Canadian Networks for Particle Physics Research

2010 Report to the Standing Committee on Interregional Connectivity, ICFA Panel
January 2010

This report describes the status and plans of the Canadian network infrastructure used for particle physics research in 2010[1]. We describe the status of the CANARIE network infrastructure. The ATLAS Tier 1 Centre at TRIUMF and planned Tier 2 upgrades for 2011 are highlighted as key projects using high-speed networks.

CANARIE Network

CANARIE Inc. is Canada's Advanced Research and Innovation Network. Established in 1993, CANARIE manages an ultra high-speed network, which facilitates leading-edge research and big science across Canada and around the world. More than 40,000 researchers at over 225 Canadian universities and colleges use the CANARIE Network, as well as researchers at institutes, hospitals, and government laboratories throughout the country. Together with 12 provincial and territorial advanced network partners, CANARIE enables researchers to share and analyze massive amounts of data, like climate models, satellite images, and DNA sequences that can lead to groundbreaking scientific discoveries. CANARIE is a non-profit corporation supported by membership fees, with the major investment in its programs and activities provided by the Government of Canada.

CANARIE Optical Infrastructure

The CANARIE optical infrastructure is comprised of two different optical technologies; SONET and ROADM.

The lower optical layer is the DWDM layer, which was built from the latest ROADM technology. Two ROADM networks, depicted in the thick blue lines in Figure 1, were built in 2005 and 2006. The first ROADM project, called the "Eastern Canadian ROADM", established a 2700 km, multi-degree optical mesh network along the busiest corridor of the previous generations of CANARIE Networks: Chicago – Toronto – Ottawa – Montreal – New-York City. The second ROADM project, called the "Western Canadian ROADM", established a 1500 km, multi-degree optical mesh network from Seattle – Victoria – Vancouver – Kamloops – Calgary, with a spur from Kamloops to Kelowna. The ROADM hardware is 40G and 100G ready as CANARIE partnered with Ciena in October 2010 to demonstrate the 100G capability bringing up a wavelength from Ottawa to Chicago.

The second layer of the optical network infrastructure is the SONET layer. In this layer, the network is composed of individual point-to-point 10G SONET wavelengths linking the network with a topology consisting of a series of rings. The ring architecture allows CANARIE to provision redundant network services, as in the case for the CANARIE IP network. Wavelengths are added to the network incrementally, based on current and medium-term projections for demand. There are three types of SONET wavelength services in the network; CANARIE-lit wavelengths, leased wavelengths, and wavelength swapping. CANARIE-lit wavelengths are possible because of the ROADM networks.

CANARIE Network Services

CANARIE offers numbers of network services to users. Lightpath and IP Network services are traditional offerings for many years. In 2010, CANARIE deployed IPv6 commercial service, and also started working on a new initiative, Content Peering Service. These new services would prepare ways to access to future networking services and knowledge.

[1] Additional information can be obtained by contacting Dr. R. Sobie, Director of HEPNET/Canada (rsobie@uvic.ca)

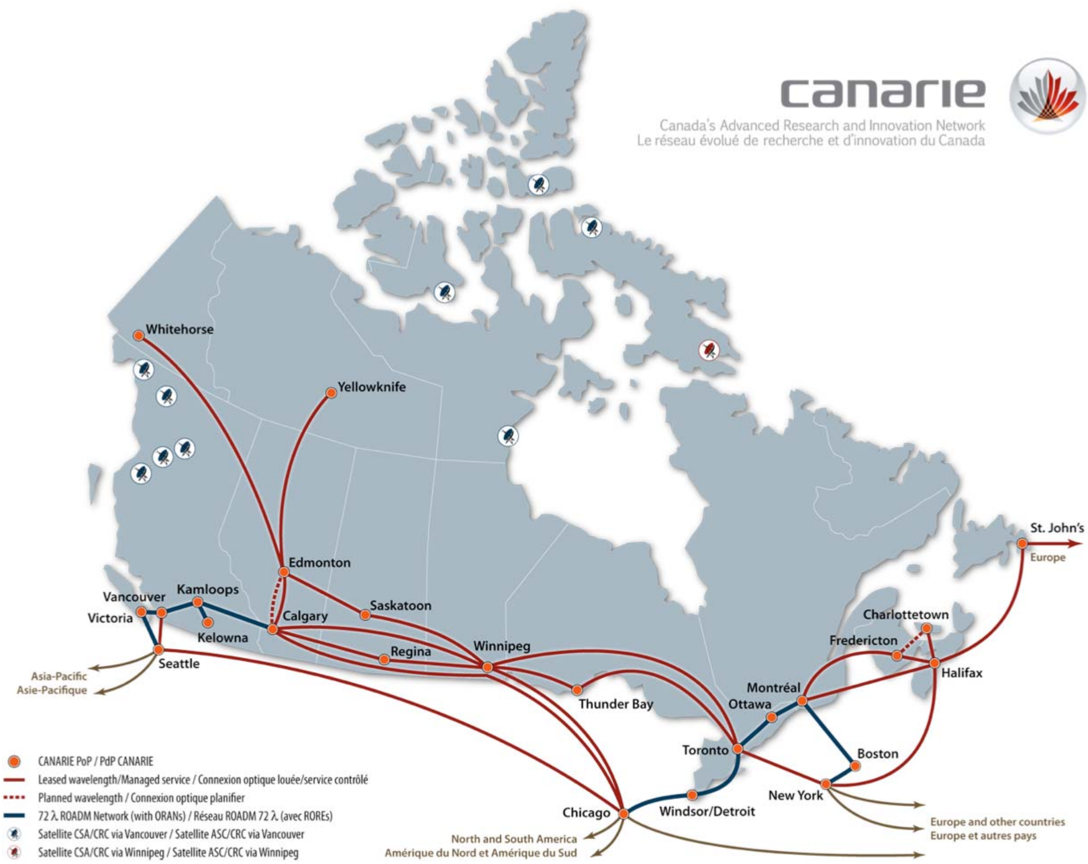

Figure 1: The CANARIE network

CANARIE Lightpath service

CANARIE Lightpath service is delivered on GbE and 10 GbE client interfaces over SONET or a dedicated wavelength. 10Gps SONET wavelengths are usually partitioned into smaller capacity channels, from 155 Mbps to GbE, and up to a full 10 GbE. With the new ROADM networks, lightpaths can also be 10GbE wavelengths that can be dropped directly into researchers' equipment thus bypassing CANARIE optical switches or routing equipment.

CANARIE IP Network service

A number of lightpaths on CANARIE's optical infrastructure are used to provide traditional IP services offering full and equal support for IPv4 and IPv6 unicast and multicast routing. Internally the IP network is comprised of five major routing nodes, which are located in Calgary, Winnipeg, Toronto, Montreal, and Halifax. A sixth smaller node was in Edmonton to aggregate IP network traffic for the Yukon Territory and Northwest Territory GigaPoPs. Seven internal segments link the routers in a partial-mesh topology and five external network segments which extend to international R&E layer 2 exchanges: the Pacific Wave in Seattle, StarLight in Chicago, and Manhattan Landing (MANLAN) in New York. Most of these external segments are operated at 10GE.

IPv6 Commercial Peering

Traditionally, CANARIE carries only R&E IPv4 and IPv6 unicast and multicast traffic. Institutions are required to have their own Internet connection for commercial traffic. CANARIE has been supporting

IPv6 routing exchanging IPv6 traffic with other R&E networks since 2002. As IPv6 routing is becoming increasingly important as the world runs out of IPv4 addresses – estimated to occur in early 2011, IPv6 connectivity is becoming critical to Canadian research institutions as some networks can only be accessed through IPv6, and lack of good IPv6 connectivity causes issues to various University Campus departments as new platforms are primary IPv6 platforms. The adoption of IPv6 by Canadian commercial service providers has been slow; in fact, only a small number of commercial ISPs in North America are currently offering IPv6 as a standard service. Most of the Canadian Regional Advanced Networks (RANs) and universities are unable to set up IPv6 peering with commercial ISPs. CANARIE saw an opportunity to support the IPv6 networking needs of the Canadian research community. In 2010, CANARIE acquired connections from a few IPv6 providers and opened its door carrying full IPv6 Internet routing to Canadian institutions.

Content Peering Service
Canadian researchers and institutions have an increasing need for fast access to content providers like Amazon and Google as their research activities involve cloud computing services more and more. Providing faster access to these clouds and services will significantly increase overall efficiency. CANARIE, as a R&E backbone provider, operates a nationwide network with a wide range of international connections, is well positioned to deliver these kinds of services to users. In 2010, funding was approved to develop and deploy Content Peering Service by July 2011. This service aims to deliver high-speed access to content providers like Amazon, Microsoft and Google. CANARIE will be working with RANs to deliver this service to all Canadian R&E institutions.

International collaboration

The CANARIE Network is connected to GLIF (the Global Lambda Integrated Facility) through three International eXchange Points (IXPs): Pacific Wave, StarLight, and MANLAN (Manhattan Landing Exchange Point). This connectivity provides CANARIE with the capacity to create dedicated lightpaths throughout most of the world, enabling researchers to perform collaborative research.

GLIF is a technical forum where lightpath-enabled network resource providers work together to develop inter-operable international infrastructure. CANARIE is a funding member of GLIF and has volunteered in creating numbers of operation papers and procedures to better define the needs of the next generation GOLE architecture and equipment. In the past year CANARIE participated in the Automated GOLE taskforce that organized a demonstration during the GLIF meeting at CERN.

GLORIAD is built on a fiber-optic ring of networks around the northern hemisphere of the earth, providing scientists, educators and students with advanced networking tools that improve communications and data exchange, enabling active, daily collaboration on common problems. Its aim is to provide a stable, persistent, non-threatening means of facilitating dialogue and increasing cooperation between nations that have often been at odds. CANARIE, is a member of GLORIAD, along with partners from China, Russia, South Korea, the United States, the Netherlands, and the Nordic countries, has been providing and continues to provide important lightpath services linking research networks across Asia, North America and Europe. Allowing scientific communities in three continents moves a large volume of valuable data effortlessly, streams video and communicates through quality audio- and video-conferencing.

ATLAS Tier 1 Computing Centre at the TRIUMF Laboratory

TRIUMF, Canada's National Laboratory for Nuclear and Particle Research, has built a Tier-1 (T1) Computing Centre for the ATLAS experiment in Canada. The TRIUMF Centre is linked to the LHC Worldwide Computing Grid (WLCG) and provides an interface to a grid of computing resources at universities across Canada.

In July 2005, CANARIE signed a Memorandum of Understanding (MOU) with HEPnet/Canada, ATLAS Canada and TRIUMF to provide the high-energy physics community with a dedicated 10G APN across Canada and initial 5G lightpath to the CERN Tier-0 (T0) Centre. This lightpath became active in December 2006.

The TRIUMF T1 to CERN T0 circuit, depicted in Figure 3, runs over the CANARIE infrastructure until it disembarks North America in New York City. Each T1 site must use a small or series of small publicly routable Classless Inter-Domain Routing (CIDR) blocks as only traffic from the Large Hadron Collider Private Optical Network (LHCOPN) address space is allowed to flow over the network. Exchange of routing information is performed using Border Gateway Protocol (BGP) at the T1 and T0 institutions. The 5G lightpath transits Canada west to east. The lightpath travels over the BCNet network from TRIUMF to the CANARIE PoP at UBC, and continues along CANARIEs network, then debarks North America at the MANLAN transit exchange in New York City. The lightpath enters Europe on SURFnet in Amsterdam and then transits Geant2 network to CERN.

The TRIUMF T1 will hold 4.3% of the ATLAS data, and it is anticipated that a 5G link will be sufficient for the first few years of LHC operation. However should the demand increase beyond 5G, the lightpath capacity can be increased in 155Mbps increments up to the full 10G.

The backup link for the TRIUMF T1 passes from the Vancouver CANARIE OME to the Pacific Northwest Gigapop in Seattle, then via Chicago, Toronto and MANLAN to Amsterdam. This link provides an alternate fibre path across most of North America and the Atlantic. Even at a lower capacity of 1G it is expected to be able to temporarily handle loads while the 5G link is restored. The tertiary backup link added travels via Victoria to Pacific Northwest Gigapop and then to the Brookhaven National Laboratory (BNL). In the event of the failure of the primary and secondary links traffic will be carried via the US LHC Network to CERN. The tertiary backup also acts as T1 to T1 link which is particularly advantageous because BNL will host 25% of the ATLAS Data. In 2008 an additional T1 to T1 link was established with SARA (TRIUMF's Tier-1 partner) following the same path as the primary 5G link.

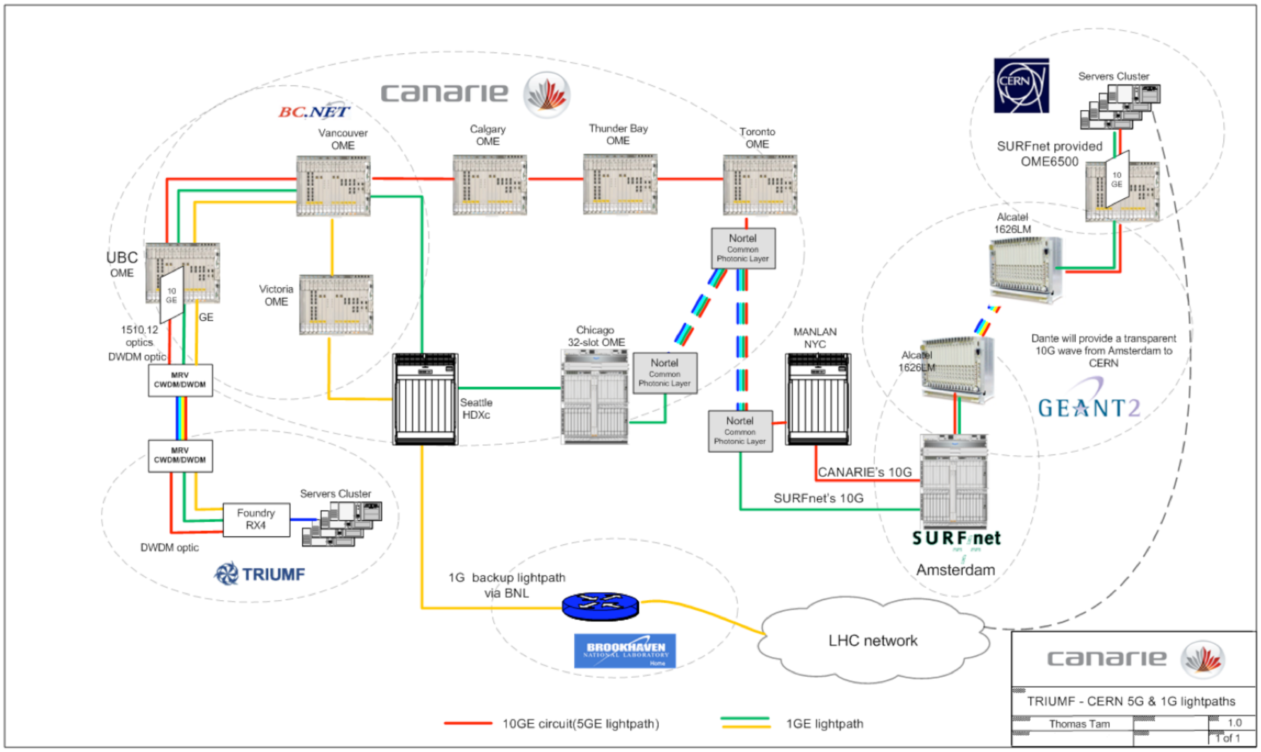

Figure 3: TRIUMF – CERN 5G and 1G lightpaths

Canada hosts Tier 2 centres at the University of Victoria, University of Alberta, University of Toronto, Simon Fraser University and McGill University. The T1 to T2 connections in Canada follow the successful LHCOPN model with TRIUMF at the centre of a star pattern. Four of the T2 institutions have a 1 G lightpath to TRIUMF while Simon Fraser University shares a 1 G circuit with WestGrid. Each path, excepting Simon Fraser University, is carried on CANARIE ROADM network.

In Canada a T2 to T1 connection either transits an Optical Regional Advanced Network (ORAN) before arriving at a CANARIE POP to be carried to TRIUMF or connects directly to a CANARIE POP. The largest challenge for the Canadian T2's has been establishing a connection with either an ORAN POP or CANARIE POP. In several cases the T2 must employ Course Wavelength Division Multiplexing (CWDM) over existing fibre to arrive at an appropriate POP. This equipment can be expensive and pose a serious financial hurdle for a T2. In addition not all T2's have routing equipment capable of supporting BGP, which is requirement for an ATLAS T2.

In addition TRIUMF has taken on a Tier 1 role for University of Melbourne Tier 2. All traffic from this institution arrives via TRIUMF regular general purpose 1G ORAN route. The addition of this traffic to the route has brought to near to saturation on several occasions necessitating some of the upgrades discussed in the following section.

Upgrades for 2011

Traffic patterns for the Canadian ATLAS Tier-2s in 2010 show signs of saturation. Figure 4 shows a typical traffic pattern from May of 2010 for the University of Toronto in which the circuit was either near its maximum capacity or essentially off. In order to address these growing demands we plan to upgrade a number of the Tier-2s to 10G lightpaths in early 2011. In addition, we will be upgrading TRIUMF's primary general purpose ORAN link to 10G.

Both the University of Victoria and Simon Fraser University are members of the WestGrid HPC consortium and have their ATLAS Tier-2 facilities located inside the WestGrid clusters. WestGrid is connected with a dedicated set of 10G lightpaths. In the first 2 months of 2011, TRIUMF will upgrade its connection to the WestGrid network to 10G thus allowing the University of Victoria and Simon Fraser to use these lightpaths for Tier-2 traffic.

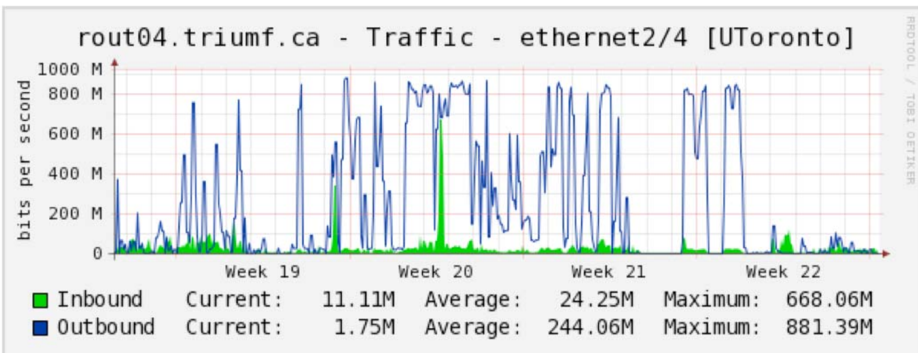

Figure 4: Typical traffic pattern for the University of Toronto Tier 2 in May of 2010

Our third target institution for upgrade is the University of Toronto SciNET facility (part of the SciNET HPC consortium). This upgrade here requires more work as the site is outside of the CANARIE Western ROADM network.

To facilitate these upgrades TRIUMF and BCNet are installing an additional 6 lit fibre pairs between TRIUMF and the CANARIE ROADM POP at the University of Victoria. This will allow the migration away from CWDM multiplexing equipment presently used to carry Tier 2 lightpaths to the local CANARIE ROADM POP. This is expected to reduce the cost both for 10G optics and the maintenance on the circuits. The maintenance and debugging of the CWDM circuits has proved excessive in terms of manpower from BCNet technical staff in the Vancouver area.