Canadian Networks for Particle Physics Research

2011 Report to the Standing Committee on Interregional Connectivity, ICFA Panel
January 2012

This report describes the status and plans of the Canadian network infrastructure used for particle physics research in 2010. HEPnet/Canada (http://hepnetcanada.ca) coordinates the networks[1] and works with the network providers and particle physics laboratories. We describe the status of the CANARIE network infrastructure, the ATLAS Tier 1 and 2 centres in Canada, and highlight other network projects in 2011.

CANARIE Network

CANARIE is Canada's Advanced Research and Innovation Network, a dedicated network of high-speed, fibre optic cable that stretches 19,000 km across Canada and links researchers and innovators throughout Canada and around the world. In addition to the network, CANARIE funds programs and tools that promote the evolution of a leading-edge digital infrastructure supporting Canadian researchers and innovators, fostering a growing, competitive, knowledge-based Canadian economy.

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 ROADM technology. Two ROADM networks, depicted in the thick blue lines in Figure 1, were built in 2005 and 2006. The first ROADM network, called the "Eastern Canadian ROADM", established a 2700 km, multi-degree optical mesh network along the busiest corridor of the CANARIE Network: Chicago –Toronto – Ottawa – Montreal – New-York City. The second ROADM network, 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. CANARIE partnered with Ciena during the Ciena Vector Summit in October, 2011 to demonstrate the system's capability of bringing up a 100G wavelength from Ottawa, Ontario to Chicago, Illinois as well as a 40G wavelength from Ottawa, Ontario to Montreal, Quebec. In November 2011, with the assistance from Ciena, CANARIE partnered with BCNET and the University of Victoria for a100G demonstration during SuperComputing 2011.

CANARIE continues its ROADM expansion plan with two new segments: the first segment is an extension of the Western ROADM from Calgary to Edmonton, Alberta. The project is underway and will be completed by April, 2012. The second project, which will commence in early 2012, is building a new ROADM segment from Winnipeg, Manitoba to Thunder Bay, Ontario.

The second layer of the optical network infrastructure is the SONET layer where CANARIE delivers most of network services. In this layer, the network is composed of individual point-to-point 10G SONET wavelengths 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. Delivering network services over the SONET layer has been a core service delivery method for a number of years. CANARIE engineering continues to evaluate other transport technologies to maximize the infrastructure investment.

CANARIE Network Services

CANARIE offers a number of network services to users: Lightpath; IPv4 Network service; IPv6 full Internet routing; and Content Delivery Service (CDS).

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

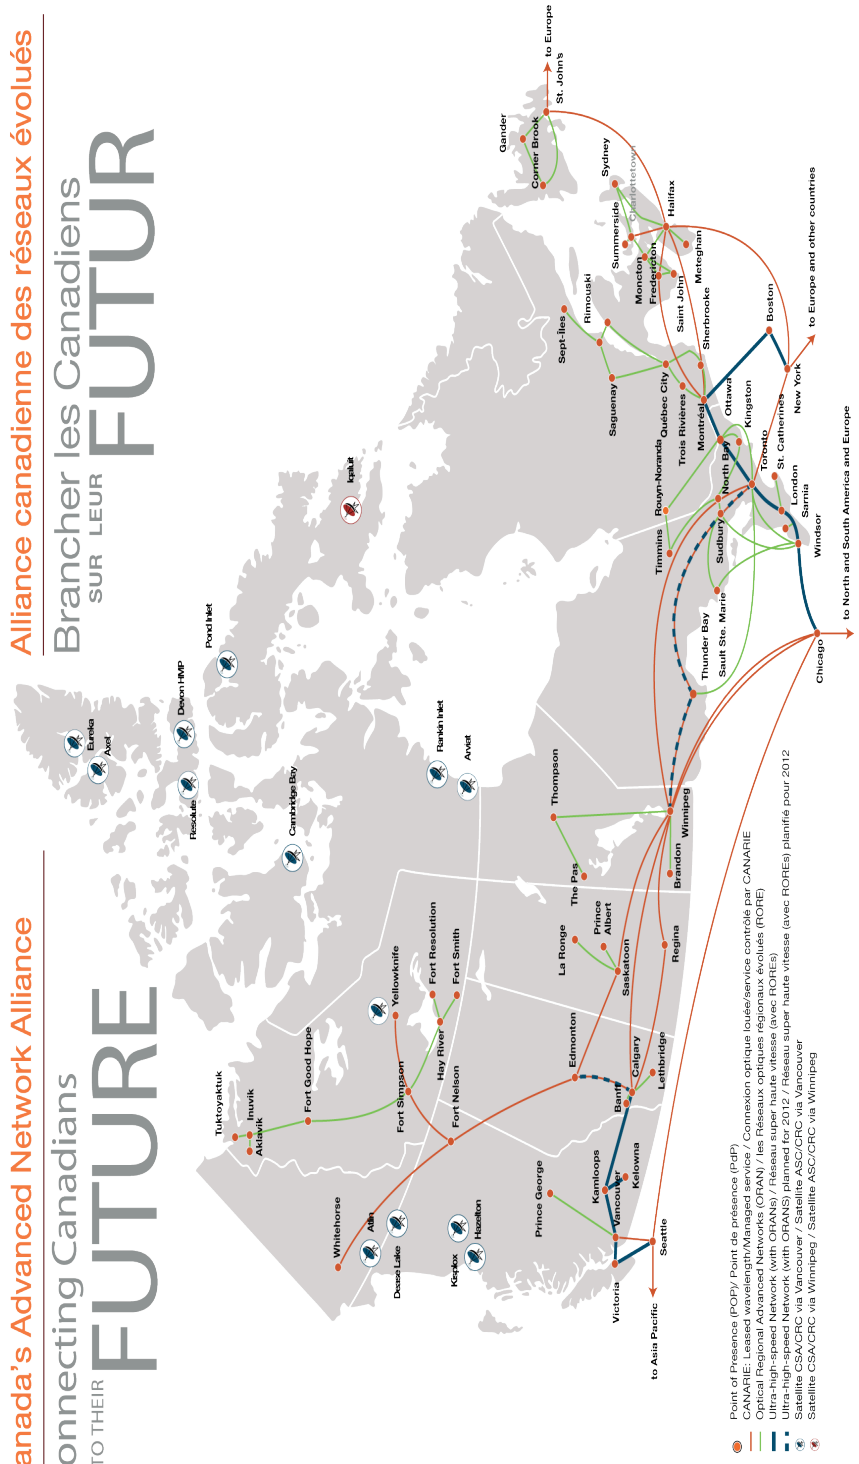

Figure 1: The CANARIE network

CANARIE Lightpath service
The CANARIE Lightpath service is delivered on 1 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 ROADM networks, lightpaths can also be 10GbE wavelengths that can be dropped directly into researchers' equipment, thus bypassing CANARIE optical switches. TRIUMF has had access to lightpaths since the early days of the R&E network, and has used the service to connect to the LHCOPN network.

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. In 2011, CANARIE made a few major changes to the IP core network: added a new node in Saskatoon as demand increased, acquired a logical routing service from BCNET to extend routing infrastructure into BC, and decommissioned the minor node in Edmonton to minimize operational expense. Currently, the IP network is comprised of six routing nodes, which are located in Calgary, Saskatoon, Winnipeg, Toronto, Montreal, and Halifax, as well as one logical router in Vancouver. Nine internal segments link the routers in a partial-mash topology and five external network segments that 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 10GbE. The international connectivity is important and critical for TRIUMF, as it exchanges LHC traffic with University of Melbourne over the IP network.

With the depletion of IPv4 addresses IPv6 full routing connectivity became critical to Canadian research institutions as some networks could only be accessed through IPv6. 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 initiated programs 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 to carrying full IPv6 Internet routing to Canadian institutions. In 2011, CANARIE funded the development of a series of IPv6 training modules to equip users for upcoming IPv6 deployment challenges and an equipment refresh funding program to replace outdated, non-IPv6 equipment.

Content Delivery Service
As Canadian researchers and institutions continue to grow their demand for cloud computing and content providers services, there is an increasing need for faster access to content providers (CPs) such as Google and Amazon. In 2010, funding was approved to develop and deploy a Content Delivery Service (CDS) by mid 2011. CDS provides reliable, high-speed access to CPs for all Research and Education institutions in Canada across the CANARIE Network. The service is delivered on a dedicated network: CANARIE's Peering Service IP network. That network, which is logically separate from the CANARIE IP Network providing the R&E service, has links to 3 different inter-exchange points: SIX (Seattle, WA), PWave (Seattle, WA), TorIX (Toronto, ON) and NYIIX (New York City, NY). Currently, 6 content providers including Amazon, Microsoft and Google are reachable through the CDS.

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). These connections provide CANARIE with the ability and capacity to create dedicated lightpaths throughout most of the world, enabling researchers to collaborate internationally. One of the prime examples is the LHCONE deployment linking Tier 1 and 2 sites over common exchange points,

most of which are the GLIF nodes in US and Europe. The connection to MANLAN allows CANARIE to deliver a dedicated network connection linking the University of Toronto Tier 2 site to the LHCONE facility in MANLAN.

ATLAS Tier 1 Computing Centre at the TRIUMF Laboratory

TRIUMF, Canada's National Laboratory for Nuclear and Particle Research, operates 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 circuit 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 2, 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 Geant network to CERN.

The 5G link has been 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.

In 2011 a 4th route to CERN was opened via the IceLink project that connects CANARIE directly to NorduNet via North Atlantic submarine cables traversing Newfoundland Iceland and then on to the Nordic countries.

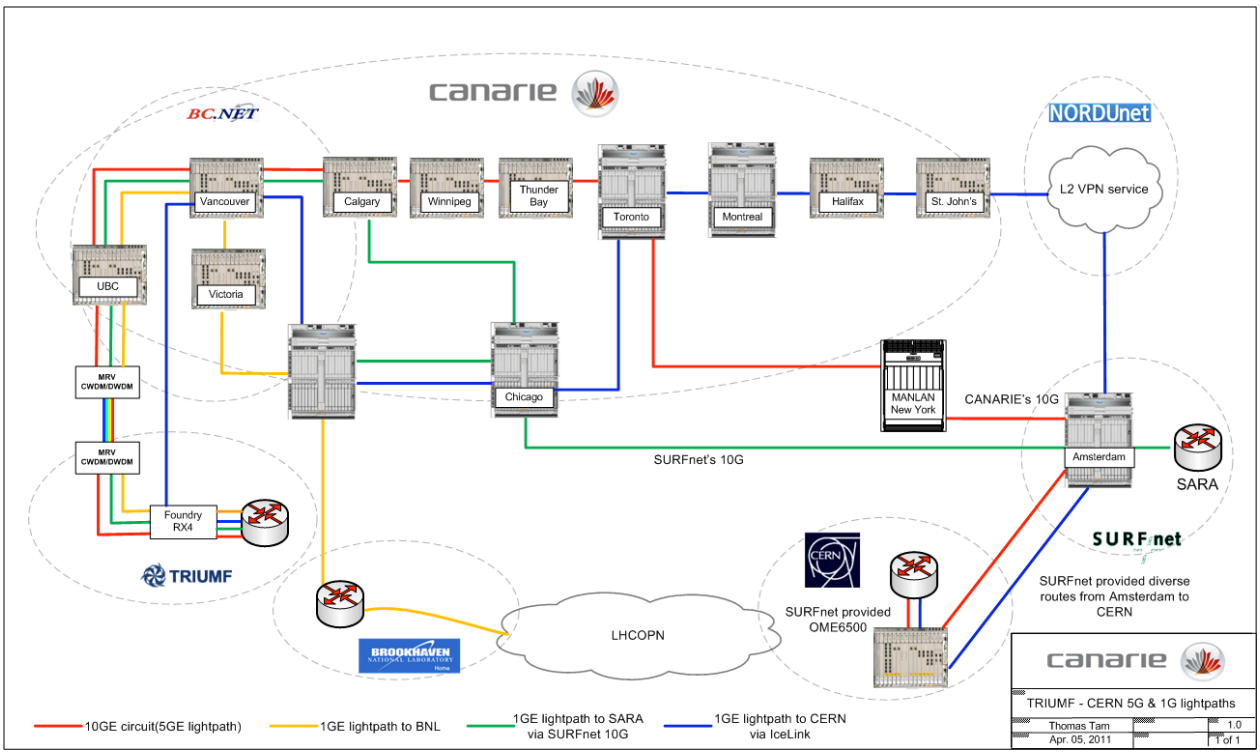

Figure 2: TRIUMF – CERN network

The TRIUMF Tier 1 hosts Tier 2 centres at the University of Victoria, University of Alberta, University of Toronto, Simon Fraser University and McGill University and the University of Melbourne (Australia). 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 10 G circuit with WestGrid. Each path, excepting Simon Fraser University, is carried on CANARIE ROADM network. The lightpaths have met the network demand but frequently the connections are saturated.

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. Improvements in the local fibre infrastructure in Vancouver have allowed the removal of CWDM equipment between TRIUMF and UBC ROADM POP. TRIUMF and BCNet have established an additional 6 lit fibre pairs between TRIUMF and the CANARIE ROADM POP at the UBC. Each T2 is now carried on it's own fibre pair. This simplification allows for lower cost upgrades to individual T2s, in addition decreased maintenance cost on the CWDM equipment.

In 2011 TRIUMF also undertook a 10G upgrade of their general purpose IP network connection to BCNet. This upgrade serves both the University of Melbourne T2 and the increasing traffic from T2s outside of Canada.

Developments in 2011

LHCONE Network

The LHCONE network was a significant focus of effort in 2011. The University of Toronto was the first North American institution online in the prototype global VLAN LHCONE design in September 2011. A dedicated 10 G circuit was established between the SciNet computing facility and the CANARIE POP in downtown Toronto. From the POP a 2G portion of CANARIE existing connection to ManLan in New York was dedicated to the LHCONE. The circuit saw maximum 1 day average of 1.88 Gbps and typical daily maximums of 600 – 800 Mbps. Typical usage of the circuit is shown in Figure 4. The circuit has already proved valuable for the University of Toronto. We plan to explore models for making LHCONE connectivity available to all Canadian T2s and TRIUMF. Possible models include the implementation of a series of VRF as an overlay to the existing CANARIE IP network or establishing dedicated connections for individual sites to LHCONE exchanges point in Seattle, Chicago and New York (the less desirable option).

Network monitoring

The increasing reliance of ATLAS on networking outside the standard tier architecture has created increase need to characterize end-to-end network performance between many ATLAS sites. In order to fulfill this demand and to prepare for the arrival of the LHCONE, HEPnet/Canada has deployed a dedicated network identical perfSonar PS monitoring boxes at all Canadian Tier 2s and at TRIUMF. All host were physically installed by the end of November and most configuration was completed by year-end. The perfSonar PS Toolkit was selected for its' ease of installation and excellent support from Internet 2. Each site has dedicated bandwidth measurement hosts with 10G Intel SFP+ direct attached network card to enable the use of either copper coax cabling or traditional optics. Each site also has a dedicated latency measurement box using a standard 1GE network card. We were able to make use of perfSonar monitoring dashboard developed at Brookhaven National Lab for quickly displaying network performance metrics.

100G Network Demo

For SuperComputing 2011 in Seattle the University of Victoria in partnership with Caltech, CANARIE and BCNet undertook a demonstration of 100G networking technology. With the industrial partners Ciena, Brocade and Dell a 100GE circuit was established over BCNet and CANARIE between the University of Victoria data centre and the Caltech exhibition booth of the show floor of the Washington State convention center. The demonstration showed a sustained full duplex throughput of 186 Gbps. In addition the team was able to achieve 60 Gbps disk-to-disk throughput using a two small clusters of machines with SSDs. The test successful demonstrated that it possible to complete saturate a 100 G connection using a modest set of properly configured hardware and software.