Canadian Networks for Particle Physics Research

2013 Report to the Standing Committee on Interregional Connectivity, ICFA Panel
January 2014

This report describes the status and plans of the Canadian network infrastructure used for particle physics research in 2013. 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 2013.

CANARIE Network

The CANARIE Network is Canada's national research and education network. CANARIE with its high performance network infrastructure supports data-intensive research, leading-edge research and big science across Canada and around the world. Operating with latest optical and routing equipment, CANARIE offers IP based network services and Lightpath services (dedicated end-to-end connections), supporting researchers, scientists and students at over 1,100 Canadian institutions, including universities, colleges, research institutes, hospitals, and government laboratories.

CANARIE, together with twelve provincial and territorial network partners, forms Canada's National Research and Education Network (NREN) depicted in Figure 1. This powerful digital infrastructure connects Canada's researchers and innovators provincially, nationally, and globally to the data, tools, colleagues, and classrooms.

CANARIE also supports the development of research software tools that enable researchers to more quickly and easily access research data, tools, and peers. In support of Canada's high-tech entrepreneurs, CANARIE offers cloud-computing services to help them accelerate product development and gain a competitive edge in the marketplace.

CANARIE Optical Infrastructure

The CANARIE optical infrastructure is comprised of two optical technologies: SONET and DWDM systems. In 2006 CANARIE acquired fibre and started building a ROADM based DWDM infrastructure, with support of 88 wavelengths up to 100Gbps per wavelength. The western Canada build covered the corridor of Seattle-Victoria-Vancouver-Calgary-Edmonton. The eastern build spanned from Chicago through Ontario to Montreal and down to New York. At the end of 2012, a standalone Winnipeg – Thunder Bay segment was added.

With CANARIE's long-term vision of building a coast-to-coast fibre network, in mid-2013 CANARIE acquired more fibre for two sections of the network: Central and Atlantic Canada. The nation-wide fibre network allows CANARIE to deliver high-bandwidth services for data intensive applications, high performance computing centres and research facilities.

SONET is an important layer of CANARIE network; it delivers most of CANARIE network services to users. The SONET infrastructure is composed of individual point-to-point OC-192 SONET wavelengths with a topology consisting of a series of rings. The ring architecture allows CANARIE to offer redundant network services, as in the case of the CANARIE IP network. Delivering network services over the SONET infrastructure has been a core service delivery method for a number of years, however CANARIE engineering continuously evaluates transport technologies to better serve institutional users and maximize infrastructure investment.

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

CANARIE Network Services

CANARIE offers a number of network services such as: Lightpath service, IP Network service, Content Delivery Service.

CANARIE Lightpath service
CANARIE Lightpath service is standardized in Ethernet technology. With service capacity ranging from 155 Mbps to up to a full 100Gbps circuit, a lightpath can be dropped directly into researchers' equipment; connecting high-performance computing centres across Canada or in other continents. TRIUMF has been using lightpaths since the early days to create LHCOPN. In 2012 CANARIE partnered with other NRENs to deploy LHCONE, a VRF based technology infrastructure over the IP network, connecting TRIUMF and the Canadian Tier-2 centres to the world.

CANARIE IP Network service
CANARIE IP network service remains the biggest workhorse of CANARIE's service offerings. The core network, consisting of 7 core routing nodes, provides full and equal support for IPv4 and IPv6 unicast and multicast routing. There are eleven internal segments that link the routers in a partial-mash topology and five external network segments that extend to international R&E exchanges: Pacific Wave in Seattle, StarLight in Chicago, and Manhattan Landing (MANLAN) in New York. With an anticipation of substantial traffic growth in the next few year, in 2014 CANARIE plans a major overhaul of IP network, deploying a full 100Gbps IP infrastructure across Canada from Vancouver to Montreal and down to New York.

Content Delivery Service
CANARIE's Content Delivery Service provides institutional users high-speed access to major content providers, like Amazon, Microsoft, Google, Yahoo, Facebook and Box.net. It has become an important service of CANARIE service offerings. The Content Delivery Service IP network, which is logically separated from the CANARIE R&E IP Network, links to SIX (Seattle, WA), Pacific Wave (Seattle, WA), TorIX (Toronto, ON) and NYIIX (New York City, NY) to source content.

International connectivity
The CANARIE Network is connected to GLIF (Global Lambda Integrated Facility) through three R&E International eXchange Points (IXPs): Pacific Wave, StarLight, and MANLAN (Manhattan Landing Exchange Point). These connections provide CANARIE with the ability and capacity to establish dedicated lightpaths to link researchers around the world, enabling international collaborative projects seamlessly. One of the prime examples is the LHCONE deployment linking Tier 1 and 2 sites over common exchange points to other Tier 1/2 nodes in US and Europe.

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.

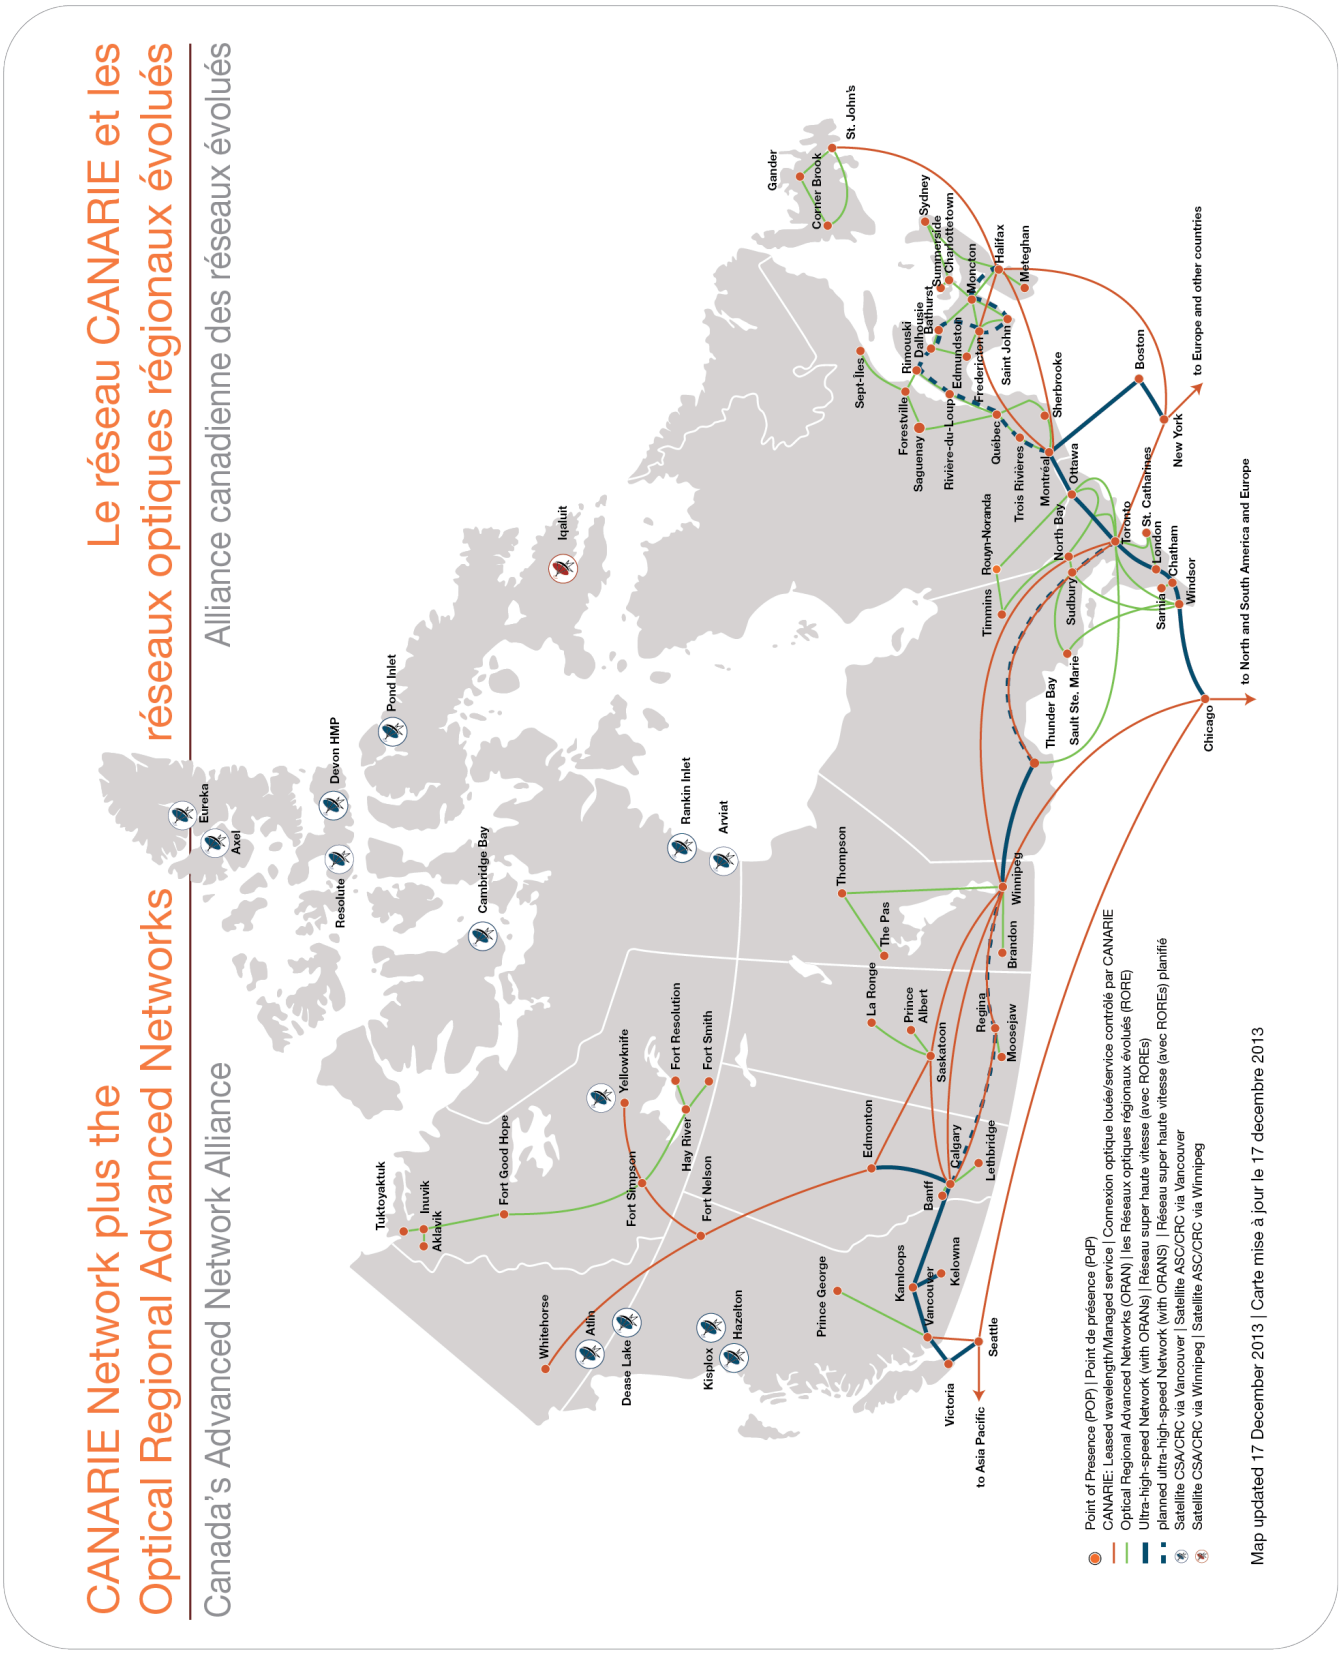

Figure 1: The CANARIE network

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, the lightpath capacity can be increased in 155Mbps increments up to the full 10G. The T0-T1 connection has been adequate and static for a number of years.

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 to Europe, and is expected to be able to temporarily handle loads while the 5G link is restored.

The TRIUMF Tier 1 supports the Tier 2 centres at the University of Victoria, University of Toronto, Simon Fraser University (SFU) and McGill University and the University of Melbourne (Australia). From 2007 to 2012 the Canadian T2 centres were connected directly to TRIUMF using point-to-point 1G lightpaths. This was model was extremely successful during this period, however, clear signs of network saturation had developed in 2011 and the 1G lightpaths were clearly inadequate in 2012. In order to leverage the infrastructure developed for the LHCONE, the Canadian T2-TRIUMF networking was completely redesigned in 2012 (see next section).

LHCONE Network

The LHCONE network is deployed as purpose built VRF overlay network for use by Canadian T1 and T2 sites spanning all CANARIE Juniper MX series routers. All Canadian T2s and the T1 were connected to this service at 10G by December 2012. The LHCONE serves HEP connectivity requirements within Canada in addition to providing connectivity to international sites. This approach allowed us to eliminate a number of point-to-point circuits between TRIUMF and the T2s. The CANARIE LHCONE VRF has International peerings have with Internet2, and GEANT at ManLan (New York City) and Starlight (Chicago), and with ESnet at PacificWave (Seattle).

During the development of the LHCONE in Canada the decision was taken early to avoid any type of Policy Based Routing (PBR) on the advice of the Canadian R&E networking community. We concluded that the best option for sites without dedicated hardware for the LHCONE would be to create a VRF at the site in a configuration shown in Figure 3. The Site Local LHCONE VRF is created on the available data centre edge equipment and configured to peer with the CANARIE LHCONE VRF. The Campus Router is set as the default route for the site Local LHCONE VRF to obtain regular R&E IP network connectivity. The CANARIE VRF is configured to accept only pre-negotiated subnets that are known to contain LHC equipment. The Site Local LHCONE VRF does not re-advertise the remote LHCONE routes to the Campus Router to avoid any asymmetric routes.

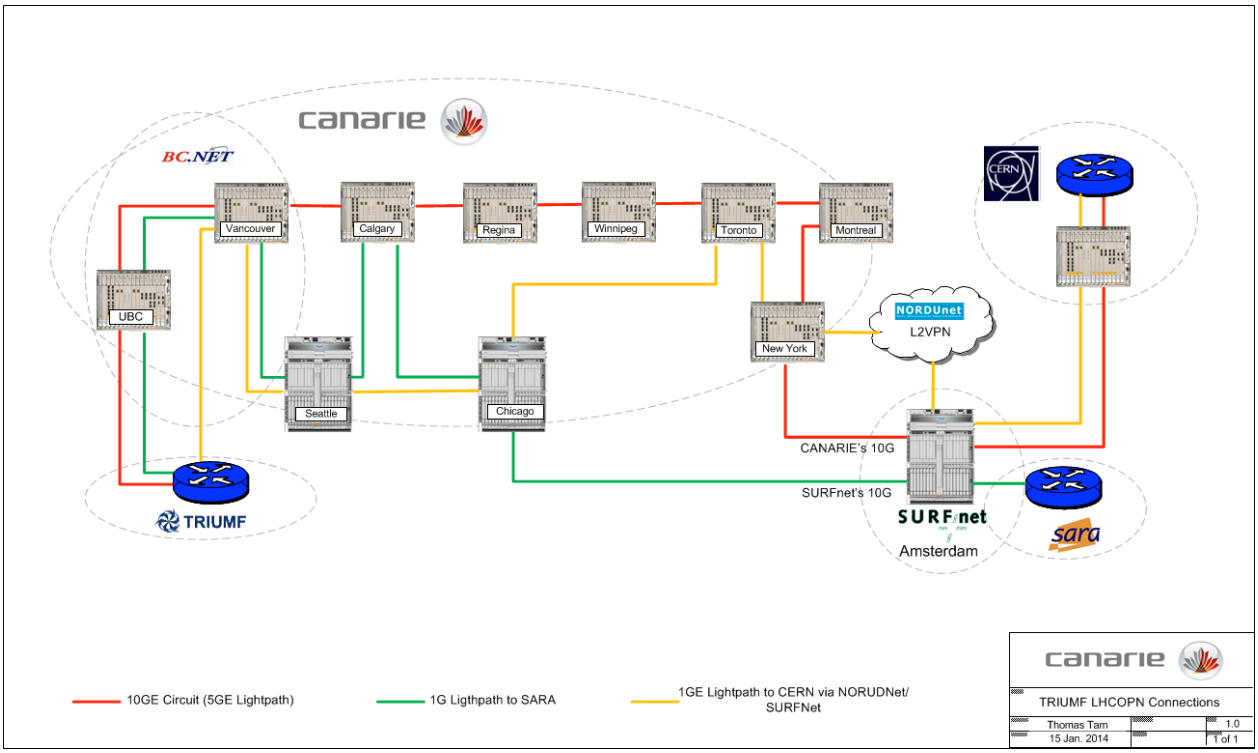

Figure 2: TRIUMF – CERN network

Developments in 2013

LHCONE

The LHCONE L3VPN was very stable and successful during 2013 accounting for a significant fraction of the total traffic across CANARIE. A contributing factor to the stability was the initial design choices to avoid PBR and use VRFs at each LHCONE site were dedicated equipment was not available (see Figure 3). Due to high traffic on the TRIUMF LHCONE (see Figure 4) a second 10G link will be put in place in Jan 2014. Because of contention between LHCONE traffic and other IP network traffic CANARIE added a second 10G path across the country. We also expect that some T2s in Canada may require the addition of a second 10G circuit in 2014 or prior to the restart of the LHCONE. For example, Figure 5 shows a T2 peaking over 9 Gbps.

Software Defined Networking

Knowledge of the Software Defined networking (SDN) within the Canadian HEP community was very limited at the start of 2013. In an effort to familiarize the community with SDN we have undertaken a project in partnership with multiple Canadian R&E networks to deploy a network of OpenFlow capable switches spanning Victoria to Ottawa with dedicated breakable 1G network in addition to 10G segment from Victoria to Internet2 Advanced Layer 2 Services (AL2S) switch in Seattle. Switches are located at Victoria, Vancouver, Edmonton and Calgary. The network became operational in late October 2012.

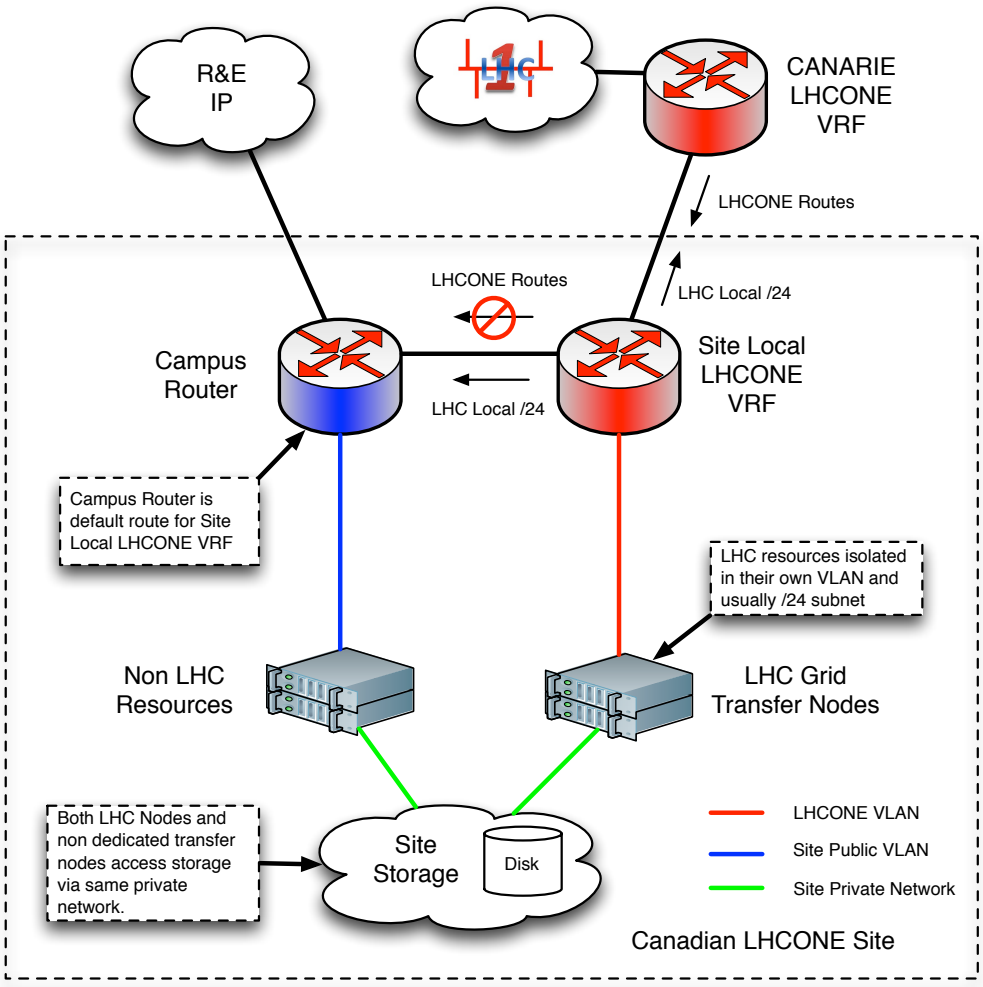

Figure 3: LHCONE network design at a typical Canadian HEP site.

Network monitoring

HEPnet Canada deployed 10G capable perfSonar boxes to all sites during 2011. In 2013, the monitoring boxes were refreshed with the latest CentOS 6 based version of the perfSonar toolkit. A complete solution for displaying and alerting on the results of all perfSonar boxes remains elusive. Significant development effort was invested in developing dashboard style tools by other laboratories, however no solution has yet to arise which covers the full set of requirements. For further details of the hardware configuration please refer to the 2011 report.

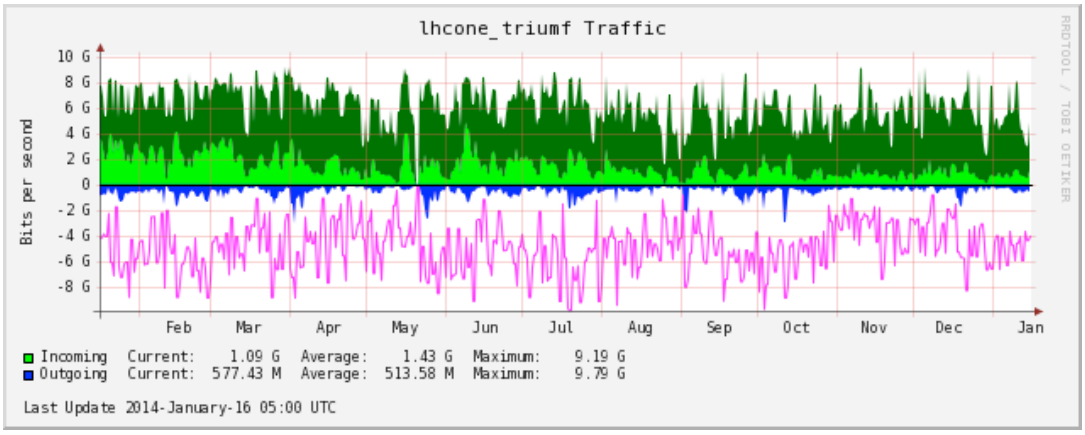

Figure 4: Heavy utilization throughout 2013 on the TRIUMF LHCONE circuit that is used both for Canadian T2 and other international LHCONE sites.

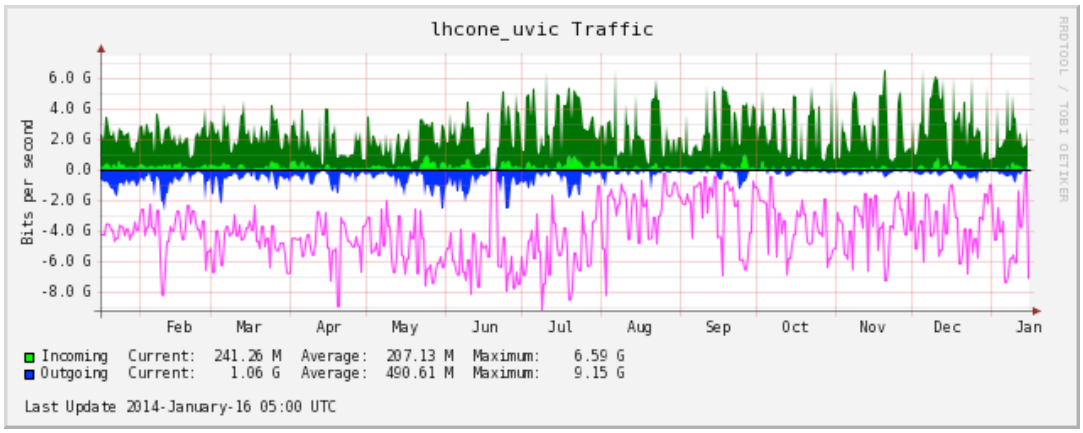

Figure 5: LHCONE traffic from the University of Victoria.

Intercontinental OpenFlow demo at SC13

At Super Computing 2013 in Denver we participated in a demonstration of an Intercontinental OpenFlow tested spanning Victoria – Denver – San Paulo – Chicago – New York – Amsterdam – CERN (in Geneva). The Western most node of the Canadian SDN testbed at UVic was connected with a 10G link leveraging the Internet2 AL2S switch in Seattle. Two diverse paths were carried over Internet 2 terminating on two separate switches on the Super Computing show floor. Physics data was transferred over this network using OpenFlow link layer multipath switching to efficiently utilize the multiple network paths to form a loop free topology. The multipath utilization is completely transparent to the end hosts of the network and accomplished using the Caltech developed OLiMPS OpenFlow controller that resided at CERN in Switzerland. This methodology allows an automated in-network load balancing of various flows to increase network utilization and efficiency. Moreover, multiple network links can be arbitrarily added to a network to increase capacity and redundancy without the traditional challenges of Layer 2 link aggregation using LACP.[2]

[2] For further information visit http://supercomputing.uvic.ca and http://supercomputing.caltech.edu