

# Update on Canadian T1/T2 and R&E Network

R.Seuster(UVic), for  
Leslie Groer (UofT), Darcy Hodgson (Sharcnet),  
Sergey Chelsky (SFU), Simon Nderitu (McGill),  
Vitaliy Kondratenko(TRIUMF),Simon Liu(TRIUMF),  
Di Qing(TRIUMF),Ryan Taylor (UVic),Lixin Liu (SFU),  
Thomas Tam(Canarie), Richard Klinger(Canarie)

2019 Spring HEPiX meeting  
SDSC University of California in San Diego  
25-29 March 2019

# T1 at SFU/TRIUMF

- all T1 services previously at TRIUMF moved to SFU (~30km or 1ms RTT away)
- new hardware at T1 @ SFU fully commissioned and in production since 2018
  - 7680 cores ~166.6k HEPSPEC06 running SL7
- some funding from CFI(Canada Foundation for Innovation) left, will buy more storage + CPU
- old hardware at T1 @ TRIUMF out of warranty, but still running production (only simulation)
  - 4744 cores ~ 72.3k HEPSPEC06 running SL6
- storage: 10.6PB disk (1.3PB soon out of warranty)  
30PB tape (20 Ito8 / 12 Ito7 drives)
- site operation still responsibility by TRIUMF Tier1 personnel, good coordination with SFU data center

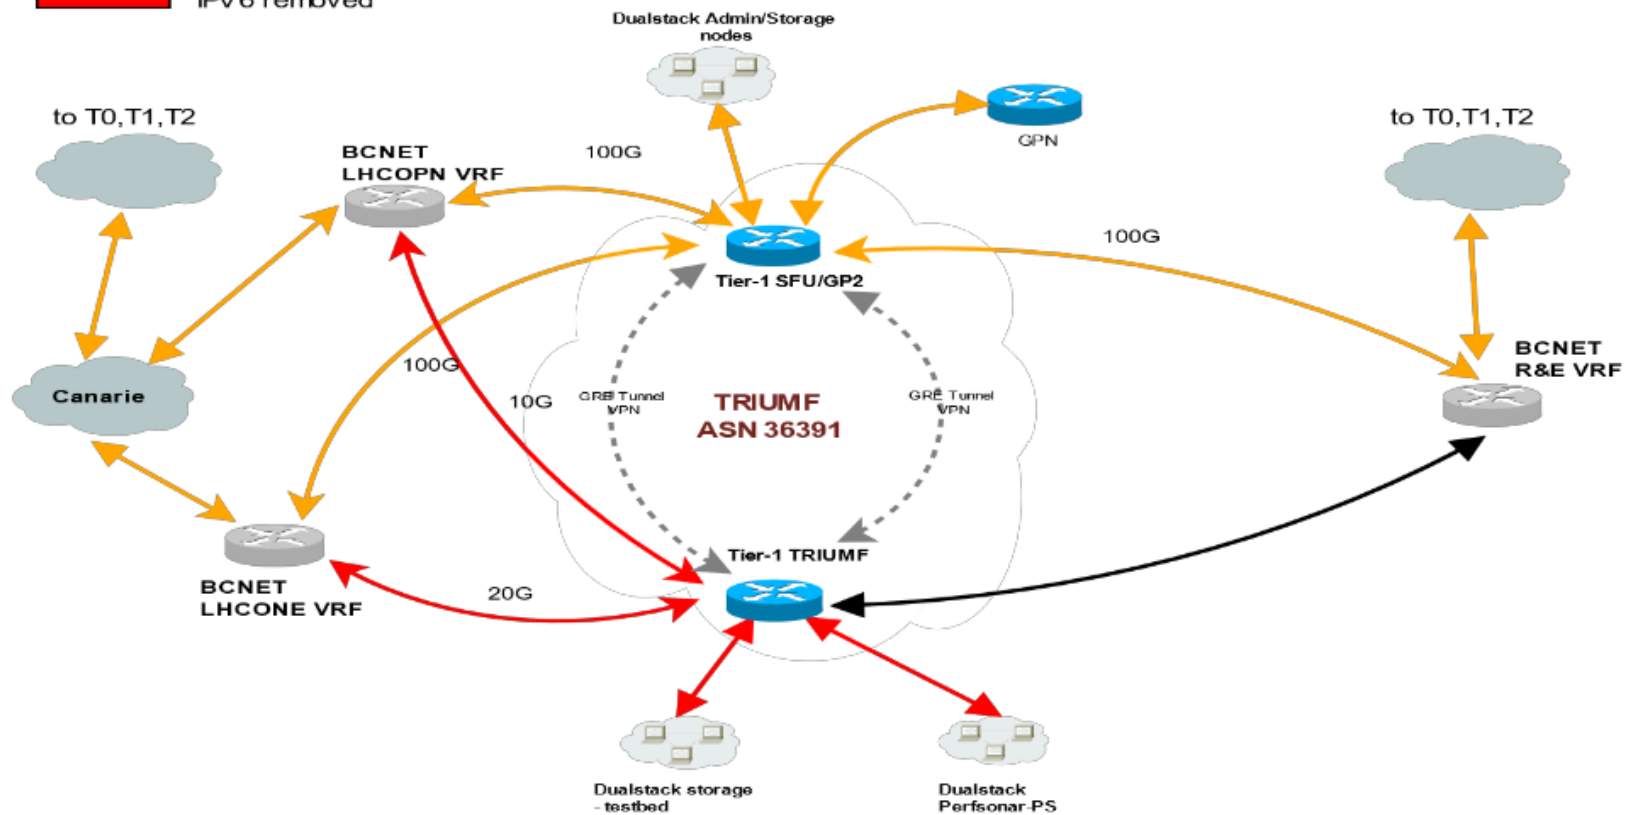
# Network and IPv6 @ SFU/TRIUMF

**IPV6 was enabled on all WAN links BCNET, CANARIE at SFU WTB:**

-Fully implemented

-**TRIUMF site:** Due to BCNET cannot advertise IPV6 routes smaller than /48 to CANARIE routes we decided to turn off IPV6 on the old TRIUMF site

 IPV6 is enabled  
 IPV6 removed



# new T2s in Canada

- Compute Canada runs 4 computing sites for general research computing needs, will all be 100G connectivity
- researchers / experiments can apply for allocations:  
in 2018: ATLAS, Belle II, DEAP, IceCube, SNO+, T2K

	Arbutus	Cedar	Graham	Beluga
start	2018 <sup>(1)</sup> , 2016 <sup>(2)</sup> , 2014 <sup>(3)</sup>	2018 <sup>(4)</sup> , 201X <sup>(5)</sup>	2017	soon
CPUs	2816xGold 6130 <sup>(1)</sup> 14000xE5-2680-V4 <sup>(2)</sup> 3328xE5-264/50-v2 <sup>(3)</sup>	30960xPlatinum 8160 <sup>(4)</sup> 24064xE5-2683-v4 <sup>(5)</sup>	36064x E5-2683-v4	28000 core
GPUs	2x Tesla K80 <sup>(2)</sup>	584 NvidiaP100 4 per node	160 dual Nvidia P100	688 Nvidia V100
storage	4PB Ceph storage	~4.2 PB disk	5 PB Lustre	13PB disk 50PB tape

# T2 at UVic (a.k.a. Arbutus)

- openstack cluster
- Arbutus cluster to absorb old machines from previous “West” cloud
- external: 100Gbps Huawei switch
- internal: 10Gbps connections to worker nodes
- ATLAS is on LHCONe
- IPv6 not fully deployed
- exploring running kubernetes cluster on cloud (potentially bare metal) for ATLAS T2

# T2 at SFU (a.k.a. Cedar)

- shared facility: ATLAS, SNO+, T2K(soon)
- new cluster: #190 in Top500 list (Nov.2018)
- Centos7.6 on worker nodes
  - ATLAS runs in singularity containers
- external: 100Gbps Huawei switch
- internal: 100Gbps OPA
- ATLAS is on LHCONE
- IPv6 not fully deployed
- SLURM batch system

# T2 at Waterloo (a.k.a. Graham)

- shared facility:
- Centos7.5 on worker nodes
  - ATLAS runs in CentOS6 singularity containers
- external: 100Gbps Huawei switch
  - now connected to ORAN (ORION)
- internal: Infiniband (FDR and EDR)
- ATLAS is on LHCONE (BGP is off for now)
- SLURM batch system
- small openstack cluster for some services

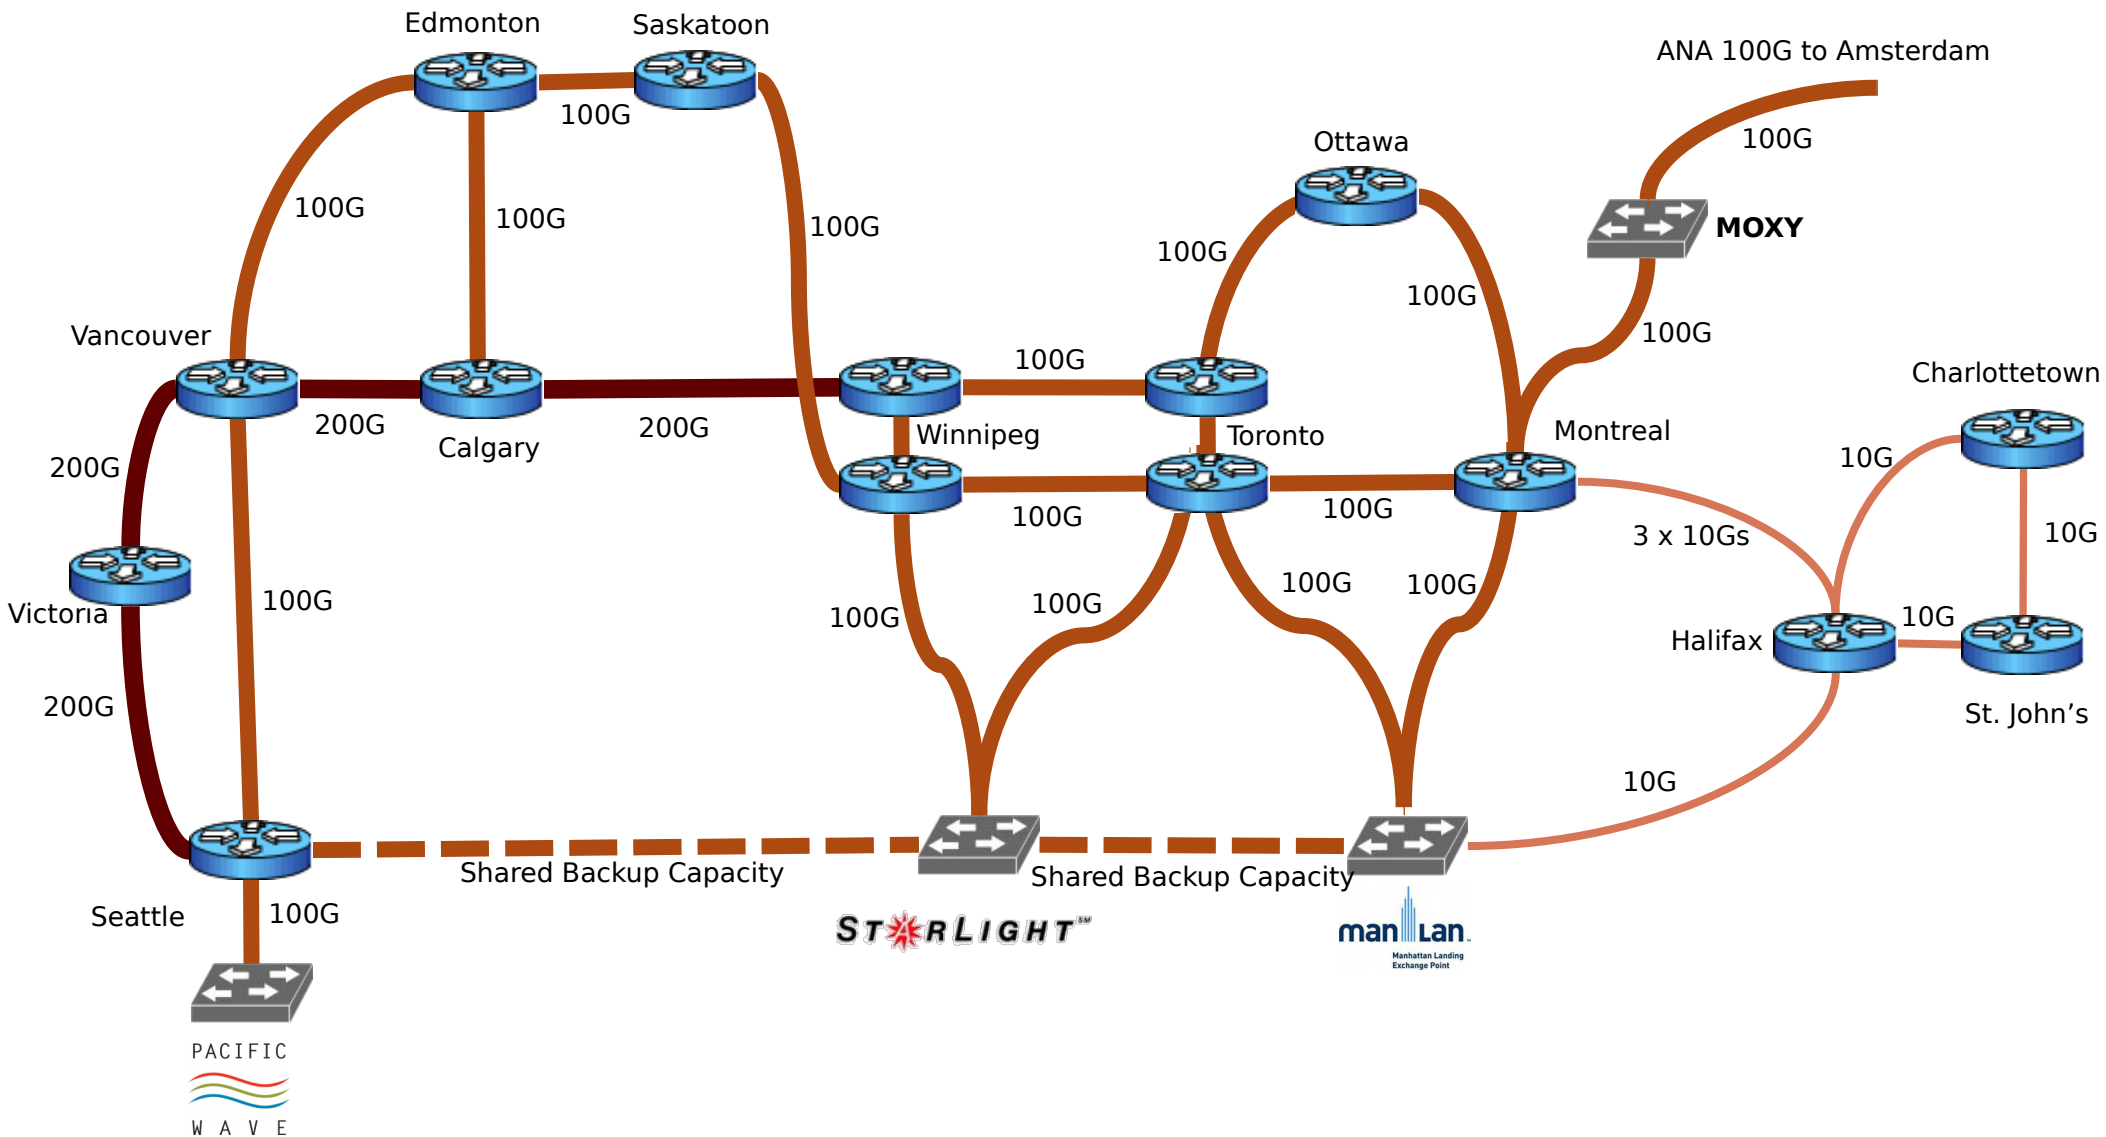


# Beluga

- T3 facility, might be upgraded to T2
- shared facility for ATLAS, Belle-II, SNO+, IceCube, DEAP
- installed & being commissioned
  - but still under NDA ...
- most nodes will have 4GB/core, rest 2GB/core plus small number of 16GB/core



# CANARE IP Core Network



# Research Network in Canada

- 200G on Canarie's own fibre via second wavelength, added second 100G path from carrier on Winnipeg ↔ Toronto
- international connectivity unchanged:  
currently 100G to Asia  
multiple 100G to Europe
- ETA: weeks from now

# Summary

- new T1/T2s up and running
- T3 @ McGill close to open doors
- network upgrade to 200G on main paths