only on Scientific Linux 6 (based on RedHat Enterprise) or variants thereof (e.g. CentOS 6). We have worked with existing Compute Canada sites to ensure that ATLAS jobs can run there, but care will have to be taken in the future to ensure that this remains the case. In particular the services need to scale well with the addition of CPU and storage resources at the sites and to evolve with the changes in the middleware and ATLAS data distribution model and policies. It is important to provide sufficient redundancy and fail over capability for certain critical services to maintain the up-time commitments to the ATLAS experiment and to comply with the WLCG MoU and Service Level Agreements (SLA's) with the Grid Regional Operations Center (ROC). At the moment, there is one supported method for the ATLAS software distribution: local caching mechanism on each compute node (using the CERN VM File-system or CVMFS). This requires on the order of 25 GB of local storage/partition per compute node at each site that is set aside for caching. A solution where CVMFS can be exported via NFS was also developed and tested. This solution is deployed at Simon Fraser and seems stable after extensive tuning; all of the other Canadian sites are using a physical partition on the compute nodes.

In addition, due to the increasing luminosity (i.e. intensity) of the LHC beam the individual event sizes are growing and ATLAS requests at least 2 GB RAM and 4 GB virtual memory per running process. A typical Grid job downloads a file from central storage to a local disk and then runs the analysis code producing an output file to local disk (scratch). This output file is copied back to central storage once the job is finished and local scratch space is cleaned. Data files residing at the Tier-2 sites are currently of the order of 1 GB in size. Each ATLAS job therefore requires up to 20 GB of local scratch disk space as multiple input files are actually read for efficiency and in order to reduce overhead.

ATLAS Storage requirements

ATLAS computing requires efficient data transport and fast access to the data storage. Common computing tasks (i.e. production) generate a very large amount of derived data samples. These data sets are accessed by a large number of users (in the thousands) for their specific analyses, in turn creating their own data sets. The derived data sets have to reside on disk media only and be available at all times and have to be accessible directly by running jobs. We cannot use scratch space for these and they do not need to be archived or

backed-up. Scratch space is only used for transient storage on the compute nodes. Because of the dynamic data placement and pre-placement of popular data sets, there are always multiple data sets replicated on the Grid for efficient access at the global scale. In the event of corrupted or lost files at a Tier-2 site, they can be recovered from another site for the most part, except for users' output data sets that would be lost if not replicated soon enough to another site or a local Tier-3 storage. In principle, a user can always resubmit his/her jobs and reproduce his/her output, though this can require significant CPU resources and time. Simulation files can also be lost if they are not replicated yet to the Tier-1 centre.

The storage allocation is effectively fully utilized at all times; some free space is required for efficient operations and transient jobs. ATLAS automatically applies a disk cleanup procedure when the disk areas are getting full, but the disks typically get repopulated quickly with newly placed or replicated data. ATLAS disk usage is effectively persistent and ATLAS files are already compressed. The storage system needs to be of very good performance as data files are constantly being accessed by local jobs and being transferred over the wide area network. User analysis jobs in particular are I/O limited; therefore add a higher stress on the storage systems than simulation jobs, which are CPU limited. The current performance of the sites has been satisfactory. In the future it is critical to maintain scalability while adding more storage and processing capabilities.

ii. **T2K**

The Tokai-to-Kamioka (T2K) experiment studies neutrino oscillations by sending an intense beam of muon neutrinos produced by the J-PARC complex on the east coast of Japan to the Super-Kamiokande (SK) detector 295 km away. T2K is an international effort, with over 400 collaborators from 10 nations, with approximately 40 collaborators from Canada from 6 universities and TRIUMF. It has been operational since 2010 and is now making the leading measurements of neutrino oscillation parameters. One of the goals of T2K is to find evidence of CP violation in the lepton sector over the next several years. Canada has made major contributions to the beam line and the near detectors, and has led seminal developments to the SK analysis effort.

Canada has also played an essential role in the computing for the collaboration, providing effectively half of the "Tier-2" computing resources through Compute Canada resources at Scinet and WestGrid. T2K currently does not (nor expects in the future to) require Tier-1 processing resources but does require Tier-1 storage. Currently TRIUMF hosts one of two of the T2K Tier-1 storage systems. T2K also uses Compute Canada resources for processing and storage of analysis data.

Currently T2K uses approximately 610 core-years for Tier-2 activities. A large portion of this is devoted to running the reconstruction algorithms for the Super-Kamiokande detector, which is expected to increase only modestly as beam exposure is accumulated. Compute Canada resources have also been used heavily for the relatively computing-intensive statistical analyses (Feldman-Cousins or Markov Chain calculations) that determine the oscillation parameter regions favored by the data. CPU requirements for this will also grow modestly as the number and sophistication of these analyses increase. For the near detector, the processing of data and Monte Carlo simulation will scale with beam exposure. Combined, we expect roughly a 10% increase in CPU usage each year until 2018,

when an upgrade to the accelerator complex should roughly double the beam exposure rate.

For storage, approximately 600 TB of Tier-1 storage is situated at a dedicated server at TRIUMF, with another 100 TB stored on Compute Canada resources. We expect this usage to increase by approximately 40-50% each year, with a larger increment in 2018 resulting from the higher beam intensity. Approximately 200 TB of storage is used for analysis-level Tier-2 processing. The T2K Tier-1 storage uses the same WLCG software suite as ATLAS and Belle-II in order to integrate our resources with our international colleagues. Currently, the Tier-1 network transfer rates are 1-2 Gbit/sec; in the future we expect higher rates of at least 10 Gbit/sec or higher will be necessary in order to keep up with file transfers from KEK (the "Tier 0" site for T2K) and to other sites. Currently these resources are managed directly by T2K personnel. Given the relatively specialized and complicated software suite, management of the T2K Tier-1 storage is somewhat burdensome and may be challenging for the T2K group in the longer term. We are interested in options for merging these resources with other facilities.

T2K operations are currently scheduled to continue till approximately 2022, when the currently approved exposure will be achieved.

iii. **Belle-II**

The Belle-II experiment is a new project designed to study high-energy electron-positron collisions at the SuperKEKB collider at the KEK Laboratory in Tsukuba, Japan. The goal of the project is to study CP violation (matter/antimatter symmetry) and search for evidence of physics beyond the Standard Model. The Belle-II project is an international collaboration involving over 500 researchers from 23 countries including McGill, Montreal, UBC and Victoria. Belle-II is currently constructing its detector for first collisions in 2016.

The Canadian Belle-II group will require computing and storage resources for the next 10-15 years as part of its international commitment to the project and for the computing requirements of the Canadian group. Although the experiment does not begin recording collision data until 2016, the 2014 and 2015 period will be used to generate simulated data samples for the development of reconstruction and analysis algorithms and understanding the expected performance of the detector.

The Belle-II computing model is a distributed model with a number of large centres that will store a second copy of the raw data. These centres are nearly equivalent in function to the ATLAS Tier-1 facilities, although Belle-II does not have a requirement for near-real transfer of the raw data. These large centres will also provide resources for the reprocessing of the raw data, store a fraction of the physics samples and provide computing resources for the generation of simulation samples.

Canada is expected to provide the capability for storing 10% of the raw data starting in 2019 (the raw data in 2016-2018 will be stored at Pacific Northwest National Laboratory). In addition, Canada is expected to host 10% of the physics data samples. The Belle-II computing centres are often co-located with existing WLCG centres for ATLAS or CMS. Belle-II will use the WLCG infrastructure, hence simplifying the local management requirements. The sites will either use a traditional cluster or a cloud-based system (both are in operation in Canada).

By 2021, the Belle-II project will require 1500 core-years of processing per year for raw data reconstruction and another 1500 core-years of processing per year for simulation and physics analysis computing. Belle-II will also need 2400 TB of disk storage and 7000 TB of tape storage by 2021.

iv. **Jefferson Lab projects**

Jefferson Lab hosts a high-luminosity continuous-wave superconducting electron linac, which has just been upgraded to 12 GeV maximum beam energy. Several groups of Canadian experimentalists and theorists perform research there, making it the largest centre for offshore Canadian nuclear physics research.

The experiment with the most significant computing demands is GlueX, which is currently in its commissioning phase and is scheduled to start running in Hall D in 2015. A long-standing goal of hadron physics has been to understand how the quark and gluonic degrees of freedom that are present in the fundamental QCD Lagrangian manifest themselves in the spectrum of hadrons. Of particular interest is how the gluon-gluon interactions might give rise to physical states with gluonic excitations. One such class of states is a hybrid meson, naively pictured as a $q\text{-}\bar{q}\text{-}g$ state. QCD-based calculations support this notion by predicting the existence of several light hybrid mesons that have exotic quantum numbers that cannot be formed from a simple quark/anti-quark pair and need instead to carry explicit gluonic degrees of freedom. To achieve this goal, all possible decay modes of conventional and hybrid mesons must be studied systematically, in a spectroscopic fashion. The primary goal of the GlueX experiment in Hall D is to search for and study these exotic hybrid mesons, which do not mix with other $q\text{-}\bar{q}$ states and can be uniquely identified based on their $J(PC)$ quantum numbers.

The GlueX computing model is for the primary data to be permanently stored at Jefferson Lab and first level processing on those data to be performed there. Monte Carlo simulations and analysis will be performed off-site, with approximately 30% of the simulations to be performed on Canadian facilities. For the Canadian fraction of the simulations an average demand of 1500 Intel cores (i7 at 2GHz) is projected for 5 years starting in 2016, with a typical usage pattern of one large job cluster of several million cpu hours, several times per year, and 300TB of grid-accessible storage starting in 2016, increasing to twice that by 2018.

Other Jefferson Lab experiments include the Moller experiment, whose data taking is projected to be from 2020-25, plus an additional 4 years of data analysis afterward, and experiments in Hall C of deep exclusive electron scattering reactions starting in 2016. Both of these experiments would need Compute Canada resources to run large simulations and perform data analysis. Once Moller data taking starts in 2020, storage needs are projected to increase from 200 TB to 2 PB. Another Canadian group performing work at the Triangle Universities Nuclear Laboratory (USA) has more modest computing requirements.

v. **SNOLAB projects**

SNOLAB is an underground science laboratory specializing in neutrino and dark matter physics. Located 2 km below the surface in the Vale Creighton Mine located near Sudbury Ontario Canada, SNOLAB follows on the important achievements in neutrino physics achieved by SNO.

There are many projects at SNOLAB with significant Canadian involvement: DEAP-3600, HALO, PiCo, SNO+, and Super CDMS. HALO is a search for galactic supernova neutrinos. SNO+ is a search for neutrinoless double-beta decay, a rare decay which if existent would shed considerable light on the nature of neutrinos and mass. SNO+ also will measure neutrinos from the earth, the sun and galactic supernova. DEAP, PiCo, and SuperCDMS are direct searches for dark matter.

Next generation experiments are being considered including DEAP-50t, a scale up of DEAP from 3600kg to 50 tonnes of liquid argon; nEXO, the next scaling of the Enriched Xenon Observatory for neutrinoless double beta decay; and Pico-250L, a 500kg-scale experiment looking for spin dependent interactions with cosmological dark matter.

The computing needs of SNOLAB experiments are currently dominated by the DEAP-3600 and SNO+ experiments. As the computing requirements for DEAP-3600 and SNO+ are quite well known, they will be used to estimate computer usage by SNOLAB experiments. Both experiments will average data-taking rates of approximately 5 to 10 Mbytes/second. Thus a Tier-2 system is needed for data analysis and simulation. For DEAP-3600 we expect 250 TB of disk space used per year (cumulative to a total of approximately 1 PB and approximately 350 TB/year for SNO+). Other computing needs at SNOLAB are small.

Next generation experiments will also require computational resources. The design stages will require a few hundred core-years over the next three years. Then operation will require approximately double the needs of DEAP-3600 and SNO+. Larger experiments require more background simulation but the scale up of DEAP to 50t will use liquid argon depleted in the isotope Ar-39, thus reducing the data rate. Approximately 2 to 4 PB of data storage and 800-1000 core-years of processing will be required in the second half of the decade.

vi. **TRIUMF projects**

TRIUMF is Canada's national laboratory for nuclear and particle physics research and is operated as a joint venture by a consortium of 18 Canadian universities. TRIUMF provides the national research community with access to research infrastructure and opportunities for global impact far beyond the capacity of any individual institution.

Prominent among the research facilities at TRIUMF is ISAC, the only source in Canada of rare isotope beams and the world's highest power facility utilizing the Isotope Separation On-Line (ISOL) production technique. The ISAC complex at TRIUMF is currently being further expanded to include the new Advanced Rare Isotope Laboratory (ARIEL), which will come on line in the next few years. Fourteen individual experimental facilities are permanently stationed in ISAC and use the rare-isotope beams to further our understanding in the areas of nuclear structure, nuclear astrophysics and fundamental symmetries. We highlight three projects, GRIFFIN, TIGRESS and PiENu, which have the largest computing requirements.

GRIFFIN

Gamma-Ray Infrastructure For Fundamental Investigations of Nuclei (GRIFFIN) is a state-of-the-art new gamma-ray spectrometer for nuclear structure, nuclear astrophysics, and fundamental symmetries research with low-energy radioactive ion beams at TRIUMF-ISAC. The first \$8.9M phase of GRIFFIN was funded in the 2009 CFI competition.

GRIFFIN has now been installed in ISAC-I during the first half of 2014 and will begin scientific operation in the fall of 2014. The spectrometer consists of sixteen large-volume, hyper-pure germanium clover detectors, coupled with a suite of ancillary detector sub-systems, which all view the implantation location where the low-energy radioactive beam from ISAC is stopped. The radiations emitted, when the beam nucleus decays, are read out using a custom-designed digital data acquisition system. These observed radiations reveal the properties of the nucleus in order to understand the fundamental interactions, which bind nuclei, as well as the evolution of nuclear structure phenomena as a function of proton and neutron number. The decay properties of neutron-rich nuclei are especially important to furthering our understanding of the astrophysical rapid neutron capture process responsible for the generation of the heavy elements in the universe.

GRIFFIN has been designed to perform experiments using radioactive beams, which are produced at ISAC with intensities from 0.01 particles per second to 10^8 particles per second. In the former case, the data collection rates will be small due to the low intensity of the rare isotope beam and typically a few TB per year. In the latter case with high-intensity rare isotope beams the important physics is often revealed in the observation of very weak gamma-ray transitions from excited nuclear states, which have a relative intensity 5 orders of magnitude lower than the strongest decay branch. The study of these very weak transitions necessitates a very high throughput data acquisition to collect very high-statistics datasets, followed by intensive analysis to identify specific coincidence events. In order to limit the collected data to only the most useful coincidences, the data is filtered for various detector multiplicity and temporal coincidence conditions in real time before being stored on disk. Despite this real-time filtering of the detector signals, the GRIFFIN data acquisition system is capable of writing filtered data at a rate of 300 MB per second and will therefore collect datasets of around 250 TB in typical one-week experimental runs. It is anticipated to perform an average of two of these high-rate runs per year and accumulate 500 TB of new data per year, which will be stored for a minimum of five years while the analysis is completed. The processing requirements of GRIFFIN is listed in Table 1.

TIGRESS

The TRIUMF-ISAC Gamma-Ray Escape Suppressed Spectrometer (TIGRESS), a state-of-the-art position-sensitive gamma-ray spectrometer for research with accelerated rare isotope beams from ISAC-II was funded through a 6-year \$8.1M NSERC Research Tools and Instruments (RTI), Category 3 award in 2003.

The spectrometer is utilized in a diverse scientific program in nuclear structure and nuclear astrophysics research. TIGRESS consists of up to sixteen highly segmented hyper-pure germanium clover detectors, which surround a reaction target where a nuclear reaction is induced by the accelerated radioactive beam from ISAC-II. Charged particles and heavy ions following the reaction are detected in various ancillary detector systems located inside the vacuum chamber at the same time as emitted gamma rays are detected in the TIGRESS detectors. Coincidence triggering conditions between detector sub-systems can be made very selective with such a setup, and due to the typically low intensity of accelerated radioactive ion beams the number of individual events recorded is not as large as in the high-rate decay spectroscopy experiments with GRIFFIN. However, in order to make full use of the position sensitivity of these highly segmented TIGRESS clover detectors it is common to record a short waveform sample from the detector along with the other event data, which is processed in the offline analysis. This waveform collection can increase

dramatically the size of the recorded dataset. The TIGRESS digital data acquisition system was custom designed and built as part of the original installation. It is anticipated that the full digital data acquisition system will be upgraded in the 2015/2016 period with the more modern digitizers that have been developed for GRIFFIN project. Once this new DAQ system is in place the expected data collection rate will be typically 100 TB per year. The data will be stored for a minimum of five years while the analysis is completed. The processing requirements of TIGRESS is listed in Table 1.

PiENU

The PIENU experiment at TRIUMF aims to precisely measure the branching ratio of rare pion decays to electrons and muons to search for new forces of nature with mass scales up to 1000 TeV. The PIENU experiment aims to improve the precision of the pion branching ratio measurement by a factor of five, confronting the precise Standard Model theoretical prediction to better than 0.1%. In addition to the measurement of the pion branching ratio, the PIENU experiment searches for evidence of heavy neutrinos, which may occur rarely in pion decays to electrons. A large data set accumulated in the period 2010 to 2012 is now under intense analysis. PIENU is performed by an international group of 25 scientists from eight institutions and five countries led by UBC and TRIUMF. CPU utilization will increase due to additional GEANT4 based Monte Carlo Simulations that are required. In addition, participation in the CERN experiment NA62, is envisioned. The CERN NA62 experiment is designed to study rare kaon decays. Specifically, NA62 will measure the rate at which the kaon decays into a charged pion and a neutrino-antineutrino pair among many other reactions. A large amount of Monte Carlo Simulations and data analysis is foreseen.

The computing resources, currently at 30-core-years, are expected to grow to 60 core-years in 2015-16 and 90 core-years in 2017-2020. Currently we utilize 200 TB with the PIENU experiment data and some reference Monte Carlo Simulations. The project will require storage of approximately 800 TB.

3. Theory Projects

i. Lattice QCD (Lewis and Maltman, York)

Lattice QCD is a computational method to obtain a quantitative understanding of strong interactions. Two major thrusts are proposed here. The first will focus on vector mesons and vector current two-point functions and will also yield a precise determination of the muon's anomalous magnetic moment, thereby reducing the uncertainty on the quantity that has the most tension with its SM prediction to date. The second will focus on light nuclei, including the H-dibaryon.

For the vector meson project, we aim to compute the current-current correlation functions in the vector channel to high precision. This entails pushing the lattice calculation to physical quark masses, or short chiral extrapolations, and removing the lattice cut-off by taking the continuum limit. In addition we will set up a basis of interpolation operators that includes also two-pion states. We will carry out a finite volume scaling and subsequently extract the scattering phase shifts, i.e. infinite volume results. The successful conclusion of this project will yield the first complete and precise result on the vector meson spectrum ab initio calculations, with all lattice systematics under control. The project needs about 4500 core-years of CPU and at least 10 TB of storage.

For the light nuclei project a LapH/Distillation method will be implemented, which enables the computation of so-called perambulators. Storing these perambulators enables the formation of a multitude of particle systems a posteriori. This procedure, however, requires storing full propagators to disk, entailing about 70 TB of storage and more than 400 core-years of CPU.

ii. Relativistic quark-gluon plasma (Jeon & Gale, McGill)

When the nuclear temperature exceeds 170 MeV/kB or about 2 trillion kelvin, the quarks and gluons inside hadrons are liberated to form a Quark Gluon Plasma (QGP). This new state of matter existed in bulk only when the universe was about a microsecond old. Recent years have seen the culmination of decades of experimental efforts in probing the predicted phase transition to QGP by producing this hottest and densest matter at RHIC (Relativistic Heavy Ion Collider) and the LHC. The focus of the McGill group's research is on the theoretical and phenomenological investigation of the ultra-relativistic heavy ion collisions in general and the properties of QGP in particular.

The computational challenge is to simulate the entire evolution of the ultra-relativistic heavy ion collision, starting from the Lorentz-contracted nuclear initial state, to the QGP phase, and on to the final hadronic state. Each of these steps needs different computing environment: The most CPU-intensive step is the hydrodynamic evolution part, using up to 8 million spatial cells and $O(10^3)$ time steps., utilizing a parallel computing capability. Another CPU intensive step is the hadronic scattering simulation, where in a typical LHC collision, $O(10^4)$ hadrons emerge out of the QGP. Keeping track of collisions among these particles is numerically challenging. This part is not parallelized, and requires a massive number of serial nodes since $O(10^2)$ hadronic simulations need to be performed on each of the $O(10^3)$ hydrodynamic simulations. Calculations for the outgoing hard and electromagnetic probes have yet different computing needs, requiring a massive storage capability of $O(10^2)$ TB. It is estimated that at least 2,000 core-years with 3–4 GB/core per year and 250 TB of storage space is required in coming years.

iii. Ab-initio nuclear structure (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 codes are written in Fortran90 and utilize MPI and OpenMP parallelizations. At present, these calculations are performed at the parallel computers at Lawrence Livermore and Oak Ridge National Laboratories (USA), but the computations are expected to transition to Canadian facilities in the future. The computing allocation at ORNL is about 20 million core hours per year., and the calculations use up to 6000 nodes (96000 cores). On the LLNL machines, 128 nodes (2048 cores) are typically used, with CPU usage exceeding a million core hours per year. For these calculations, it is important to use nodes with a large RAM memory; most machines have 32-64 GB/node. The storage requirements needs are modest (few TB). The 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.

iv. Quantum Monte Carlo simulations of neutron stars and nuclei (Gezerlis, Guelph)

A variety of microscopic ab-initio simulation methods are used to compute the properties of neutron stars and light nuclei. For dilute neutron matter (s-wave interactions), the method of choice is Diffusion Monte Carlo (DMC). For more complicated nuclear interactions, nuclear Green's Function Monte Carlo (light nuclei) and Auxiliary Field Diffusion Monte Carlo (infinite matter and medium-mass nuclei) are used. All these methods are exact, modulo the fermion-sign problem and can be extended to become effectively variational methods. These methods use as an Ansatz a trial wave function, which embeds sufficient physical insights (from few- or many- body theory) as well as variational parameters, which are systematically varied to approach the state of chosen symmetry.

The computations are currently performed on SHARCNET clusters, and typically use 100-1000 cores. Memory and networking needs are not large, but the codes are computation-intensive. Over the next 10 years, needs are projected to increase by a factor of 50-100: the biggest issue is the number of particles, N , that can be handled in the simulation. Since the time required scales as N^3 , doubling the number of particles takes 8 times more resources. We are currently studying neutron-star related systems composed of roughly 100 particles and would need to address systems of 400-500 to be truly realistic.

v. Two-loop effects in precision electroweak measurements (Aleksjevs, Memorial & Barkanova, Acadia)

Although the Standard Model (SM) has been enormously successful to date, we know it is incomplete. A very promising method to search for physics beyond the SM is to perform high-precision electroweak experiments at the intensity/precision frontier, such as such as parity-violating Møller scattering, e^+e^- collisions or electron-nucleon scattering. These studies can provide indirect access to physics at multi-TeV scales and play an important complementary role to the LHC research program. These lower-energy experiments tend to be less expensive than experiments at the high-energy colliders, but they are more model-dependent and require significant theoretical input. In order to match the proposed electroweak experimental precision, it is necessary to provide theoretical predictions on the observables of the SM with $\sim 1\%$ precision. Since the most of the SM electroweak theoretical predictions employ perturbative expansion in orders of α (fine structure constant), it will be required to consider contributions to the electroweak cross sections of up to $\sim \alpha^4$, corresponding to two-loop calculations in the Feynman diagram approach. For QED processes, higher order calculations (up-to five loops) have been carried out already with a high degree of accuracy for some specific processes, however, in the case of electroweak processes full two-loop calculations face dramatic complications due to the presence of massive vector bosons (W^\pm and Z) in the two-loop integrals.

Consider the Møller process proposed for high precision study at Jefferson Lab as an example. For the calculation of the parity-violating electroweak asymmetries up to two-loop level, it is required to compute $\sim 127,000$ Feynman diagrams. Our estimate of the computing time for a single two-loop graph is of order of several minutes. In order to complete these calculations in several months (not including package development and deployment), ~ 100 dedicated nodes with the Mathematica parallel toolkit capable of

running one two-loop graph per node are required. Based on the assumption that a one loop graph analytical result requires around ~ 5 Gb of RAM, about 500 GB of allocated RAM per 100 two-loop graphs are required. Finally, since the first stage of the calculation deals with analytical computations, it is required to store the results permanently, needing of the order of 600 TB of storage. The second stage deals with numerical evaluation only, which can be completed using any available HPC resource in Canada. These two-loop calculation techniques can later be adapted for electron-proton processes, electron-positron collisions, and other low-energy experiments involving particles of the SM and new physics.

4. Technical requirements

The computing requirements of the SAP community are presented in the following four tables: Tables 1, 2 and 3 give the expected needs of the experimental projects (processing, disk storage and tape storage) and Table 4 gives the theoretical computing requirements.

Each of the first three tables for the experimental projects is split into two use cases: raw data storage and processing (Tier-1 computing) and analysis (Tier-2 computing). The primary distinction between the two types, beside their function, is their level of service: Tier-1 computing requires 24x7 support while for Tier-2 computing, the expected service level is 5x9 or best effort. Currently, only the ATLAS experiment has requirements for Tier-1 computing and the ATLAS-Canada group has operated a dedicated centre at the TRIUMF Laboratory since 2007. The facility is committed to storing its share of the ATLAS raw data within a few hours of it being recorded by the ATLAS detector at CERN. The ATLAS Tier-1 centre was funded in part by CFI, independently of Compute Canada, and is solely managed and operated by ATLAS-Canada. TRIUMF is operated on a 24x7 basis and it was deemed to be a natural location for the facility.

The ATLAS data from CERN is transferred to the TRIUMF centre where it is stored on tape. The disk system is used for large derived data sets and to act as buffer space for efficient tape operations during data reprocessing campaigns. The processors are used to reconstruct the raw data into physics analysis data sets that are sent to the other facilities around the world. In the coming years, the ATLAS project will continue to dominate use of the Tier-1 facility, however, other projects such as T2K, Belle-II, GRIFFIN and TIGRESS, will require a similar 24x7 facility for the storage and processing of their raw data.

The second use case for the experimental projects is the generation of simulated data samples and the provision of resources for the analysis of the physics data samples (Tier-2 computing). Each project uses simulated data samples for modeling their experimental apparatus, studying the response of particles to the material in the detector, and developing analysis algorithms. The generation of simulated data requires relatively little input data (few terabytes) and is CPU-intensive computing (well suited, for example, to computing clouds).

The analysis of physics data on a Tier-2 facility is chaotic, high I/O computation requiring well-designed high-throughput batch clusters with high-speed access to large enterprise-type storage systems. A user analysis job, in the ATLAS project for example, will be split into many shorter (6-12 hour) jobs each accessing between 10-100 GB of data. These types of computing uses are well suited to shared-computing facilities such as those operated by Compute Canada. All of the SAP projects, including ATLAS, have successfully used Compute Canada resources over the past few years.

The number of core-years of processing required by the experimental projects is 11,000 in 2014 and will grow to approximately 33,000 by 2021. The number of cores is nearly equally split between the Tier-1 and Tier-2 usage. The ATLAS experiment is the dominant project, corresponding to approximately three quarters of the SAP computing requirements. The experimental projects use 64-bit x86 processors with 2-4 GB of RAM per core with some local (internal to the node) disk (40-100 GB/core). In future, the RAM requirement will likely be at least 4 GB per core. Our expectation is that the CPU benchmark may increase by 50% from 2014 to 2021, based on our past experience. Significant effort is being made by projects like ATLAS to make more efficient use of multi-core nodes, which might reduce the memory requirement (and cost) per node. The usage of graphics processors (GPU's), with their massively parallel architecture, can significantly reduce the overall computing time per event for several analysis applications that can easily be parallelized. Recently, it was successfully used in ATLAS for matrix elements calculations; the demand for GPU computing in the SAP community is expected to grow in the coming years.

The disk storage requirements are approximately 13 PB in 2014 and growing to 43 PB by 2021, with a nearly equal split between Tier-1 and Tier-2 facilities. The disk systems need to be able to provide high-speed access (both read and write) to the data. The expectation is that disks will continue to increase in size per unit cost (possibly by a factor of 2-4 times), as a result the requirements in 2021 are comparable to the current requirements once technology evolution is taken into account.

The tape storage is used primarily to archive the raw data for long-term storage and security. Further, if the older data requires reprocessing then it can be staged from tape to a disk cache. Approximately 40 PB of tape storage is required in 2021, which is again comparable to the amount used today once technology evolution is taken into account.

The requirements of the theory community vary for each group and are listed separately in Table 4. Some groups require large numbers of cores in an interconnected (such as an Infiniband network) cluster whereas others require nodes with large amount of RAM (32-64 GB) per core.

Not included in the tables is the requirement for high-speed networks within the data centres and the national research network provided by CANARIE. The CANARIE network is migrating to a 100G backbone and our links to the US and Europe will soon be 100G. It is essential that all sites have 100G links and have the potential for further upgrades. Network speeds are increasing by a factor of two per year and a 1000G network may be the standard in 2021.

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.

Experimental projects CPU-core-year requirements

Year	<i>Off-shore accelerator projects</i>					<i>SNOLAB projects</i>			<i>TRIUMF projects</i>				Total
	ATLAS	T2K	BelleII	GlueX	JLAB	DEAP	SNO+	Other	GRIFFIN	TIGRESS	PIENu	Other	
Raw data (Tier 1) requirements													
2014	5,500	0	0	0	0	0	0	0	0	0	0	0	5,500
2015	6,170	0	0	0	0	0	0	0	100	0	0	0	6,270
2016	7,040	0	0	0	0	0	0	0	100	100	0	0	7,240
2017	8,020	0	0	0	0	0	0	0	100	100	0	0	8,220
2018	9,150	0	0	0	0	0	0	0	100	100	0	0	9,350
2019	10,430	0	630	0	0	0	0	0	100	100	0	0	11,260
2020	11,890	0	1,100	0	0	0	0	0	100	100	0	0	13,190
2021	13,550	0	1,500	0	0	0	0	0	100	100	0	0	15,250
Analysis (Tier 2) requirements													
2014	3,130	610	300	500	200	300	300	100	0	0	30	100	5,570
2015	3,790	670	300	500	200	300	300	100	0	0	30	100	6,290
2016	4,340	730	330	1,500	200	300	300	100	0	0	30	100	7,930
2017	5,230	790	630	1,500	200	300	300	100	0	0	60	100	9,210
2018	7,060	910	630	1,500	200	300	300	100	0	0	60	100	11,160
2019	8,800	1,030	770	1,500	200	300	300	100	0	0	60	100	13,160
2020	10,100	1,150	1,200	1,500	200	300	300	100	0	0	90	100	15,040
2021	12,160	1,270	1,500	1,500	200	300	300	100	0	0	90	100	17,520
Total requirements													
2014	8,630	610	300	500	200	300	300	100	0	0	30	100	11,070
2015	9,960	670	300	500	200	300	300	100	100	0	30	100	12,560
2016	11,380	730	330	1,500	200	300	300	100	100	100	30	100	15,170
2017	13,250	790	630	1,500	200	300	300	100	100	100	60	100	17,430
2018	16,210	910	630	1,500	200	450	450	400	100	100	60	100	21,110
2019	19,230	1,030	1,400	1,500	200	450	450	400	100	100	60	100	25,020
2020	21,990	1,150	2,300	1,500	200	450	450	400	100	100	90	100	28,830
2021	25,710	1,270	3,000	1,500	200	450	450	400	100	100	90	100	33,370

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.

Experimental projects disk storage (TB) requirements

Year	<i>Off-shore accelerator projects</i>					<i>SNOLAB projects</i>			<i>TRIUMF projects</i>				Total
	ATLAS	T2K	BelleII	GlueX	JLAB	DEAP	SNO+	Other	GRIFFIN	TIGRESS	PIENU	Other	
Raw data (Tier 1) requirements													
2014	7,680	700	0	0	0	0	0	0	10	0	0	0	8,390
2015	8,090	1,000	0	0	0	0	0	0	500	10	0	0	9,600
2016	9,660	1,300	0	0	0	0	0	0	1,000	200	0	0	12,160
2017	9,830	1,600	0	0	0	0	0	0	1,500	300	0	0	13,230
2018	10,250	2,100	0	0	0	0	0	0	2,000	400	0	0	14,750
2019	10,780	2,600	200	0	0	0	0	0	2,500	500	0	0	16,580
2020	12,010	3,100	400	0	0	0	0	0	2,500	600	0	0	18,610
2021	13,370	3,600	700	0	0	0	0	0	2,500	600	0	0	20,770
Analysis (Tier 2) requirements													
2014	3,040	200	50	300	100	200	250	100	5	0	200	100	4,545
2015	3,160	250	150	300	100	400	500	100	50	1	200	100	5,311
2016	4,310	300	200	300	100	600	750	100	100	20	400	100	7,280
2017	5,640	350	400	300	100	800	1,000	100	150	30	400	100	9,370
2018	6,370	450	400	600	200	1,000	1,250	100	200	40	600	100	11,310
2019	6,740	550	1,000	600	200	1,200	1,500	100	250	50	600	100	12,890
2020	9,180	650	1,600	600	2,000	1,400	1,750	100	250	60	800	100	18,490
2021	11,940	750	2,400	600	2,000	1,600	2,000	100	250	60	800	100	22,600
Total requirements													
2014	10,720	900	50	300	100	200	250	100	15	0	200	100	12,935
2015	11,250	1,250	150	300	100	400	500	100	550	11	200	100	14,911
2016	13,970	1,600	200	300	100	600	750	100	1,100	220	400	100	19,440
2017	15,470	1,950	400	300	100	800	1,000	100	1,650	330	400	100	22,600
2018	16,620	2,550	400	600	200	1,000	1,250	500	2,200	440	600	100	26,460
2019	17,520	3,150	1,200	600	200	1,200	1,500	1,000	2,750	550	600	100	30,370
2020	21,190	3,750	2,000	600	2,000	1,700	1,750	1,500	2,750	660	800	100	38,800
2021	25,310	4,350	3,100	600	2,000	2,000	2,000	2,000	2,750	660	800	100	45,670

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.

Experimental projects tape storage (TB) requirements

Year	<i>Off-shore accelerator projects</i>					<i>SNOLAB projects</i>			<i>TRIUMF projects</i>				Total
	ATLAS	T2K	BelleII	GlueX	JLAB	DEAP	SNO+	Other	GRIFFIN	TIGRESS	PIENu	Other	
	Raw data (Tier 1) requirements												
2014	5,500	0	0	0	0	0	0	0	0	0	0	0	5,500
2015	7,150	0	0	0	0	0	0	0	0	0	0	0	7,150
2016	9,240	0	0	0	0	0	0	0	1,000	200	0	0	10,440
2017	11,880	0	0	0	0	0	0	0	1,500	300	0	0	13,680
2018	13,660	0	0	0	0	0	0	0	2,000	400	0	0	16,060
2019	18,430	0	2,000	0	0	0	0	0	2,500	500	0	0	23,430
2020	23,820	0	4,000	0	0	0	0	0	2,500	600	0	0	30,920
2021	30,620	0	7,000	0	0	0	0	0	2,500	600	0	0	40,720
	Analysis (Tier 2) requirements												
2014	0	0	0	0	0	0	0	0	0	0	0	0	0
2015	0	0	0	0	0	0	0	0	0	0	0	0	0
2016	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
2018	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
2020	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
	Total requirements												
2014	5,500	0	0	0	0	0	0	0	0	0	0	0	5,500
2015	7,150	0	0	0	0	0	0	0	0	0	0	0	7,150
2016	9,240	0	0	0	0	0	0	0	1,000	200	0	0	10,440
2017	11,880	0	0	0	0	0	0	0	1,500	300	0	0	13,680
2018	13,660	0	0	0	0	0	0	0	2,000	400	0	0	16,060
2019	18,430	0	2,000	0	0	0	0	0	2,500	500	0	0	23,430
2020	23,820	0	4,000	0	0	0	0	0	2,500	600	0	0	30,920
2021	30,620	0	7,000	0	0	0	0	0	2,500	600	0	0	40,720

Table 4

Computing requirements of the SAP theory community.

The requirements 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

5. Software development

CFI has indicated that, in addition to funding hardware and operational support, it will consider providing support for software development. Typically such work in SAP has been funded by NSERC and CANARIE. We anticipate that the SAP will request some small level of funding for the development of software services and tools (in addition to the manpower required for operation of the centres).

There are a number of examples of Canadian contributions that have benefited the global SAP community but also other fields of research. The areas include cloud computing, software and data management, and high-speed networks. We list a few examples in the following text.

BaBar-Canada had a leading role in the development of the BaBar Long Term Data Access system (in operation at SLAC), a cloud-based platform for long-term access to the BaBar data and software. International Committee on Future Accelerators (ICFA) has set up a working group to investigate technologies for the long-term storage and access of HEP data. This has evolved into the establishment of Data Preservation in HEP (DPHEP) collaboration and IPP has agreed to participate in this collaboration.

Both ATLAS and Belle-II are using cloud computing resources with software developed by the Canadian SAP community. The computing cloud at CERN is operated using software developed in Canada. The Belle-II-Canada group is studying the use of commercial cloud providers (such as Amazon) for cloud bursting when our traditional resources are fully utilized.

The integration of distributed computing centres relies on high-speed networks. The Canadian SAP community is involved in international projects (including industry) providers to understand how to fully exploit the emerging 100G networks. In 2014, the SAP community was the first to test a 100G transatlantic link from Canada directly to CERN.

Software development offers a unique opportunity to train highly qualified personnel. We observe that staff with software expertise are in very high demand both in the global SAP world and by industry. A modest amount of resources could strengthen Canada's impact in the areas of cloud computing and high-speed networks.

6. Summary

Cyberinfrastructure is a critical component of SAP research. This document has described the impact of computing and how it has enabled Canadians to play a leading international role in SAP research projects.

The needs of the SAP community are dominated by the ATLAS experiment but many other projects have significant computing requirements. Currently, SAP has access to the ATLAS Tier-1 centre at TRIUMF (funded directly from CFI) and the shared Tier-2 facilities operated by Compute Canada. It is imperative that the resources be renewed. Further, there are a number of new projects in this decade requiring large computing resources.

The SAP computing requirements have been well-served by the shared resources of Compute Canada but it is critical that the ATLAS Tier-1 centre, with its 24x7 service level requirement, be maintained. The near-real-time requirement was the motivation for building a dedicated facility at TRIUMF. Approximately half of the projected SAP computing resources will need to be operated in the similar manner in the future.

The use of the shared resources, operated by Compute Canada, has been a success and enabled the SAP community to meet its international commitments and to produce world-class science. It is important that the SAP community have access to multiple Compute Canada centres for redundancy and failover, as well as making efficient use of the manpower. Hardware refresh cycles, as well as cluster designs and specifications should be done in consultation with the community.

Further, it is critical that all the centres be connected to a world-class network during this period. Our research is reliant on links to the international community and we will continue to strongly support CANARIE. Compute Canada will need to ensure the centres are linked to the national network and that the storage facilities are able to meet in the data transfer requirements.

CFI has indicated that there are two options for funding computing: the first is through the renewal of Compute Canada resources, and second, through a domain-specific or thematic proposal. The options to consider for the SAP community include (i) a single thematic proposal for all SAP; (ii) including all the SAP computing within the Compute Canada renewal; and (iii) a thematic proposal for a Tier-1 facility combined with the renewal of the Tier-2 computing resources with Compute Canada. The final approach will depend on the details of the program from CFI and will require further discussion within the SAP community in concert with Compute Canada to ensure that the SAP computing needs are fulfilled and secured for several years.

About this document:

This document was prepared by a joint IPP/CINP task force: Adam Garnsworthy, Garth Huber, Chris Jillings, Randy Lewis, Randall Sobie (Chair), Reda Tafirout, and Hirohisa Tanaka.

Submitted on behalf of the CINP and IPP to Compute Canada on July 31, 2014.