Canadian Networks for Particle Physics Research

2014 Report to the Standing Committee on Interregional Connectivity, ICFA Panel
January 2015

This report describes the status and plans of the Canadian network infrastructure used for particle physics research in 2014. HEPnet/Canada[1] (http://hepnetcanada.ca) coordinates the network, working together with the network providers, and the Canadian universities and laboratories. We describe the status of the CANARIE network infrastructure, the ATLAS Tier-1 and Tier-2 centres in Canada, and our other network projects.

CANARIE Network

CANARIE is responsible for Canada's national research and education network. CANARIE supports data-intensive research, leading edge and big science in Canada and around the world. Operating with latest optical and routing equipment, CANARIE offers IP based network services and point-to-point (p2p) connection service (lightpath service), 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, connecting Canada's researchers and innovators provincially, nationally, and globally to the data, tools, colleagues, and classrooms.

CANARIE Optical Infrastructure

CANARIE's vision is to complete a coast-to-coast fibre network, The nation-wide fibre network allows CANARIE to deliver high-bandwidth services for data intensive applications, high performance computing centres and research facilities.

In 2006, CANARIE started acquiring fibre, building up ROADM based DWDM infrastructure capable of supporting 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. Later, a standalone Winnipeg – Thunder Bay section was added.

In 2014, CANARIE acquired fibre for two sections of the network: Central and Atlantic Canada. The central section covered Calgary-Regina-Winnipeg; the deployment was started with an anticipation of completion in January 2015. The Atlantic section extended the eastern fibre network from Montreal into Halifax. The Atlantic fibre build is underway and expected to be completed in June 2015.

In addition, CANARIE operates a SONET infrastructure, delivering most of network services to users, especially for p2p connection services. CANARIE is planning to phase out the SONET infrastructure because of technological limitation prohibits the expansion of the SONET infrastructure. The migration of SONET service into MPLS VPN service will start in early 2015.

CANARIE Network Services

CANARIE offers a number of network services such as: R&E IP Service, Content Delivery Service, p2p Connection Service.

R&E IP Service

IP network service remains the biggest workhorse of CANARIE's service offerings. The core network provides full and equal support for IPv4 and IPv6 unicast and multicast routing, with a number of internal

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

segments that link the routers in a partial-mash topology and 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 anticipated traffic growth in the coming years, CANARIE has started a major overhaul of IP network, augmenting the infrastructure with a number of 100GE links and creating a redundant 100Gbps IP infrastructure across Canada and to three international R&E exchanges.

Content Delivery Service
Content Delivery Service provides institutional users with 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), Pacific Wave (Seattle), TorIX (Toronto) and NYIIX (New York City) to source content.

p2p Connection Service
CANARIE p2p Connection Service is based on standardized Ethernet technology. Servicing up to a full 100Gbps connection, a p2p connection can be dropped directly into researchers' equipment, connecting high-performance computing centres or research facilities across Canada or in other continents. TRIUMF has been using p2p connection services since the early days to establish LHCOPN. Currently LHCOPN is being delivered over SONET based p2p connection service and LHCONE is provisioned based on MPLS type of p2p service.

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.

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 TRIUMF T1 to CERN T0 circuit, depicted in Figure 2, runs over the CANARIE SONET infrastructure until it disembarks North America in New York City. This circuit will be replaced in 2015 with a MPLS VPN circuit that will connect with the ANA-200 transatlantic link at ManLan in New York City. The T1 to T0 circuit was previously provisioned with 5 G of available we be allocated 10 Gbps with the potential for this to be increased easily within the CANARIE network in their more flexible MPLS system and across the ANA-200 Network. We expect that 10G of T1 to T0 bandwidth will be sufficient for the needs of the LHCOPN TRIUMF connection in 2015 and 2016, as we expect much of the traffic growth to occur in the LHCONE. The redundancy previously provided by a standby 1G path will be provided within the new 5x100G network which CANARIE is commissioning as of the spring of 2015. TRIUMF is currently considering upgrade path to 100G as CANARIE will be capable of providing 100G coast to coast in early 2015.

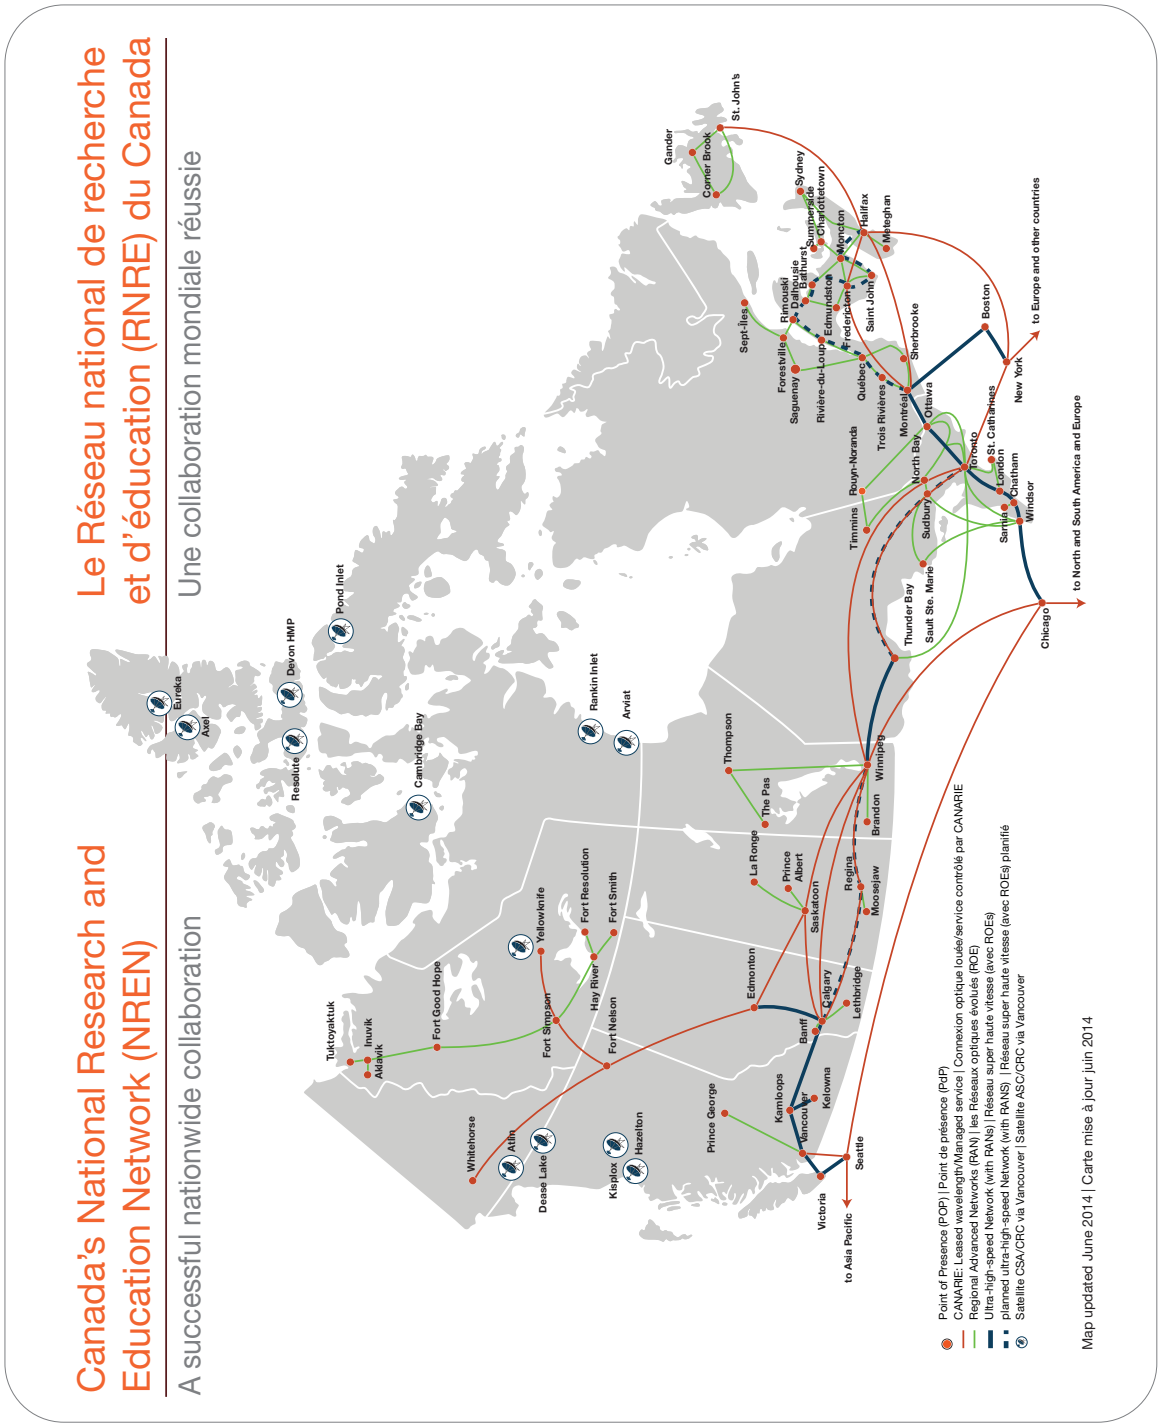

Figure 1: The CANARIE network

The TRIUMF T1 supports the T2 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, ESnet and GEANT at ManLan (New York City) and Starlight (Chicago), and with ESnet, Internet 2 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.

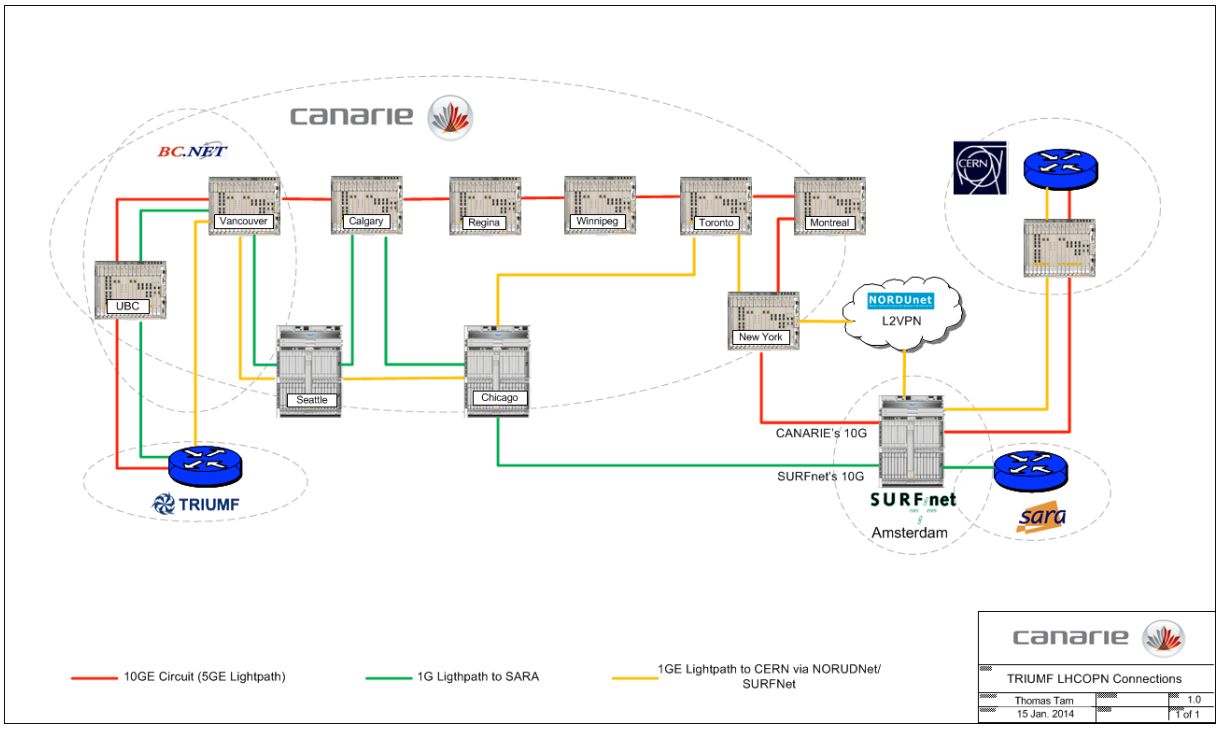

Figure 2: TRIUMF – CERN network for 2014.
This network will be replaced by a protected MPLS VPN over CANARIEs 100G network in 2015.

Developments in 2014

LHCONE
The LHCONE L3VPN was very stable and successful during 2014 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). LHCONE bandwidth across CANARIE was provision at 2x10G circuits in 2014 and the LHCONE overlay network will be moved onto the new CANARIE 100G IP network as the first customer in March of 2015. Canadian T2s are beginning to observe occasional bottlenecks with 10G links into the LHCONE, particularly in the case where the LHCONE uplink is shared with other projects (UVic and Simon Fraser). Investigation for adding a second 10G LHCONE link is being considered at most T2s. We also expect at least one Canadian T2 will move to 100G in 2015. TRIUMF will add a 3x10G capacity for the LHCONE in early 2015, while considering 100G upgrade possibilities.

BelleII on LHCONE
The BelleII experiment located at the KEK Laboratory in Tsukuba in Japan is in the early stages of developing it's computing infrastructure. There are two BelleII sites in Canada at UVic and McGill. At the fall meeting of the LHCONE/LHCOPN working group there was agreement in principal to allow BelleII to use the LHCONE. In December of 2015 LHCONE BGP peering was established between Pacific Northwest National Lab (the largest BelleII computing site), KEK Lab and UVic at PacificWave in Seattle. Traffic from this experiment is currently low in Canada but expected to grow.

100G Single Server
During the 2014 Supercomputing Conference in New Orleans UVic Caltech and U Michigan demonstrated that a single server equipped with PCIe Solid State storage devices is able to drive LAN connection using 3x40GE NICs at over 100 Gbps and work on the WAN at more then 70 Gbps. We expect that these configurations will be useful both a fast data caches at the edge of larger sites and as method to enable smaller sites to provide high bandwidth storage simply.

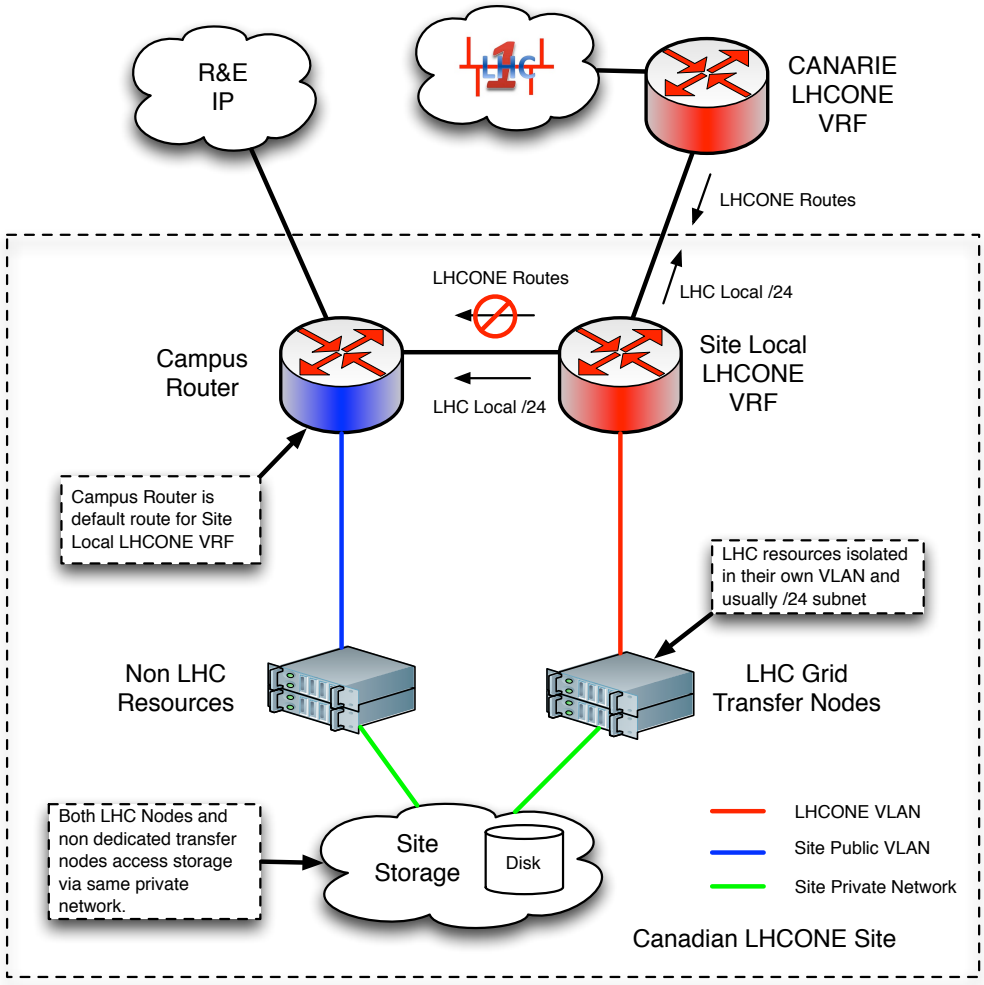

Figure 3: LHCONE network design at a typical Canadian HEP site.

Network monitoring

HEPnet Canada deployed 10G capable perfSonar boxes to all sites during 2011. In 2014, perfSonar became a mandatory installation for all ATLAS sites. In addition Canadian HPC centres and research networks adopted the tool en-mass. HEPnet nodes are now controlled using a central mesh configuration utility that allows the configuration for the tests to be pulled down from a central location. In addition we use the ESnet developed MadDash monitoring utility to graphically display a quick status of the Canadian testing mesh.

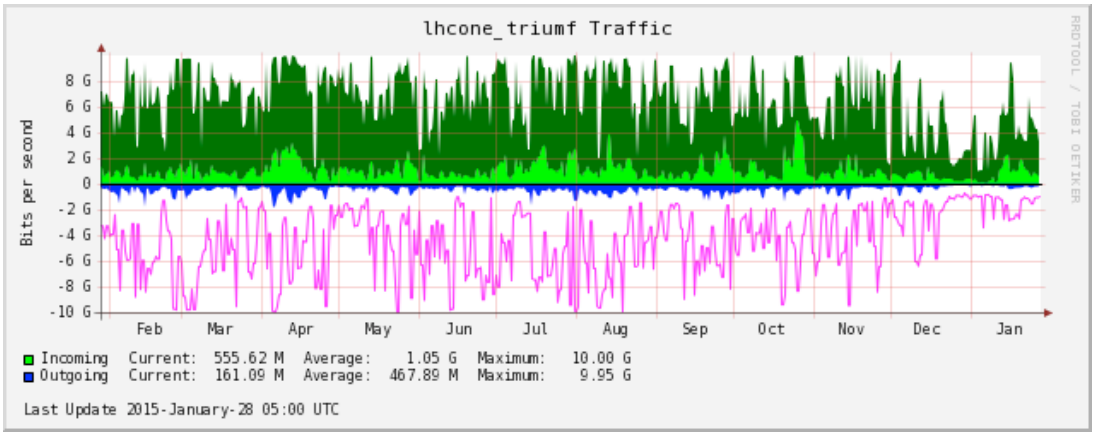

Figure 4: Heavy utilization throughout 2014 on the TRIUMF LHCONE circuit that is used both for Canadian T2 and other international LHCONE sites.

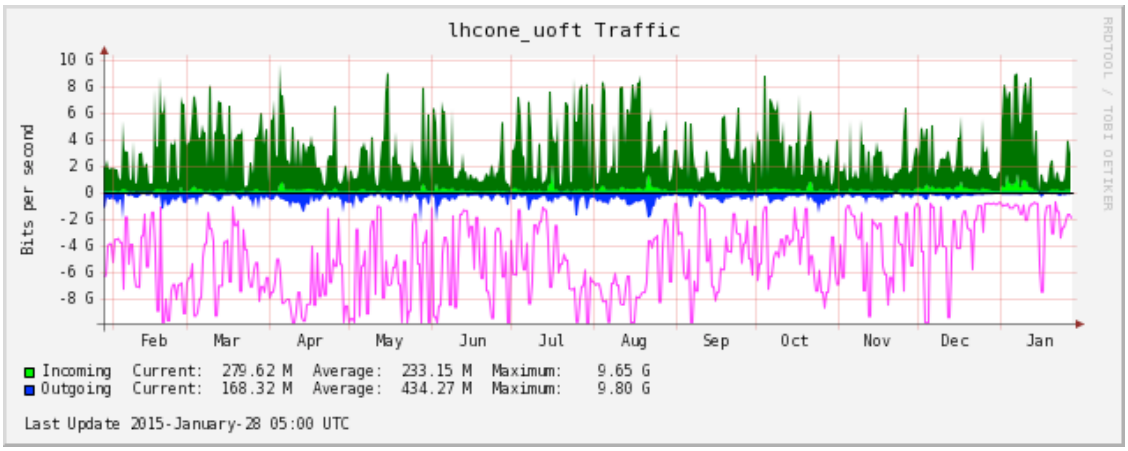

Figure 5: LHCONE traffic from the University of Toronto.