Canadian Networks for Particle Physics Research

2016 Report to the Standing Committee on Interregional Connectivity, ICFA Panel
January 2017

This report describes the status and plans of the Canadian network infrastructure used for particle physics research in 2016. 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

Twelve provincial and territorial network partners, together with CANARIE, collectively form Canada's National Research and Education Network (NREN). Canadians at universities, colleges, research institutes, hospitals, and government laboratories rely on this ultra high-speed network to collaborate in data-intensive, cutting-edge research and innovation within Canada and with colleagues in over 100 countries.

Beyond the network, CANARIE funds and promotes reusable research software tools to accelerate scientific discovery. CANARIE also supports Research Data Canada as it leads national research data management initiatives, and through the Canadian Access Federation, provides identity management services that enable secure, ubiquitous connectivity and content access to the academic community. To boost commercialization in Canada's technology sector, CANARIE offers cloud resources to startups through its DAIR service, and links a powerful community of public and private sector partners in the Centre of Excellence in Next Generation Networks (CENGN).

CANARIE Infrastructure

CANARIE continued its vision of building fibre infrastructure from coast-to-coast. The nation-wide fibre network allows CANARIE to deliver high-bandwidth services for data intensive applications, high performance computing centres and research facilities. CANARIE started the fibre system extension into eastern Canada in 2016 and will complete the fibre build in 2017. The new system brings significantly increased network capacity in eastern Canada, providing a potential linkage to connect researchers in Europe through a cable landing point in Halifax, Canada.

In 2016 CANARIE completed the decommissioning of its SONET infrastructure and fully deployed 100Gbps redundant core IP network. Also CANARIE operationalized L2 VPLS services offering highly resilient, dedicated point-to-point or point-to-multipoint services for data intensive applications. The redundant nature of the L2 VPLS technology increases the availability of the network services.

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 100Gbps redundant infrastructure significantly improves the performance and reliability of our network services, which means fewer disruptions and delays for researchers. With a number of internal segments that link

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

the routers in a partial-mash topology, the IP core network provides full and equal support for IPv4 and IPv6 unicast and multicast routing. As well, the network consists of external network segments that extend to international R&E exchanges: Pacific Wave in Seattle, StarLight in Chicago, and Manhattan Landing (MANLAN) in New York. Through these exchange points CANARIE reaches National R&E networks around the world.

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, WA), Pacific Wave (Seattle, WA), TorIX (Toronto, ON) and NYIIX (New York City, NY) to source content. With increasing demand of accessing content, CANARIE started establishing direct peering with a few big content providers.

p2p Connection Service
As stated previously, CANARIE p2p Connection Service is standardized in L2 VPLS technology. Servicing up to a full 100Gbps connection, a p2p connection can be delivered directly into researchers' equipment, connecting high-performance computing centres or research facilities across Canada or in other continents. TRIUMF and Canadian Tier 2 centres utilize this service to support LHCOPN and LHCONE infrastructure.

International Connectivity
CANARIE, partnering with Internet2, NORDUnet, SURFnet, and GEANT provisioned 3 x 100Gbps links, called ANA-300G, between five open exchange points at both sides of the North Atlantic. The ANA-300G enables researchers and scientists to transfer data between North America and Europe at speeds that were previously only possible within the continents. In 2016, ANA operation team started a plan of re-terminating one of the 100Gbps links from New York to the CANARIE PoP in Montréal. The move would offer full diversity of physical landing points in North American. The new Montréal link allows direct traffic from Canada to Europe, providing even greater flexibility and network diversity to access research data and content.

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 Optical Private 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. This circuit was replaced in 2015 with a MPLS VPN circuit that will connect with the ANA-300 transatlantic link at ManLan in New York City. The T1 to T0 circuit is provisioned with 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-300 Network.

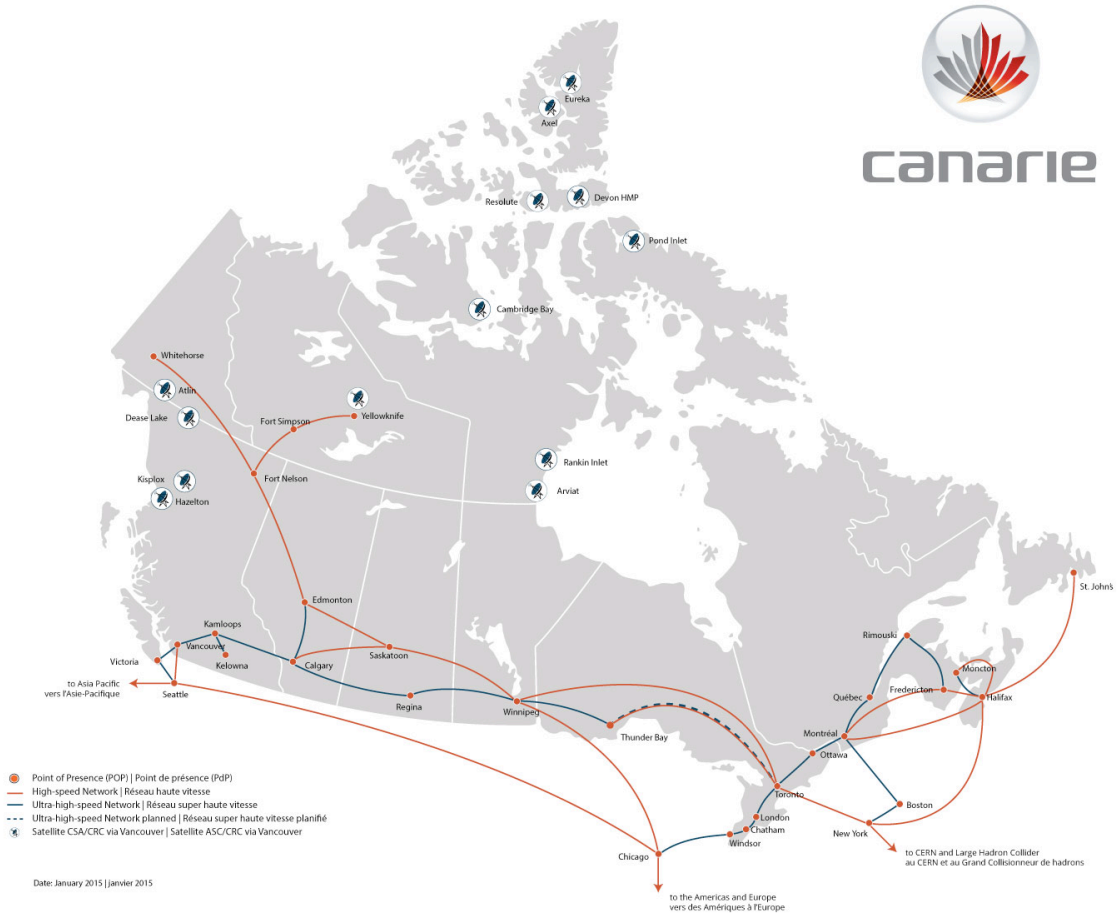

Figure 1: The CANARIE network in 2016

TRIUMF has one physical 10G LHCOPN connection to BCNET/CANARIE with two dedicated VLANs going from TRIUMF to CERN. The VLANs are designed to provide redundancy for Tier-1 site, and CANARIE achieves this by distributing them on a different set of geographically separated backbone routes. In case both VLAN are down LHCONE will be used as a last network resource to get to CERN.

LHCONE Network in Canada

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 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 2. 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.

The LHCONE L3VPN was very stable and successful accounting for a significant fraction of the total traffic across CANARIE. A contributing factor to the stability was the initial design that avoids PBR and instead uses VRFs at each LHCONE site were dedicated equipment was not available (see Figure 3). LHCONE bandwidth across CANARIE was provisioned at 2x10G circuits in 2014 and the LHCONE overlay network was moved onto the new CANARIE 100G IP network in 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 non-HEP projects. Investigation for adding additional capacity is underway.

The BelleII experiment located at the KEK Laboratory in Tsukuba in Japan is developing its computing infrastructure. The computing requirements are expected to be comparable to a smaller LHC experiment. Currently, there are two BelleII sites in Canada, at UVic and McGill. The BelleII experiment joined the LHCONE in 2015. At the moment traffic on the LHCONE from BelleII is modest in Canada but expected to grow over the next few years.

The Canadian LHCONE network was very busy in 2016. The top plot in figure 4 shows the network traffic from the Tier-1 facility at TRIUMF. The lower three plots show the traffic to and from Canada from ESNET, Internet2 and GEANT. Canada connects to GEANT from Montreal and Toronto; the plot only shows the traffic via Montreal (the traffic via Toronto is comparable).

Other activities

IPv6

With the advent of personalized cell phones and tablets, the IPv4 address range is becoming more and more limited. IPv6 with its vast address space provides a solution, but the whole affected network infrastructure outside and inside the universities and laboratories must be IPv6 enabled before worker nodes and workstations can enable IPv6 or enable IPv4 and IPv6 dual stack.

During early 2016 UVic enabled IPv6 on their perfSonar nodes, fully enabling in and outgoing perfSonar related traffic by Summer 2016, and offering help and advice for IPv6 configuration to other sites since. In September 2016 the WLCG Management Board approved the IPv6 deployment plan. Dual stack availability is mandatory for the Tier0 and Tier1s by April 2017. By April 2018 dual stack should be available in production for the Tier0 and Tier1s. By the end of Run2 a large number of sites should have migrated their storage to IPv6. At the end of 2016 TRIUMF successfully completed setup of dual stack PerfsonarPS, advertising IPV6 address space to LHCOPN and LHCONE. TRIUMF is working on their dual stack storage configuration in the timeline provided by WLCG.

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 a similar system for monitoring their network. The HEPnet nodes are participating in a central mesh including all ATLAS sites worldwide. A sub mesh including only Canadian sites will be configured, however it was not found necessary in the past. These nodes have been crucial in identifying networking problems in a timely manner on Canadian as well as on American networks.

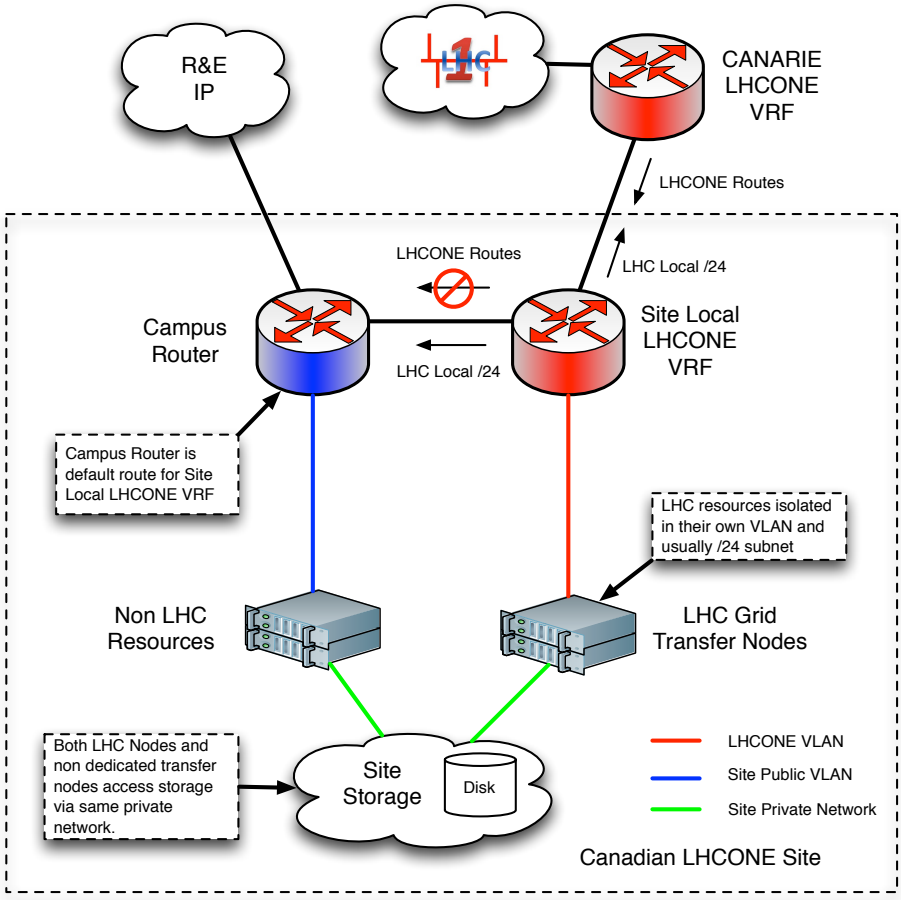

Figure 2: LHCONE network design at a typical Canadian HEP site.

Figure 4: LHCONE traffic from TRIUMF and Canadian traffic to ESNET, Internet2 and GEANT (Montreal).