Growing Advanced Networking Capacities Drive Canada's Innovation Capabilities

A Submission to the Government of Canada's Digital Economy Consultation
Reda Tafirout, Steven C. McDonald, and Timothy I. Meyer on behalf of TRIUMF
Tel: 604-222-7674 | E-mail: tmeyer@triumf.ca

Theme: Digital Infrastructure

Question: What speeds and other service characteristics are needed by users (e.g., consumers, businesses, public sector bodies, and communities) and how should Canada set goals for the next generation networks?

Abstract

TRIUMF, Canada's national laboratory for particle and nuclear physics, is host to a Tier-1 data and computing centre for the ATLAS experiment based at CERN in Geneva, Switzerland. ATLAS is one of the largest collaborative efforts ever attempted in the physical sciences: an international collaboration of over 2,000 physicists from 38 countries, and 174 laboratories and universities. Canada is a major player in the project and the success of the scientific program heavily relies on a robust and state-of-the-art networking infrastructure. This innovation-driven project not only leverages Canada's installed network capacities to enable world-class research for hundreds of Canadian scientists and students, but it also serves as a training ground for next-generation technologies such as cloud computing.

Discussion

The ATLAS experiment at the Large Hadron Collider (LHC) complex, located at CERN in Geneva, Switzerland, collects an enormous amount of data at an unprecedented scale. The ATLAS detector records proton-proton collisions at the highest energies in the world to explore the fundamental nature of matter and the basic forces that shape our universe. The experiment operates around the clock and produces several petabytes (millions of gigabytes) of data each year. For efficient and coordinated access to the data, the ATLAS Computing Model was developed to enable the 2,000 scientists around the world easy and reliable access to the complete datasets. It is based on a set of tiered computing centres which are connected via high-performance and dedicated networks: the Worldwide LHC Computing Grid (WLCG). The primary data is stored initially at the Tier-0 centre at CERN, and then distributed to ten Tier-1 centres around the world. Each Tier-1 centre has custody of a share of the data and is responsible for further data reprocessing and distribution. The data is further disseminated to the Tier-2 centres that are associated with each Tier-1 for collaboration-wide access and data analysis. In Canada, TRIUMF hosts the Tier-1 centre and there are five Tier-2 centres that are located at: The University of Victoria, Simon Fraser University, The University of Alberta, The University of Toronto and McGill University.

The Tier-1 centre was funded through a proposal led by Simon Fraser University to the Canada Foundation for Innovation in 2007.

To support the distribution of the data amongst the various tiered sites a robust, reliable, high bandwidth, long distance network needed to be established between CERN and the Tier-1 centers, as well as the Tier-1 and its associated Tier-2 centres. This tiered network model has been very successful. The data is generated in massive quantities by the ATLAS experiment at CERN's Large Hadron Collider and distributed to 10 Tier-1 sites around the world. The Tier-1 sites further distribute the processed physics data to the Tier-2 sites. In addition the Tier-2 sites send simulated Monte Carlo data back to the Tier-1 centers.

All of the above have lead to an increase in Canada's and TRIUMF's recognition as a key player and its ability to contribute to large scale science, Grid computing, and research at an international level.

This flow and distribution of the data has worked well during the commissioning phases. More recently, now that the LHC has stated producing collision data, physicists around the world are clamouring to examine the data. Consequently, a different pattern of data exchange is emerging. Everyone wants to exchange data directly with everyone else. The more structured, tiered model may now become a potential impediment. As a result a new model that involves a small number of network exchanges, or peering points, is being examined. The proposal is to locate a network exchange at each of the following locations, the USA, Europe and CERN. Tier-1, -2 and -3 sites in these regions would peer directly with their respective continental exchange. This allows for a more direct method of data exchange between the sites worldwide.

Presently, Canada's Advanced Research and Innovation Network (CANARIE) is providing 8 end-to-end light-paths (LP) to support the LHC/ATLAS project totaling 12 Gbps. For the TRIUMF Tier-1 connectivity: 5 Gbps primary for TRIUMF-CERN link along with a 1 Gbps secondary link with automatic failover, 1 Gbps for TRIUMF-BNL (Brookhaven National Laboratory, U.S. ATLAS Tier-1) which is also configured as a tertiary link for TRIUMF-CERN (for failover across the Atlantic), 1 Gbps for TRIUMF-SARA/NIKHEF (Netherlands ATLAS Tier-1) for Tier-1 peering according to the ATLAS computing model. For the Tier-2 connectivity with the TRIUMF Tier-1: 1 Gbps dedicated to each site, except for Simon Fraser light-path which does not run over the CANARIE infrastructure; because of its proximity to TRIUMF, the link runs over the local BCNET provider fiber infrastructure.

Canada is well positioned to take advantage of this new model. TRIUMF has already established direct peering with its Canadian Tier-2 sites. Furthermore, it already has direct peering with CERN, a site in Europe and the USA that are the likely candidates for the network exchanges. If such a model receives approval, it would be in the interest of TRIUMF and ATLAS Canada to request a fourth exchange point in Canada. It could potentially be as simple as upgrading the capacities of the existing infrastructure from 1 Gbps to 10 Gbps.

Without the CANARIE-funded light-paths (LPs) program, there would have been no affordable alternatives possible to establish an LHC ATLAS Tier-1 Data Centre in Canada.

The commercial Internet could not have provided the necessary capacity and security requirements for an LHC/ATLAS Tier-1 at TRIUMF. The procurement of dedicated LPs from a commercial provider would have been financially unaffordable for TRIUMF and its Canadian Tier-2 University institutions. Furthermore, the experience, flexibility and willingness of the CANARIE organization and its staff to accommodate unique network requirements for research have been invaluable. CANARIE's involvement and participation with international network partners in the U.S., Europe, and Asia-Pacific makes collaborative research with institutions in these regions a much easier process. Canadian scientists and students are just beginning to reap the benefits of the foresight of CANARIE; modern high-performance computing will continue to evolve in a truly globalized context.

In the coming years, it is anticipated that the networking demand will grow significantly as the ATLAS project evolves and expands. It will be crucial for CANARIE to remain a key player by provisioning high-bandwidth dedicated light paths in order for Canada to remain competitive and a leader on the global stage.

As is clear with the ATLAS Tier-1 Data Centre, not only must high-performance computing capacity evolve with time to maintain its utility and competitiveness, but so also must the networking capacity. The Tier-1 Data Centre serves as a centralized hub for Canadian access to the scientific riches of the ATLAS global particle-physics project; its relevance to the global effort is defined by its networked international access and its importance and utility to the Canadian university community is determined by its nationally networked access. As discussed, there are new models of connectivity emerging and Canada's ability to select, participate, and contribute to the next-generation of network speeds and connectivity will be critical.