

The ATLAS Canada Computing Model

The ATLAS Canada Collaboration*

Abstract

The ATLAS experiment is the next generation of particle physics projects. ATLAS will allow physicists to study nature at an unprecedented energy scale in pursuit of the understanding of the mechanism by which the elementary constituents of matter acquire mass. The knowledge gained will have an impact on our view of the early Universe, as well as its fate. The experiment will require significant improvements in our ability to acquire, distribute and analyze enormous data sets. Canadians already play a significant role in ATLAS, and must develop computing and analysis facilities to lead the collaboration in the extraction of new physics from ATLAS. The Canadian ATLAS computing and analysis model is described in this document.

*The ATLAS Canada collaboration consists of faculty members, scientists, engineers, technicians and graduate and undergraduate students from the University of Alberta, University of British Columbia, Carleton University, McGill University, l'Université de Montréal, Simon Fraser University, University of Toronto, University of Victoria and York University, as well as the TRIUMF laboratory and the Canadian Institute of Particle Physics. Questions concerning this document can be address to the ATLAS Canada Spokesperson Professor Robert Orr, University of Toronto (orr@physics.utoronto.ca), Deputy Spokesperson Associate Professor Robert McPherson, University of Victoria (robert.mcpherson@triumf.ca), or Computing Coordinator Professor Michel Vetterli, Simon Fraser University (vetm@triumf.ca). An electronic version of this document can be found at <http://atlas-canada.web.cern.ch/atlas-canada/documents/AtCanComp-PublicV1.pdf>.

Executive Summary

The next generation of particle physics experiments at the Large Hadron Collider (LHC) near Geneva, Switzerland, will produce enormous and complex data sets that will be analyzed in a world-wide computing system. Canada has played a critical role in the construction of the LHC and the ATLAS experiment, and exploitation of this investment requires that Canada deploy an integrated set of computing resources that will enable Canadians to lead the analysis of the ATLAS data and the extraction of some of the most significant scientific results in history. The computing and analysis facilities will need to process several Petabytes (10^{15} bytes) of complex data several times per year. The global ATLAS collaboration will have nearly 2000 physicists, including about 100 researchers in Canada alone. ATLAS Canada is a microcosm of ATLAS as a whole, including researchers from nine different Canadian universities and the TRIUMF laboratory with interests in widely differing areas of particle physics.

The ATLAS computing and analysis model has been developed by an international team of computing and physics experts. The model matches the sequential analysis pattern of particle physics experiments, where raw detector signals are first processed into “objects” that represent primary physics quantities (*e.g.*, electrons with a given energy and direction produced in an interaction), and then the objects are analyzed both for measurements of known processes and for searches for new physics (*e.g.* is there an anomalous peak in the distribution of the invariant mass of electron pairs?). These different steps in the analysis lead naturally to a set of “tiered” centres in the international computing system. The initial processing of the data, called “reconstruction”, is a continuous operation running 24 hours/day 7 days/week and requires facilities dedicated only to the ATLAS experiment, while the later detailed physics analyses are best performed on shared facilities. We know from our experience with large experiments that the resource needs of physics analysis have many “peaks” and “valleys” making shared-use facilities ideal. These experiments, including CDF and D0 at Fermilab near Chicago and Babar at SLAC near San Francisco, have already demonstrated successful use of existing shared CFI facilities. All users benefit from the resource sharing, since they have access to resource peaks above their average allocation.

The ATLAS Canada computing model is designed to give Canadians the ability to lead the many different ATLAS scientific analyses in which Canadians have interest. It includes a “Tier-1” centre dedicated to ATLAS and a distributed “Tier-2” computing federation utilizing shared facilities. The use of grid tools will allow the computing resources to function effectively together. Particle physicists, led by members of ATLAS Canada, pioneered the use of a Canadian computational grid, GridX1, for distributed analysis. GridX1 has already been used successfully for ATLAS data simulation exercises. Funding for the dedicated Tier-1 centre is being requested from the Canadian Foundation of Innovation (CFI) Exceptional Opportunities Fund, while the Tier-2 centres are being requested as part of the shared-use facilities in the CFI National Platforms Fund request.

The ATLAS experiment is our next step in the understanding of the basic building blocks of matter and the interactions among them. The analysis of the ATLAS data will be a daunting task, and Canadians are well positioned to lead the way. Both computing and human resources are critical to its success. This document describes the ATLAS Canada computing and analysis model.

Contents

1	Introduction to ATLAS	2
2	The ATLAS World-Wide Computing Model	3
3	The ATLAS Canada Computing Model	6
3.1	The Canadian Tier-1 Centre	7
3.1.1	Canadian Tier-1 functionality	8
3.1.2	Canadian Tier-1 resources	8
3.1.3	Canadian Tier-1 resource constraints	9
3.1.4	Canadian Tier-1 Personnel	9
3.2	The Canadian Tier-2 Federation	10
3.2.1	Canadian Tier-2 functionality	10
3.2.2	Canadian Tier-2 resources	10
3.2.3	Canadian Tier-2 resource constraints	12
3.2.4	Canadian Tier-2 Personnel	13
3.3	University and desktop capacity	13
3.4	Network	14
4	Summary	14
A	ATLAS-Canada Collaboration	16

1 Introduction to ATLAS

Particle physics is the study of the fundamental constituents of matter and the forces by which they interact. Particle physicists ask some of the most basic questions about nature, which have far-reaching implications for our understanding of the origin, current properties, and eventual fate of the universe. In particular, ATLAS at the CERN Large Hadron Collider (LHC) will collide high-energy protons allowing physicists to study nature at an unprecedented energy scale in pursuit of the understanding of the mechanism by which the elementary constituents of matter acquire mass. Fits to data from current experiments strongly indicate the existence of a new particle, called the Higgs boson, which performs this role. Theoretical considerations also demand additional new physics processes accessible to ATLAS, and the experiment has been optimized for their detection. Possibilities of new phenomena include supersymmetry, extra dimensions beyond the four space-time coordinates we know, and the search for substructure in particles we currently believe to be point-like. Discovery of any of these signs of new physics would revolutionize our understanding of the nature of the universe.

ATLAS is a collaboration of over 1800 physicists and engineers from 150 laboratories in 34 countries. The Canadian group, a founding member of ATLAS, includes scientists from the University of Alberta, the University of British Columbia, Carleton University, McGill University, Université de Montréal, Simon Fraser University, the University of Toronto, TRIUMF, the University of Victoria, and York University. The group consists of over 30 faculty members, about 20 engineers and technicians, 15 research associates and 25 graduate students. This represents about 30% of the Canadian experimental particle physics community and is expected to grow to nearly 50% by the time data taking starts in 2007.

The ATLAS detector must determine the direction, momentum, and energy of all the measurable products of the proton-proton collisions. The innermost parts of the detector measure the direction and momentum of charged particles (*e.g.* electrons), the central components measure the total energy of all particles that can be stopped in these components, and the outermost parts measure the direction and momentum of the few particles (typically particles called muons) that were not stopped. The Canadian ATLAS group concentrated its construction efforts on these energy-measuring components called “calorimeters”. These are massive detectors suitable for stopping the particles created by the proton collisions. These particles can traverse several metres of copper before finally giving up all their energy. ATLAS has five such detectors enabling it to fully measure nearly all the particles from the proton collisions. Canadians took a leading role in the design and construction of two of these massive detectors: one called the Hadronic End-Cap and another called the Forward Calorimeter. In addition to providing a significant fraction of these detectors, the group was part of the development and construction of the associated electronics, as well as the cryogenic feedthroughs that pass the signals from the calorimeters that reside in liquid argon (LAr) at -185°C to the room temperature electronics located outside the cryostat.

2 The ATLAS World-Wide Computing Model

The result of a proton-proton collision is commonly referred to as an event. For each ATLAS event, millions of detector elements must be read out, and these events occur almost 100 million times per second. The production of interesting events can be overwhelmed by more mundane events produced by well-understood physics processes. Sophisticated electronics, called the “trigger”, analyze in real-time the signals from the detector and pre-select certain events to be recorded to permanent storage, which greatly reduces the amount of data recorded; however, enormous amounts of data are still permanently recorded. ATLAS will collect about 3.2 Petabytes (3.2×10^{15} bytes) per year of data that will need to be stored and analyzed. In addition, several Petabytes of storage will be needed for secondary data sets.

An integral part of the analysis of a particle physics experiment is the production of simulated data, which are then run through the same analysis procedure as the real data. The real data are compared in detail with the simulated data to determine the physics processes that actually governed the proton-proton collision. In particular, differences between the real and simulated data may indicate that new physics processes are being observed.

The computing challenge is therefore to analyze several Petabytes of data, 1-3 times per year, as well as to produce and analyze the associated simulated data. An additional challenge is that the ATLAS detector will be taking data almost continuously for a decade or more, with maintenance periods of only a few months per year. Thus, processing of new data must be carried out simultaneously with the analysis of existing data sets.

CERN is coordinating an international network of high-performance computing centres to provide the resources necessary for the analysis of data from the LHC. This network will use grid tools to manage the data and to make efficient use of these world-wide resources. Because of the progressively refined nature of data reconstruction and analysis, a “tiered” structure of processing and analysis centres is being adopted, with a Tier-0 centre at CERN for initial data reconstruction, a network of Tier-1 centres for data reprocessing, a network of Tier-2 centres for large-scale data analysis and Tier-3 systems for final data visualization. A schematic of the LHC Computing Grid (LCG) is shown in Fig. 1.

Over 100 sites of varying sizes in 31 countries currently participate in LCG. ATLAS will have 10 Tier-1 centres and ~ 40 Tier-2 centres around the world. Canada will provide both Tier-1 and Tier-2 resources. The resources required by ATLAS for all Tier-1 and Tier-2 data centres is illustrated in Figures 2 and 3, respectively.¹ The units used are kSI2k for the CPU and Terabytes for the storage. SI2000, often abbreviated as SI2k, is a unit used to benchmark CPU performance for integer operations; kSI2k is 1000 SI2k. It has been found that this, rather than the floating point performance, is the best indicator of requirements for the analysis of particle physics experiments. For reference, a typical ~ 3 GHz Intel or compatible CPU corresponds to about 1 kSI2k.

The LCG resources will be contributed by the member countries on a fair-share basis. Canada’s minimum share is about 5% of the total ATLAS computing needs. The computing network built from each country’s minimum contribution will provide a data processing, reprocessing, simulation and analysis system for collaboration-wide use, with only the uni-

¹Technical details on the LCG can be found on the web at <http://lcg.web.cern.ch/LCG/>.

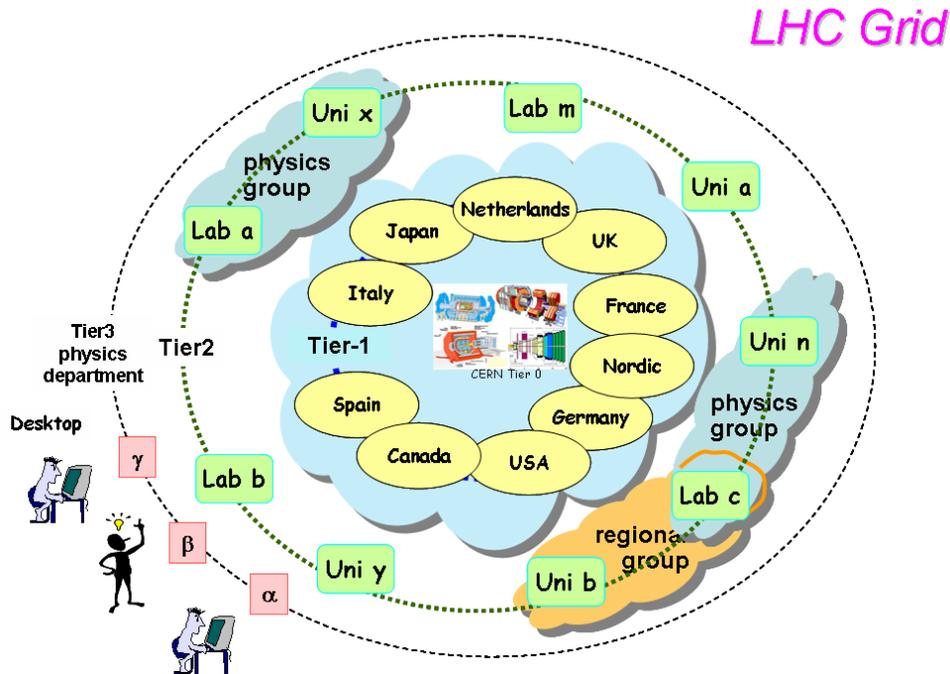


Figure 1: LHC Computing Grid (LCG)

versity Tier-3 layer requiring additional funding; however, performing early and pioneering analysis will require additional resources, and ATLAS Canada is designing a national computing model to provide these resources.

Extensive modeling and simulation of the ATLAS analysis chain have been done. Preliminary versions of the data analysis code exist and are being continually refined as more is learned about the actual performance of the detectors from detector simulation and from on-going analyses of data with prototype tests. Quantities such as data size and analysis time per event are therefore known to a reasonable degree of accuracy. A working group was charged with carrying out a detailed study of how analyses will be done in ATLAS based on simulation work and on experience at other experiments.²

The processing, reprocessing and simulation of ATLAS data will use many different data formats. A summary of the data types used for ATLAS is given in Table 1. The acronyms defined in this table will be used throughout this document. The RAW data recorded by ATLAS will be initially processed at the CERN-based Tier-0 centre and distributed to the ten Tier-1 centres world-wide. The RAW data are processed at the Tier-0 and Tier-1 centres to make the ESD. Each Tier-1 will host a share of the ATLAS RAW data and the corresponding set of ESD. Access to the ESD and RAW data will remain essential for detector calibration tasks and for developing early physics analyses; however, the ESD contains detailed detector data that is large and requires significant additional analysis to produce summary quantities

²The report of this group can be found at <http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/OO/computing-model/Comp-Model-January10.doc>

Cumulative Tier-1 Resources at all sites

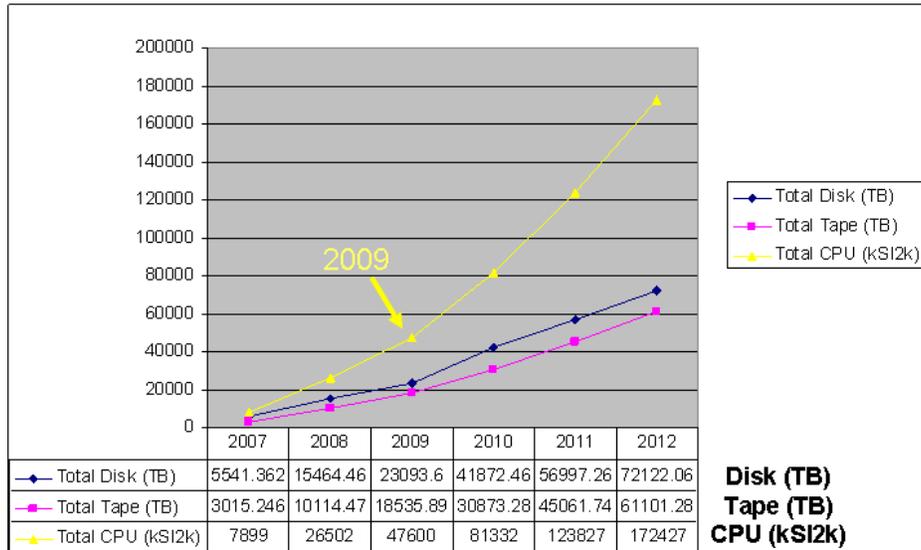


Figure 2: Resources required at all Tier-1 centres for ATLAS. The numbers are cumulative.

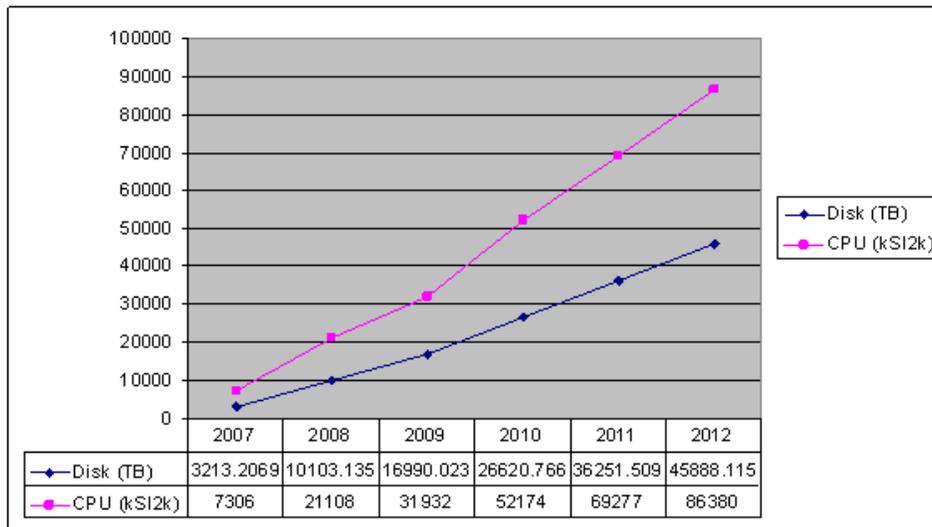


Figure 3: Resources required at all Tier-2 centres for ATLAS. The numbers are cumulative.

RAW	Real Raw Data, as recorded after the 3 rd Level Trigger
SIM	Simulated Data
ESD	Event Summary Data (after reconstruction)
AOD	Physics Analysis Object Data
DPD	Derived Physics Data (analogous with today's n-tuples)
TAG	Event tags, short event summaries primarily for event selection

Table 1: Acronyms used in the description of the ATLAS Data Structure.

useful for final physics analysis. The ESD are therefore further processed to the smaller-sized AOD, which will be used as the basis for most physics analyses.

3 The ATLAS Canada Computing Model

The ATLAS Canada collaboration must play a central role in the global ATLAS computing system while also addressing the additional needs of Canadian physicists analyzing the ATLAS data. It should be noted that in the first 2-3 years of an experiment significant progress is usually made in the early stages of the analysis chain because of the rapid development of the understanding of the detector. For ATLAS, this means that it will be crucial, especially in the early years, to be part of the Tier-1 system where RAW data and ESD can be accessed. The Canadian computing model is based on a dedicated facility providing our share of the Tier-1 resources for ATLAS, and the use of university facilities for the Tier-2 and Tier-3 layers.

The Canadian ATLAS computing and data analysis system is illustrated here through a set of typical use-cases:

- 1) **RAW and ESD data handling:** The RAW data for events surviving the ATLAS trigger system will be processed, when possible, on the Tier-0 system at CERN. RAW, and first-pass ESD, AOD and TAG data will be sent to the Tier-1 centres. Canada's Tier-1 centre will be responsible for hosting a share of the RAW and ESD data, plus additional samples of particular interest to Canadians. During ATLAS running there might be periods where the Tier-0 capacity is insufficient. In these cases, the Tier-1 centres will also perform initial reconstruction producing early ESDs. As better detector calibrations and reconstruction code become available, the Tier-1 centres will reprocess the RAW data to produce higher-quality ESDs. The ESDs will be further processed to build AOD and TAG data, which will be distributed to all ATLAS Tier-1 centres, and also to Canadian Tier-2 sites. The demands on a Tier-1 centre are exceptional. In addition to the processing and storage capabilities, the Tier-1 centre must continuously receive RAW and ESD data from CERN, as well as AODs from other Tier-1 centres. It must also distribute AODs to other Tier-1 centres and to the Canadian Tier-2 sites.
- 2) **Simulation-data handling:** One of the primary uses of the Tier-2 resources will be to produce the large amounts of simulated data needed for analyses. The simulation

is strongly limited by processing capacity, and the Tier-2 sites will have large CPU requirements. As the simulated data at the Canadian Tier-2 sites are produced, they will usually be copied to the Canadian Tier-1 centre for collaboration-wide use.

- 3) **Access to data on Canadian sites:** Canadian facilities will be used by the ATLAS Collaboration for analyzing the RAW, ESD, AOD and simulated data hosted in Canada. The RAW and ESD data, as well as the data simulated in Canada, will be accessed via the Tier-1 centre. The AOD data hosted in Canada for common ATLAS-wide analysis will be hosted at Tier-2 centres, which must also provide the corresponding processing capacity. AOD data for specific Canadian analysis use will be stored at, and analyzed with, additional resources at Tier-2 centres. These different types of data access place differing constraints on the sites hosting the corresponding data sets.
- 4) **Detector calibration:** Detector calibration will require direct access to RAW data. Canadians are responsible for the calibration of the detectors that we built. We are developing calibration techniques, in particular, for the ATLAS Liquid Argon (LAr) Calorimeters. This puts extra constraints on the Canadian Tier-1 centre, which must host additional LAr calorimeter data needed for detector calibration, and also supply significant additional CPU resources for their analysis by Canadian physicists.
- 5) **Physics analysis:** Individual Canadian physicists will typically use desktop systems for preparing and finalizing data analysis; however, every analysis will require significant access to Tier-2, and in some cases Tier-1, resources. Many analyses will be based on AOD data at the Tier-2 sites. However, initial analyses will also sometimes require access to ESDs, particularly in cases where the AODs are not optimized for specific analysis tasks. In such cases, the ability for Canadian physicists to run jobs directly on the Canadian Tier-1 centre will also be needed, requiring additional processing and storage capability at the Tier-1. Results from the different analysis stages will typically be collected as DPDs at storage elements dedicated to Canadian physics analysis at the Tier-2 sites. Some limited data sets (*e.g.* ntuples) will be stored on local systems at universities and laboratories. The results will be visualized using standard tools on desktop systems.

3.1 The Canadian Tier-1 Centre

Canada will have one Tier-1 centre in the ATLAS computing system, located at the TRIUMF laboratory in Vancouver. The specifications of this centre have been established by ATLAS Canada based on the ATLAS computing model. ATLAS Canada also strongly endorses the location of this centre and the corresponding support from the TRIUMF laboratory. Support for this unique facility is currently being sought through the CFI Exceptional Opportunities Program.

Note that the requirement and constraints on the Tier-1 resources are extreme. This can only be a success if the centre is dedicated to ATLAS-only use.

Capacity	2007	2008	2009	2010	2011	2012
CPU (kSI2k)	570	1910	3425	5855	8915	12415
Disk (TB)	400	1115	1660	3015	4105	5200
Tape (TB)	215	730	1335	2225	3245	4400

Table 2: Resources required for the Canadian Tier-1 centre. Numbers are cumulative.

3.1.1 Canadian Tier-1 functionality

The Tier-1 centre must perform a set of roles, serving both the ATLAS Collaboration and Canadian community. It must:

- Receive and archive 1/10th of the ATLAS RAW data as the primary host centre in the world for this data fraction.
- Receive, archive, and provide access to the ATLAS collaboration to approximately 1/5th of the ATLAS ESD using LCG middleware.
- Reprocess 1/10th of the ATLAS RAW data with new detector calibrations and reconstruction software, typically ~ 2 times per year, producing new ESD and AOD.
- Host simulated data produced by the Canadian Tier-2 resources for access by the ATLAS collaboration.
- Provide access to RAW and ESD data for Canadian physicists performing detector calibration tasks, and for new analyses for which the AOD is insufficient. This requires resources in addition to those necessary for centre's role in hosting the 1/10th and 1/5th fractions for ATLAS mentioned above.
- Host a copy of the entire ATLAS AOD set for access by Canadian Tier-2 sites.
- Maintain a complete replica of all ATLAS databases needed for data reconstruction and analysis. In some cases, this may require continuous two-way synchronization of the Canadian Tier-1 centre with the CERN Tier-0 and online systems, and the other Tier-1 centres.

3.1.2 Canadian Tier-1 resources

The Canadian Tier-1 centre will require significant and increasing CPU and storage capacity with the resource profile outlined in Table 2. Taking into account the periodic revisions to the ATLAS Computing model, we include a modest 25% contingency in our Tier-1 estimates. As outlined in the previous discussions it will also be essential to have capacity on the Tier-1 dedicated to Canadian detector calibration and analysis, and we have increased the minimum resource requirements by 20% to account for this. In addition, the network capacity from the Tier-1 must be large, with connectivity to many different locations, outlined in Table 3.

Location	Capacity
CERN Tier-0	3 Gbit/sec
Other Tier-1	1–2 Gbit/sec
Canadian Tier-2 sites	1 Gbit/sec

Table 3: Minimum network capacity required for the Canadian Tier-1 centre.

3.1.3 Canadian Tier-1 resource constraints

Successful operation of the Tier-1 centre will require that certain constraints on those resources be respected.

- 24/7 operations.
- Ability to run all versions of the ATLAS analysis software. In practice, this currently requires that the processing capacity must be in the form of Intel or compatible CPUs running a restricted set of releases of the Linux operating system. It is likely that similar constraints will exist in the future, requiring control of the computer operating system and other installed software.
- Ability to control resource use. In particular, it will be necessary to have control of the fractions of processing power and storage being allocated to different tasks. This will be particularly challenging at the Tier-1 centre, which will be running simultaneous reconstruction jobs, analysis jobs from ATLAS members, and additional Canadian analysis.
- Tight coupling between CPU and storage. In particular, during RAW data reprocessing and ESD data analysis, jobs can be limited by the ability to access storage elements efficiently from the CPUs running the jobs. The Tier-1 centre will be used to simultaneously receive new RAW, ESD and AOD data from CERN, reprocess existing RAW data to produce new ESD and AOD, distribute AOD to the Canadian Tier-2 sites, and analyze RAW and ESD data for detector calibration and physics analysis use.
- Ability to install, run, update and use grid toolkits for connection to the ATLAS/LHC world-wide grid, and additionally have the ability to maintain grid connectivity to the Canadian Tier-2 sites. This will require deployment of LCG middleware and opening the centre resources to the ATLAS Virtual Organization (VO).

3.1.4 Canadian Tier-1 Personnel

Approximately five people with hardware expertise will be needed to develop and maintain the Tier-1 centre. Included in this complement is one dedicated technician who will be responsible for routine swapping of faulty equipment and for dealing with vendors on warranty-related issues. The other four will be computing professionals. While all four will have knowledge of the whole system, they will be expected to have specialized knowledge of

one or two parts of the system, including CPU/OS, storage hardware and access, networking and security, database and conditions data management, and the LHC computing and other grid specialization.

In addition to the computing professionals, successful operation of the Tier-1 and utilization by Canadian physicists will require approximately four additional personnel with specific knowledge of ATLAS software and other analysis tools. These people will have a physics and analysis background, and will provide support for ATLAS and related software at the Tier-1 while also leading and supporting ATLAS analysis activities across Canada.

3.2 The Canadian Tier-2 Federation

Canada will also host ATLAS Tier-2 resources. In addition to our minimal fraction of the overall ATLAS Tier-2 resources used for central analyses within the framework of the ATLAS physics working groups, these resources will also be used for primary analyses by Canadian physicists. The Canadian Tier-2 resources are being requested as part of the upcoming CFI National Platforms Fund request. We request that the resources be split among the CLUMEQ, SciNet and WestGrid consortia.

The requirement on the CPU and storage are less extreme for the Tier-2 than the Tier-1. These facilities will fit well into shared-use models, so long as certain minimum constraints (*e.g.*, operating system version) are respected.

3.2.1 Canadian Tier-2 functionality

The Tier-2 federation must perform a set of roles, serving both the ATLAS Collaboration and Canadian community. It must:

- Simulate ATLAS events for the full ATLAS collaboration.
- Provide additional simulation capacity for Canadian analyses.
- Provide resources for common ATLAS-wide analysis.
- Provide additional resources for Canadian analysis, including hosting special AOD streams needed by Canadian groups.
- Provide grid-accessible storage for Canadian physicists. There will be some analysis jobs that will be distributed to the Canadian Tier-1, or even to non-Canadian sites, often to access simulated data from other centres. AOD and DPD (*e.g.* ntuples) produced by the distributed jobs will be copied to storage elements at the Canadian Tier-2 sites.

3.2.2 Canadian Tier-2 resources

The Canadian Tier-2 sites will require significant and increasing CPU capacity. Part of the Tier-2 resources in the ATLAS computing model are dedicated to production analysis passes through both AOD and simulated data managed by the ATLAS physics working groups. Canadians will participate in these groups, and these efforts represent a significant fraction

Capacity	2007	2008	2009	2010	2011	2012
CPU (kSI2k)	700	2025	3065	5010	6650	8290
Storage (TB)	300	970	1630	2555	3500	4405

Table 4: Resources requested for the Canadian Tier-2 federation. Numbers are cumulative. The Tier-2 storage is accessed via random access and therefore must be primarily disk-based.

Location	Capacity
Canadian Tier-1	1 Gbit/sec
Other Canadian Tier-2	1 Gbit/sec
Other Tier-1 and Tier-2	1 Gbit/sec

Table 5: Network capacity required for each Canadian Tier-2 centre.

of Canadian analysis needs; however, to be involved in early and pioneering analysis efforts, it will be completely essential to have dedicated Canadian analysis resources. The dedicated additional resources needed for pioneering Canadian physics analyses will be the equivalent of having resources for one extra physics working group, of which there will be about 20 in ATLAS, for Canadian use. This will double our minimum ATLAS Tier-2 fraction. The additional resources above our strict minimal ATLAS commitments using this “additional working group” model are consistent with the additional resource fractions that have given Canadian groups the ability to do competitive and pioneering analyses at the Tevatron³ experiments. The networking demands will also be somewhat lower than the Tier-1, as outlined in Table 5.

The distribution of the Canadian Tier-2 resources among the ATLAS Canada institutions is not rigorously fixed. The resources will be integrated into “federations” using grid tools. Canadian groups have demonstrated significant progress in the application of grid technology to building federations for ATLAS use. Alberta, NRC, Simon Fraser, TRIUMF, Victoria, and later McGill and Toronto, formed GridX1, a stand-alone Canadian computational grid. All of these resources were used to produce simulated data to test the ATLAS computing model, as well as to provide pseudo-data for the development of the analysis tools. The use of GridX1 was unique in that the whole Canadian grid showed up on LCG as a single node. Jobs were sent first to a gateway at TRIUMF where they were matched and then forwarded to the best available site on GridX1. This grid federation was the first successful exercise of its kind on such a scale, and gives us confidence in the use of distributed resources for the ATLAS Canada Tier-2 sites. The Tier-2 centres can be naturally grouped into distinct types.

- (1) Data analysis sites with significant processing and storage capabilities for common ATLAS-wide analysis. The CPU and disk must be tightly coupled for efficient analysis use, and full LCG middleware must be deployed for use by the ATLAS Virtual

³The Tevatron at the Fermi National Accelerator Laboratory near Chicago collides protons with anti-protons at the highest-energies currently achieved. The experiments are the closest existing analogies to ATLAS.

Consortium	Resources	
	CPU (kSI2k)	Disk (TB)
CLUMEQ	1650	500
SciNet	2500	1500
WestGrid	2500	1500
Total Cdn 2011 Tier-2	6650	3500

Table 6: The preferred distribution of ATLAS Canada Tier-2 resources being requested from the CFI National Platforms Fund integrated to 2011. The expected scope of the current CFI National Platforms Fund request is 2011, so the “Total” line shows the full ATLAS Canada Tier-2 request to CFI.

Organization (VO).

- (2) Data analysis sites with significant processing and storage capabilities for Canadian analysis. The CPU and disk must still be tightly coupled for efficient analysis use, but the sites have more flexibility in their grid toolkit deployment and need only be accessible to the members of ATLAS Canada.
- (3) Simulation sites with significant CPU capacity but smaller storage requirements. Because the simulated data will normally be copied to the Tier-1 centre for analysis access, the grid middleware requirements on simulation sites are not as severe as for the analysis sites.

These different types of centres not only have different hardware requirements, but will require different levels of personnel support. In deciding the optimal distribution of Canadian Tier-2 resources, local research group size, university commitments to space and other infrastructure, and computing centre management experience at the different institutions all require consideration. The ATLAS Canada collaboration has reviewed the capabilities of our different member institutions, which are in turn members of the different computing consortia. Our preferred distribution of the ATLAS Canada resources being requested from the CFI National Platforms Fund for 2011 is summarized in Table 6. We show the 2011 breakdown because that is the expected scope of the current CFI National Platforms Fund request.

3.2.3 Canadian Tier-2 resource constraints

Successful operation of the Tier-2 centres will require that certain constraints on those resources be respected.

- Efficient operation for both simulation and analysis.
- Ability to run all versions of the ATLAS analysis software. In practice, this currently requires that the processing capacity must be in the form of Intel or compatible CPUs running a restricted set of releases of the Linux operating system. It is likely that similar constraints will exist in the future, requiring control of computer operating system and other installed software.

- Ability to control resource use. In particular, it will be necessary to have control of the fractions of processing power and storage being allocated to simulation and to Canadian analysis tasks.
- Tight coupling between CPU and storage will be required for AOD analysis, but is less essential for simulation where relatively small volumes of simulated data are created for a given CPU usage.
- The Tier-2 resources for common ATLAS-wide analysis will have to be in cluster(s) that implement the full LCG middleware, and are open to the full ATLAS Virtual Organization (VO). The Tier-2 resources for Canadian analysis will have to be in cluster(s) that implement grid connectivity accessible to all ATLAS Canada members. The Tier-2 resources used for simulation could be accessible primarily via a gateway such as GridX1 (see Section 3.2.2), since it is foreseen that central ATLAS simulation will be performed by a limited number of production managers.

3.2.4 Canadian Tier-2 Personnel

Sufficient personnel with hardware expertise will be needed to develop and maintain each Tier-2 site. The number of personnel will depend on the size of the centre, but several experts will be required at each site. It is expected that many of the system support personnel will be shared with other CFI and university facilities, and receive a combination of CFI and university support. The sites will also require experts in ATLAS application support, including tasks like the installation of ATLAS software, maintenance of grid tools and database installation and replication. Both system and application support expertise can, to some extent, be shared among the different centres.

3.3 University and desktop capacity

Individual physicists will work at workstations, monitors or laptop computers, much as they do today. While the main AOD and ESD access will occur at Tier-1 and Tier-2 facilities using that CPU and storage, final analyses will still typically be “driven” from university physics department clusters and desktops. These systems will also be used for the visualization of both individual events and physics distributions of ensembles of events at different analysis stages. Physicists will require some local computing capacity to run ATLAS software and other standard analysis programs, including storage for limited sets of ntuples, histograms, or other derived physics data sets. The precise setups used by individual physicists will vary widely, and will likely include single powerful workstations, multi-cpu systems with common storage shared by faculty members and students at single universities, and other similar models. Typical local institute-based resources will be about 1 Terabyte of storage, and processing power equivalent to a few current computers per active physicist. A common feature of all models is that it will be necessary for users to submit analysis jobs using grid tools to Tier-2 sites, and receive some type of DPD back to their university centre from the Tier-2 sites. We foresee about 100 active physicists (faculty members, postdocs, students) analyzing ATLAS data across Canada, and the sum of these resources is significant.

3.4 Network

The network requirements of the Tier-1 and Tier-2 sites were shown in the corresponding sections. These centres depend critically on dedicated high-speed connections from the Tier-1 centre to CERN, between Canada's Tier-1 centre and other ATLAS Tier-1 centres, and between the Canadian Tier-1 and Tier-2 sites. The network between Canadian universities and the Tier-2 sites is also critical for efficient Canadian physics analysis.

HEPNET/Canada is responsible for the national and international network for the high energy physics community. The current Director of HEPNET/Canada is R.Sobie at the University of Victoria. HEPNET/Canada has a technical manager, and is funded by an NSERC MFA grant.

The CA*net 4 network is Canada's national research and education network which is operated by CANARIE through funding provided by the Government of Canada. The underlying architecture of CA*net 4 is three 10 Gbps lambdas from Halifax to Vancouver. International connections are made at the Pacific Wave in Seattle, the STAR LIGHT in Chicago and the MAN LAN in New York City.

The TRIUMF Tier-1 facility will be connected to CERN via a 10 Gbps lightpath network that will be established in early 2006. TRIUMF is currently participating in the LCG Service Challenges where large-scale data transfers are being tested over extended periods. TRIUMF will use the 10 Gbps lightpath for the Service Challenges in 2006 and 2007. The TRIUMF-CERN connection is provided by CANARIE and HEPNET/Canada.

The Canadian Tier-2 facilities will be connected by 1 Gbps lightpaths to TRIUMF. Development of the first lightpath link between UVic (HEPNET/Canada) and TRIUMF is currently underway (early 2006) will be used to learn how to connect the Tier-2 sites. Once this lightpath connection is in place, then the other Canadian Tier-2 sites will be integrated into TRIUMF. Figure 4 shows the proposed Canadian HEP network.

4 Summary

The ATLAS Canada computing model is designed to allow the full Canadian participation in the reconstruction and analysis of the most exciting data foreseen in particle physics. The challenge is enormous, with the simultaneous acquisition of several peta-bytes of data per year, detailed simulation of the data, reconstruction of current data alongside further reconstruction of earlier data sets, and the development and execution of final analysis. Approximately 100 Canadian faculty members, research associates and students will be part of this effort. The computing model is built around the use of relatively inexpensive Intel or compatible CPUs and storage systems, using large serial clusters interconnected with grid tools and techniques. While the system will clearly evolve as technologies evolve, it will allow Canadian physicists to lead efforts within ATLAS extracting physics results from the earliest data-taking.

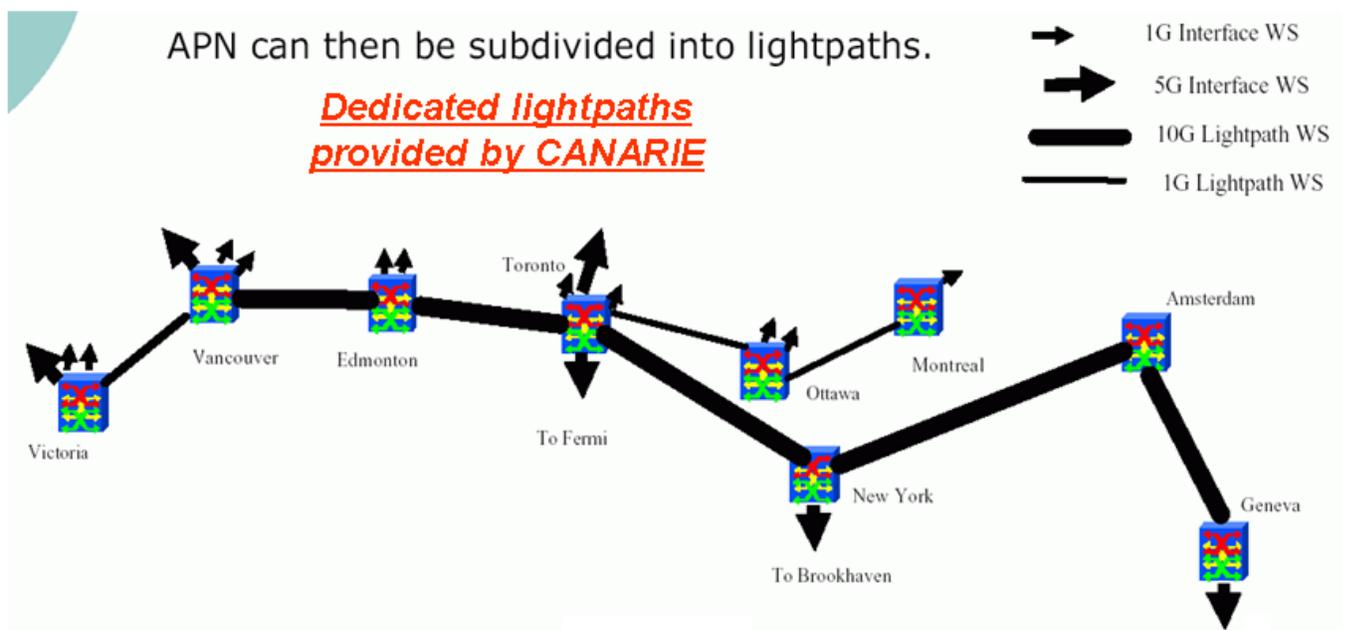


Figure 4: The Canadian network configuration used for the ATLAS Canada computing environment.

A ATLAS-Canada Collaboration

Faculty	
Alan Astbury	Victoria
David Asner	Carleton
David Axen	UBC
Georges Azuelos	Montréal/TRIUMF
David C. Bailey	Toronto
Sampa Bhadra	York
François Corriveau	McGill/IPP
Gilles Couture	Montréal/UQAM
Matt Dobbs	McGill
Douglas Gingrich	Alberta/TRIUMF
Richard Keeler	Victoria
Robert V. Kowalewski	Victoria
Peter Krieger	Toronto
Leonid Kurchaninov	TRIUMF
Michel Lefebvre	Victoria
Claude Leroy	Montréal
Mike Losty	TRIUMF
Jean-Pierre Martin	Montréal
Robert McPherson	Victoria/IPP
Roger W. Moore	Alberta
Gerald Oakham	Carleton
Dugan O'Neil	SFU
Chris Oram	TRIUMF
Robert S. Orr	Toronto
James Pinfold	Alberta
Steven Robertson	McGill/IPP
Pierre Savard	Toronto/TRIUMF
Pekka Sinervo	Toronto
Randy Sobie	Victoria/IPP
Reda Tafirout	TRIUMF
Wendy Taylor	York
Richard Teuscher	Toronto/IPP
Isabel Trigger	TRIUMF
William Trischuk	Toronto
Brigitte Vachon	McGill
Michel C. Vetterli	SFU/TRIUMF
Manuella Vincter	Carleton
Andreas Warburton	McGill

Collaborators

Sergey Chekulaev	RA TRIUMF
Pierre-Antoine Delsart	RA Montréal
Margret Fincke-Keeler	RA Victoria
Petr Gorbounov	RA Toronto
Yoshio Ishizawa	RA TRIUMF
Mohsen Khahkzad	RA Carleton
Christophe Le Maner	RA Toronto
Jiansen Lu	RA Alberta
Rachid Mazini	RA Toronto
Rashid Mehdiyev	RA Montréal
Chris Potter	RA McGill
Rolf Seuster	RA Victoria
Richard Soluk	RA Alberta
Kai Voss	RA Victoria
Rod Walker	RA SFU
Zhaoyu Yang	RA Carleton

Graduate Students

Cristen Adams	M.Sc. Toronto
J.-P. Archambault	Ph.D. Carleton
Camille B.-Champagne	Ph.D. McGill
Frank Berghaus	Ph.D. Victoria
Marco Bieri	M.Sc. SFU
B. Brelief	Ph.D. Montréal
Sebastien Charron	M.Sc. Montréal
S.L. Cheung	M.Sc. Toronto
Claudiu Cojocaru	Ph.D. Carleton
Robert Dumoulin	Ph.D. Toronto
Keith Edmonds	M.Sc. Victoria
Jonathan Ferland	M.Sc. Montréal
Marie-Helene Genest	Ph.D. Montréal
Yan Guo	M.Sc. Toronto
Sang Hee Han	M.Sc. Alberta
Louise Heelan	Ph.D. Carleton
Ahmed Hossain	Ph.D. Alberta
J. Idarraga	Ph.D. Montréal
Tayfun Ince	Ph.D. Victoria
Gustavo Kertzsch	M.Sc. McGill
J.Klamka	M.Sc. Toronto
Céline Lebel	Ph.D. Montréal
Kalen Martens	Ph.D. Toronto
Victoria Martynenko	Ph.D. York
Audrey Mcleod	M.Sc. McGill
Erfan Rezaie	M.Sc. SFU
Gabriel Rosenbaum	M.Sc. Toronto
Warren Shaw	M.Sc. Victoria
Doug Schouten	M.Sc. SFU
Malachi Schram	Ph.D. Carleton
Logan Sibley	M.Sc. Alberta
Yushu Yao	Ph.D. Alberta
Wei-Yuan Ting	Ph.D. Alberta
Dan Vanderster	Ph.D. Victoria

Engineering and Technical Infrastructure

Ashok Agarwal	Computing Technologist Victoria
Paul Birney	Senior Technician UVic/TRIUMF
J. Berichon	Technician Montréal
Mircea Cadabeschi	Engineer Toronto
Bryan Caron	Scientist Alberta/TRIUMF
Wen Chao Chen	Computer Infrstrctr. Montréal
Denice Deatrich	Large cluster manager TRIUMF
Howard Peng	System Manager Victoria
Phillipe Gravelle	Technician Carleton
Leslie Groer	Computer Infrstrctr. Toronto
Alisa Dowling	Engineer UVic/TRIUMF
Robert Henderson	Scientist TRIUMF
Lars Holm	Sr. Electronic Tech Alberta
Wade Hong	Network Infrstrctr. Carleton
W. Jack	System Manager Carleton
Roy Langstaff	Engineer UVic/TRIUMF
Mark Lenckowski	Designer/Draftsman UVic/TRIUMF
Drew Price	Elec. Tech. Alberta
Paul Poffenberger	Scientist Victoria
Bill Roberts	Electrical Engineer TRIUMF
Yun-Ha Shin	Computer Infrstrctr. Carleton
Jan Soukup	Engineer Alberta
Kenneth Vincent	Technologist Toronto
Peter Vincent	Technician TRIUMF
Len Wampler	Elect. Tech Alberta
