

Status and Concerns from



 **HEPnet Canada**

Canadian T2s


LHC Networking Workshop at CERN

2014-02-10


Ian Gable

Quick Overview

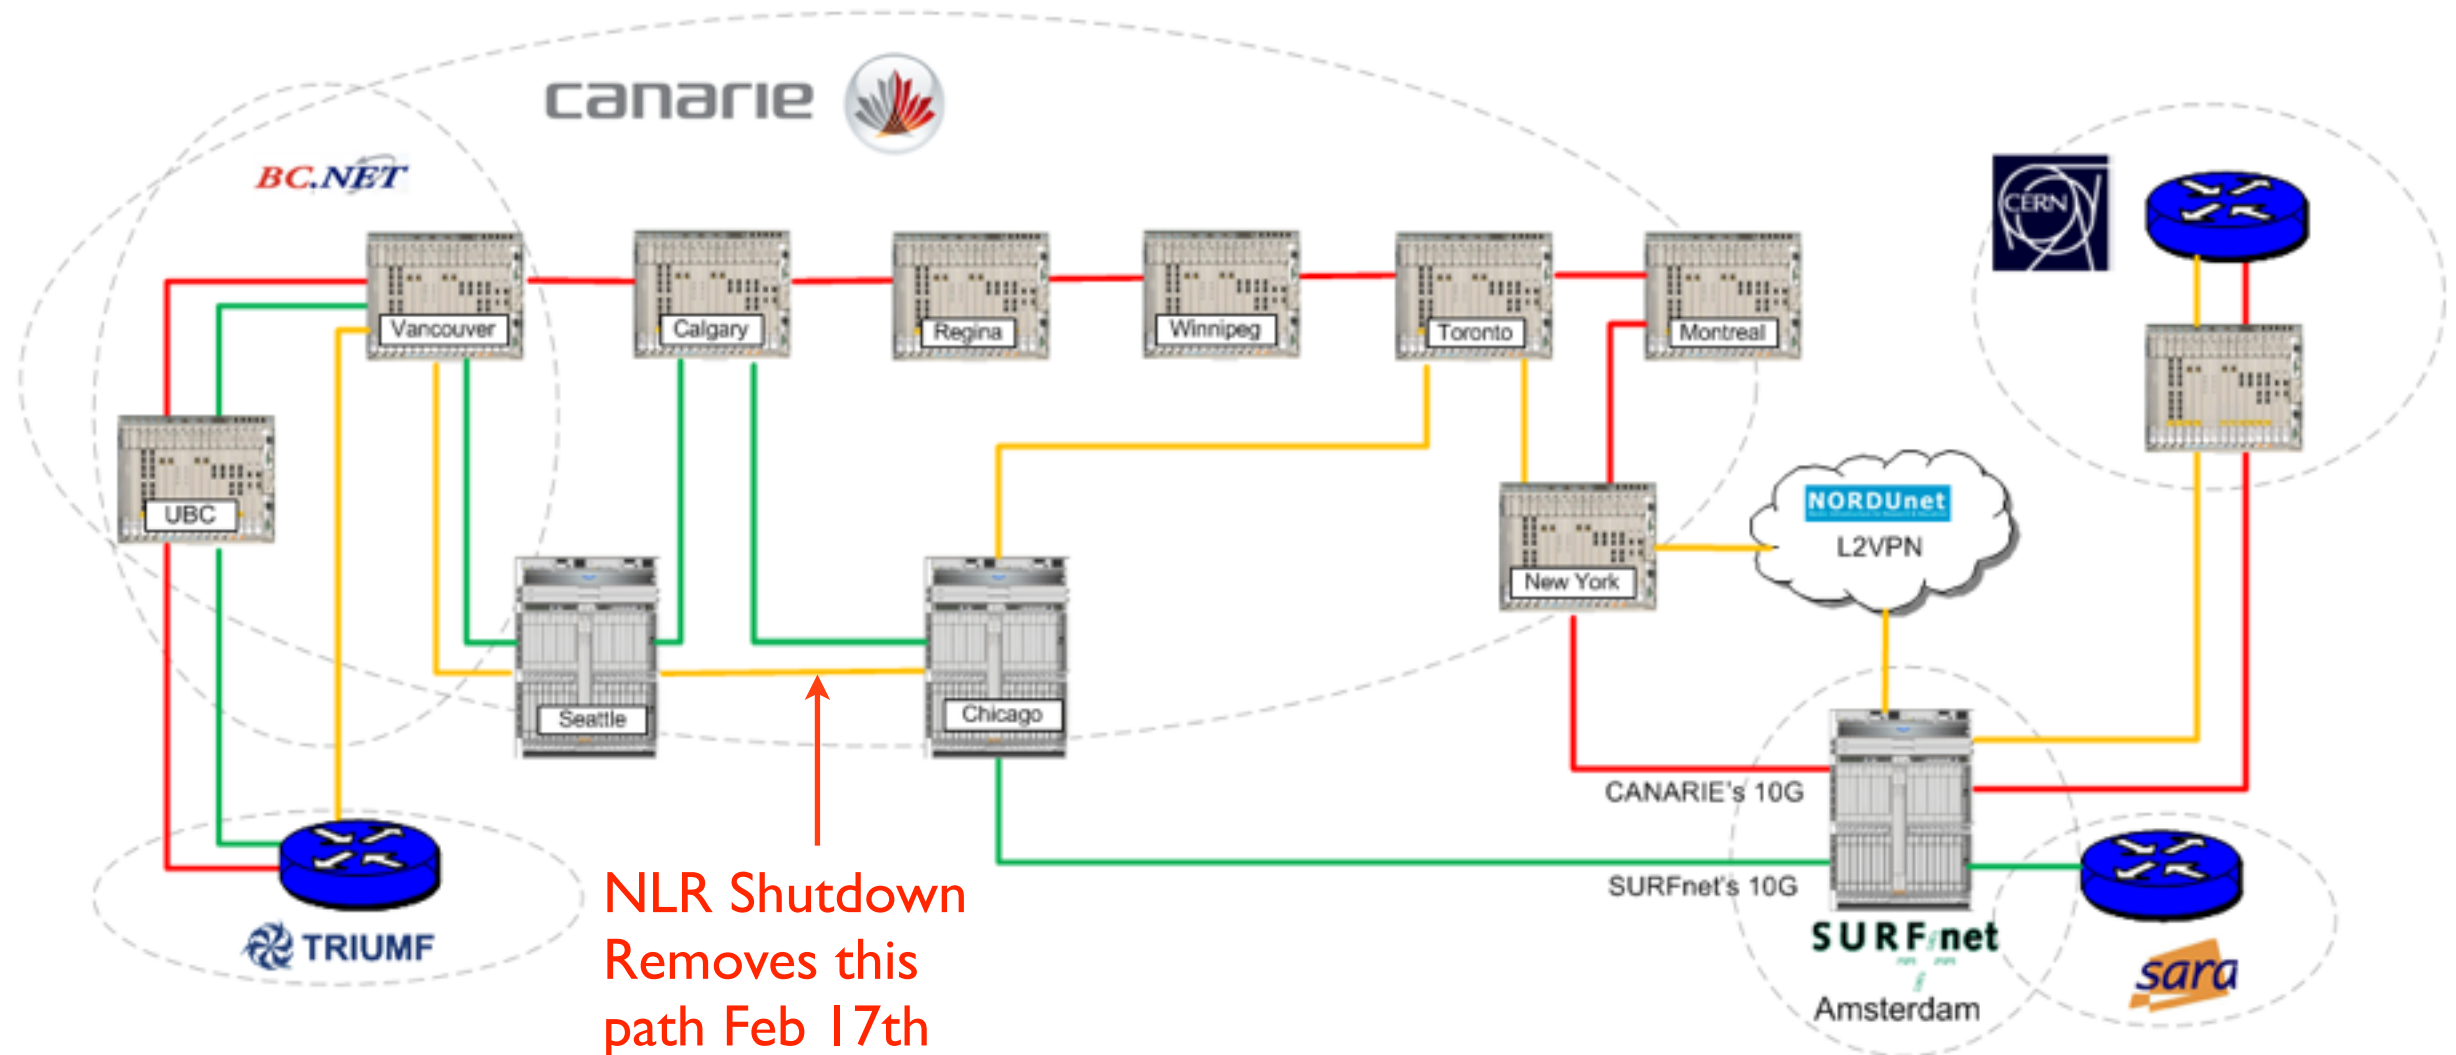
ATLAS Tier I at TRIUMF with 10% of ATLAS Data 8400 km from CERN
in a straight line, 175 ms RTT



4 ATLAS Tier 2s at University of Victoria, Simon Fraser University, University of Toronto, and McGill University



TRIUMF to CERN



— 10GE Circuit (5GE Lightpath)

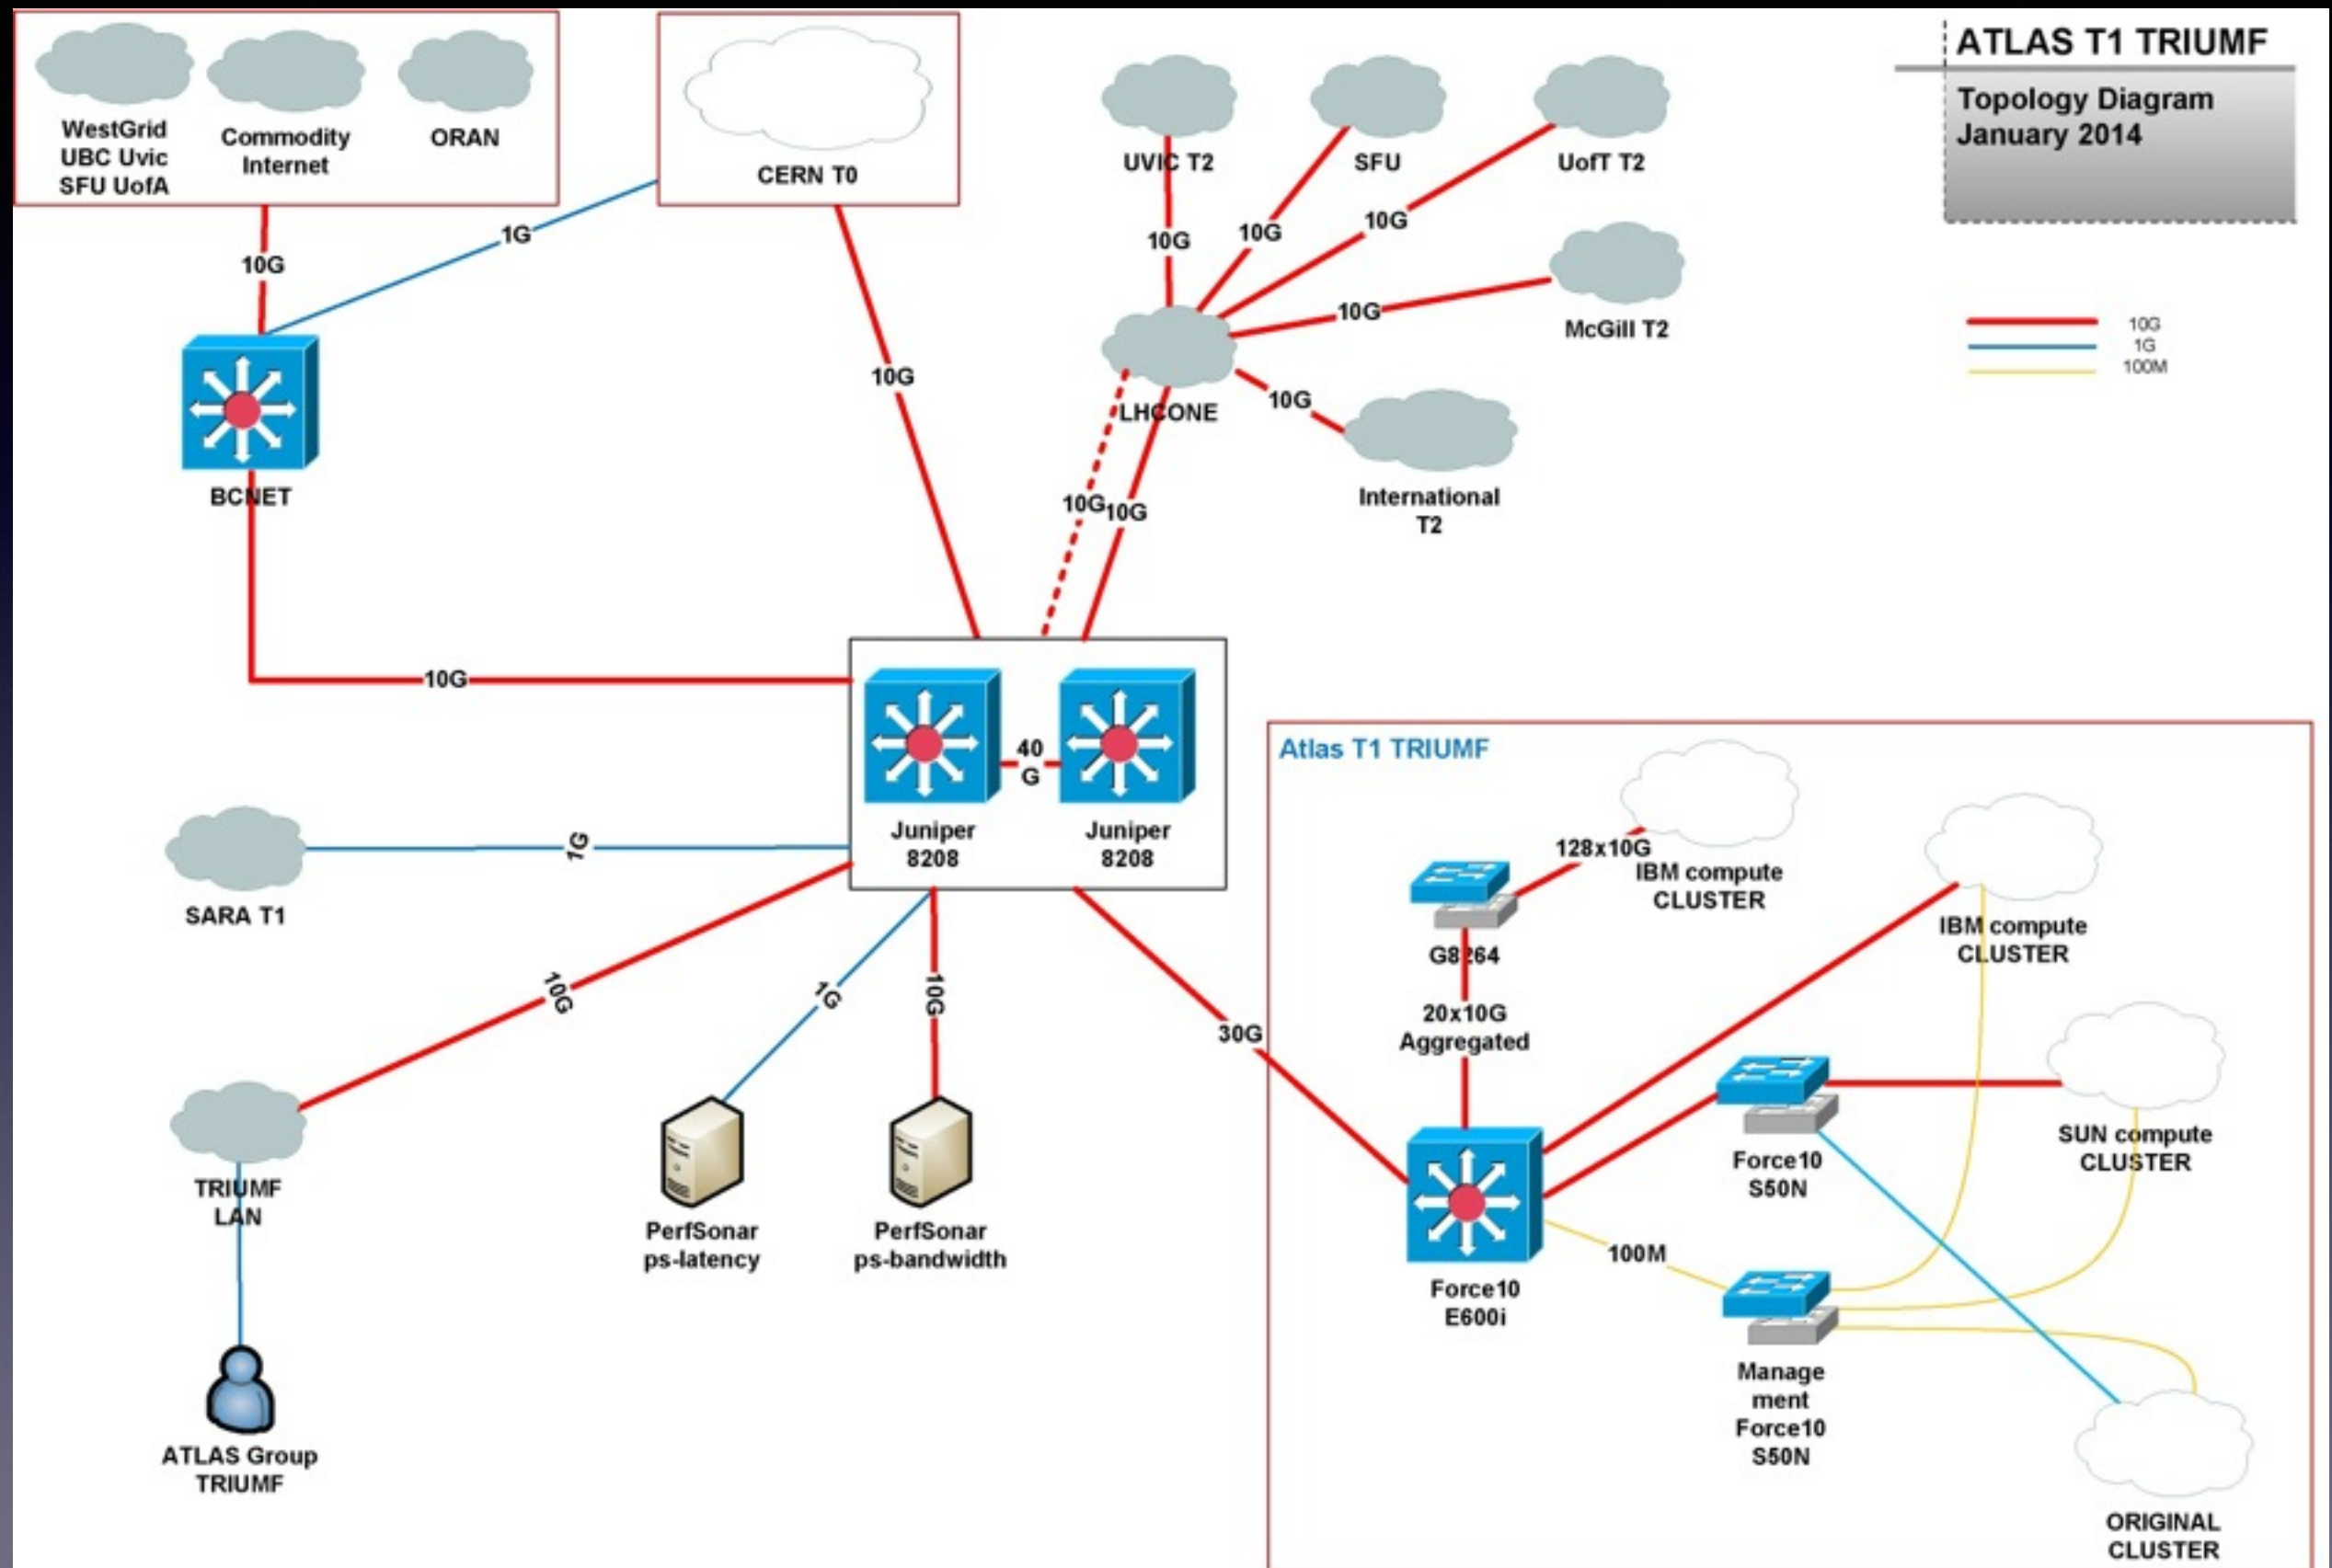
— 1G Lightpath to SARA

— 1GE Lightpath to CERN via NORDUnet/SURFnet

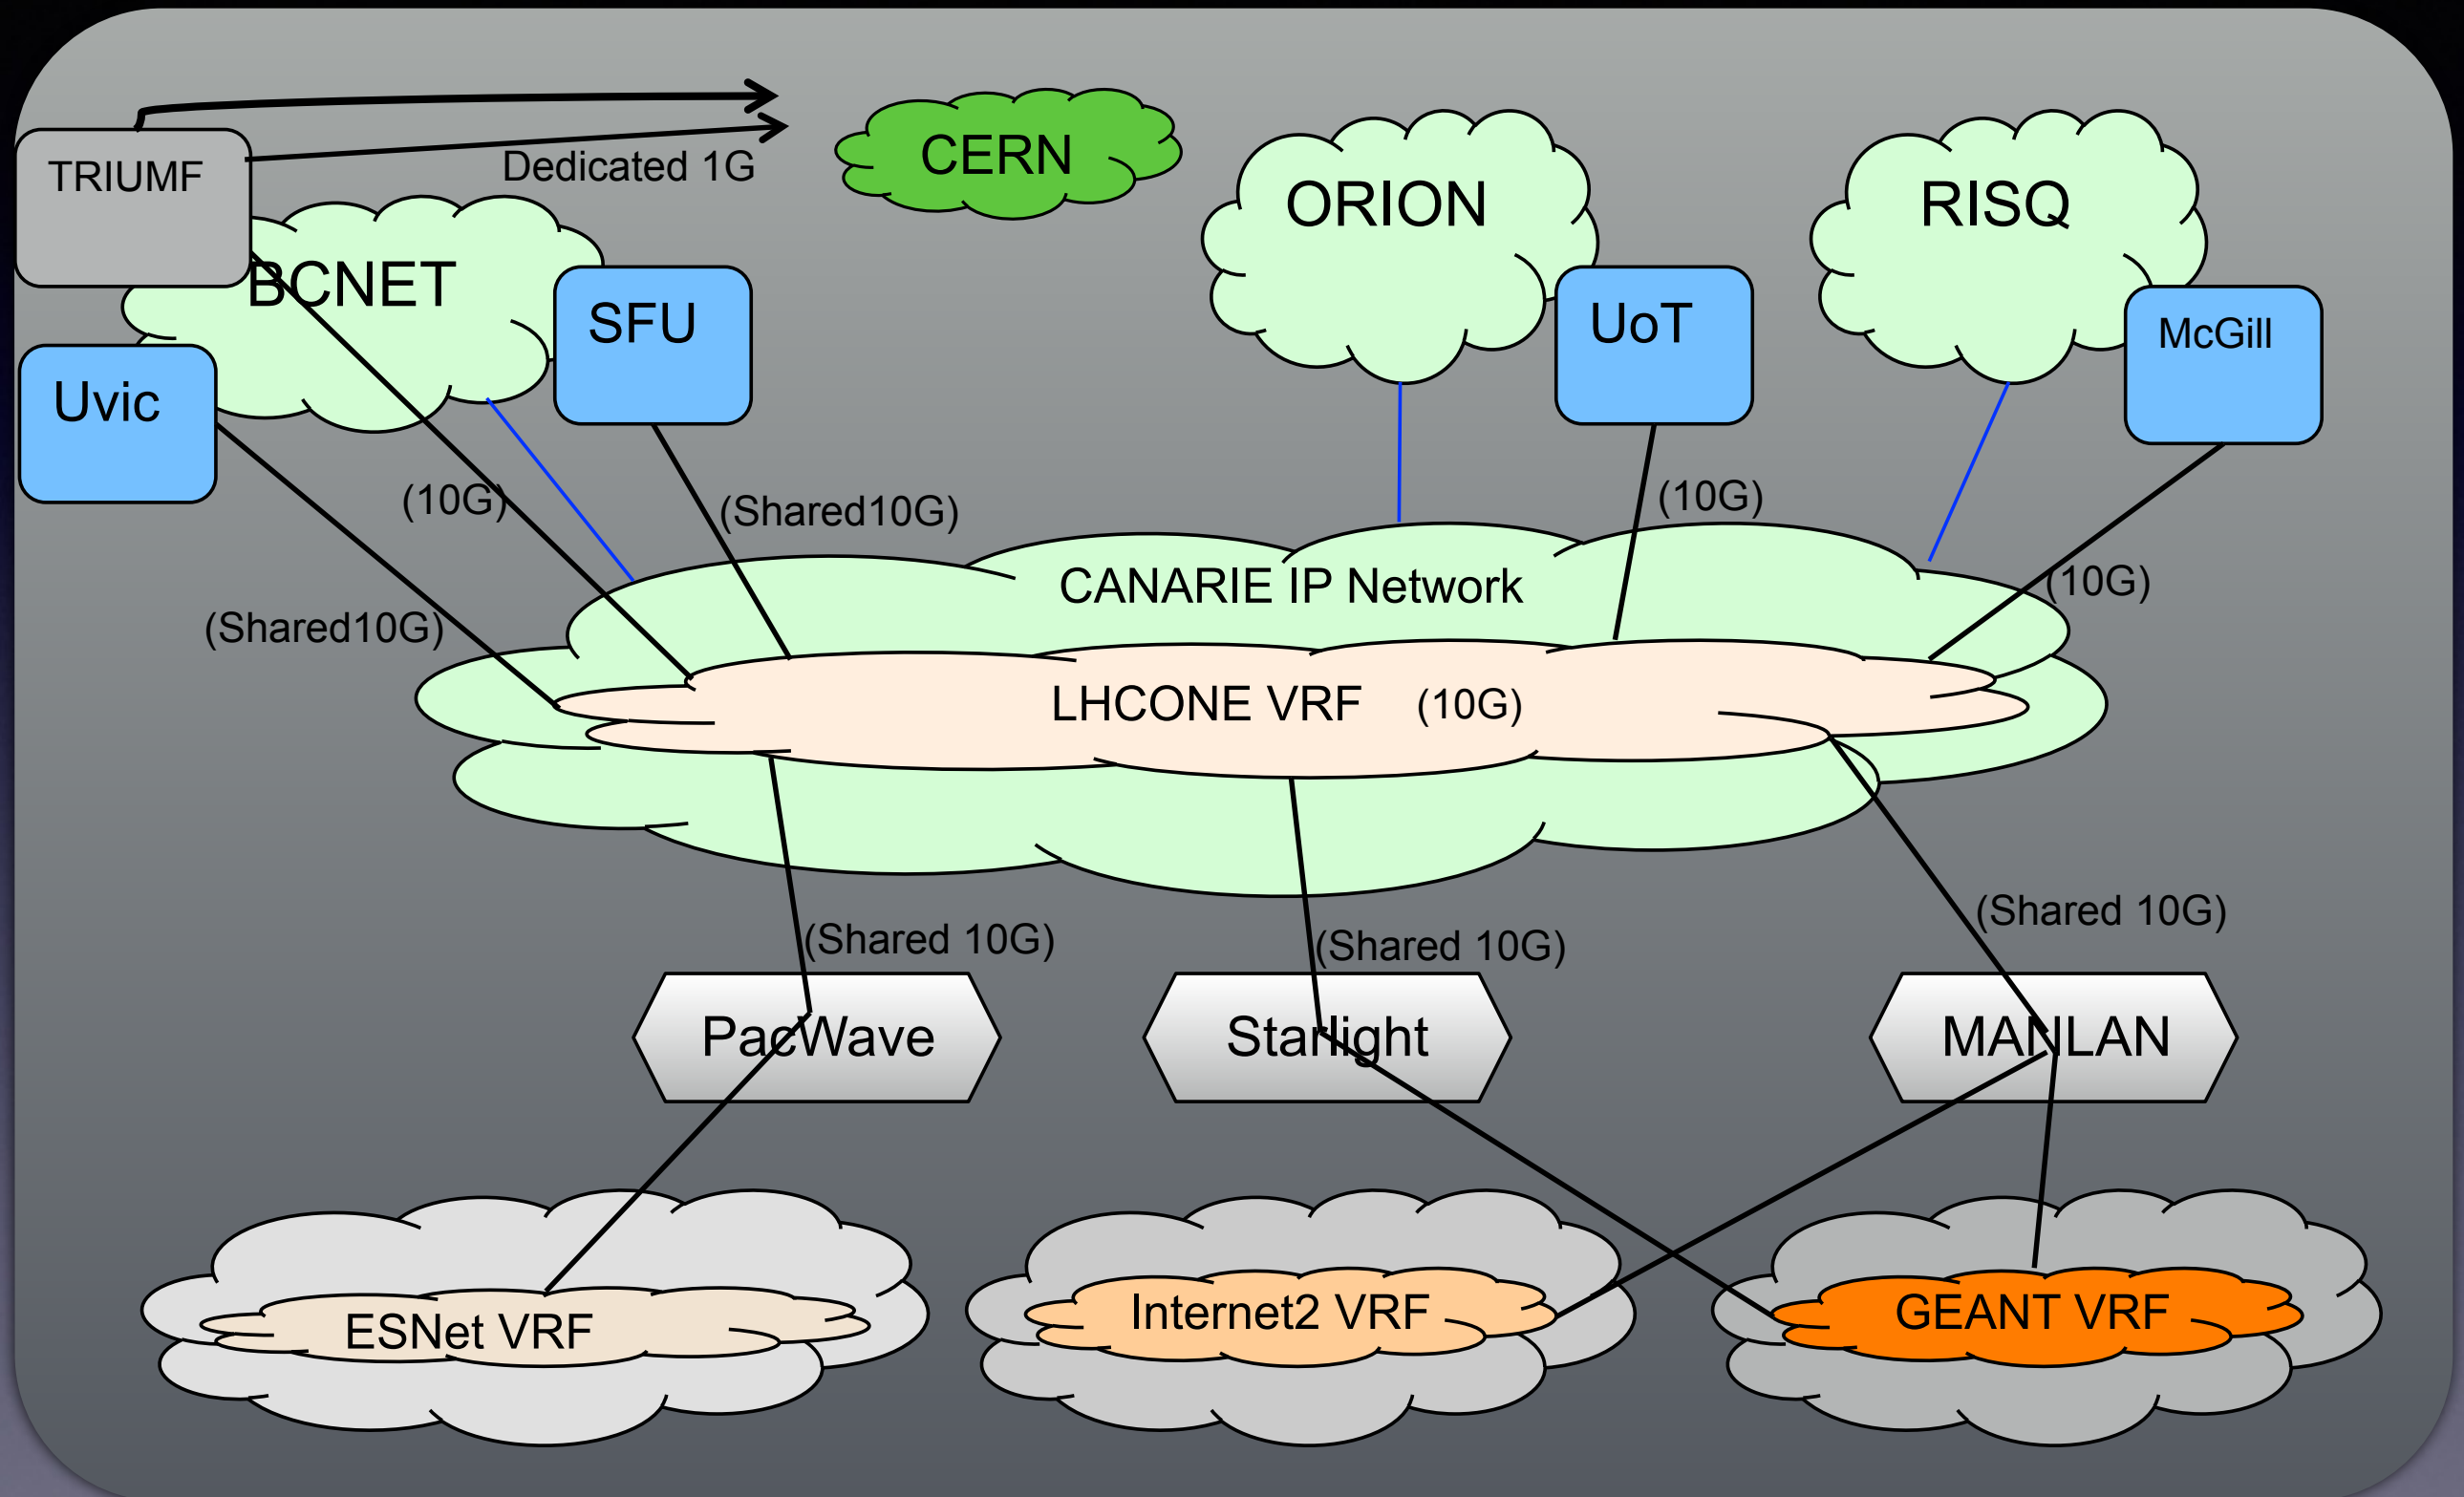
canarie		
TRIUMF LHCOPN Connections		
Author	Thomas Tam	1.0
Date	15 Jan. 2014	1 of 1



ATLAS T1 at TRIUMF



LHCONE + Canada Internal LHC Network



Canadian Sites: happy so far

- All sites feel well served by the R&E Networking community.
- In 2012 we moved from using point-to-point circuits to connect TRIUMF - Canadian Tier2s to using the LHCONE within Canada.
 - immediately boosted path utilization and increased performance
 - prevented East coast T2s from communicating with each other via TRIUMF 4000 km away.
 - 2013 CANARIE provisioned a second, dedicated 10G circuit for LHCONE in Canada
 - Additional 10G to TRIUMF LHCONE being added now.



Concerns for the Future: TRIUMF Backup

- Today backup is provided by a 1G circuit on a separate path from primary 10G.
- 1G no longer a workable capacity
- 10G backup fail over path is expensive Vancouver to Geneva
- TRIUMF LHCONE capacity continue to grow
- Can we use LHCONE as backup?

Last month primary 10G was down
9 times, totalling 11h 38m on backup path

Start	End	Duration
1/2/2014 16:53	1/2/2014 17:03	0d 0h 10m 10s
1/2/2014 22:50	1/2/2014 23:00	0d 0h 10m 10s
1/11/2014 16:07	1/11/2014 22:02	0d 5h 55m 50s
1/21/2014 22:30	1/22/2014 0:00	0d 1h 29m 38s
1/22/2014 0:00	1/22/2014 3:15	0d 3h 15m 2s
1/23/2014 23:39	1/24/2014 0:00	0d 0h 20m 8s
1/24/2014 0:00	1/24/2014 0:10	0d 0h 10m 22s
1/28/2014 1:06	1/28/2014 1:27	0d 0h 20m 20s




Networking for Infrastructure-as-a-Service

ATLAS Cloud Operations
already running on ~15 IaaS
Cloud around the world

Effort from personnel at
many sites

No specialized networking
available at these sites, and
the **number of sites is sure
to grow**

I/O already limits the types
of jobs we can run at these
sites.



How to we deliver the LHCONE to IaaS sites on demand?



Grab Bag of collected concerns

- How soon will 100G be the requirement to be a 'first class' site. Some T2s already there.
- If a point-to-point service evolves will it be practical to implement for T2s where human resources are thin on the ground?
- People are interested in the P2P discussions, but are having a hard time following the discussions on the mailing lists.
- CANARIE and Regional networks would like to know what they need to prepare for if P2P functionality becomes a requirement for sites.



Grab Bag continued

Leslie Groer Canada Tier 2 Sites Coordinator:

“A concern I have that is not always reflected in any explicit resource planning from the experiments is what the network capability needs to be for a T2 site. **This is not just pure bandwidth but the peer-to-peer performance on individual and aggregate file transfers (which is what really matters) to all the other sites.** The T2D requirements from ATLAS a few years ago were somewhat ad-hoc and there were never any specifications we could have used in planning the various facilities. We have projections that go out a few years for CPU and storage; there should be similar projections for network requirements and performance.”



Summary

- Canadian sites have been well served thus far by R&E Networking community
- LHCONE L3VPN been very beneficial to sites within Canada, especially internally
- Backup over long distance for T1 is expensive, looking for better solutions
- How can we deliver specialized LHC networking to IaaS sites?
- Better specifications of Tier 2 to disk-to-disk transfer performance



Backup Slides

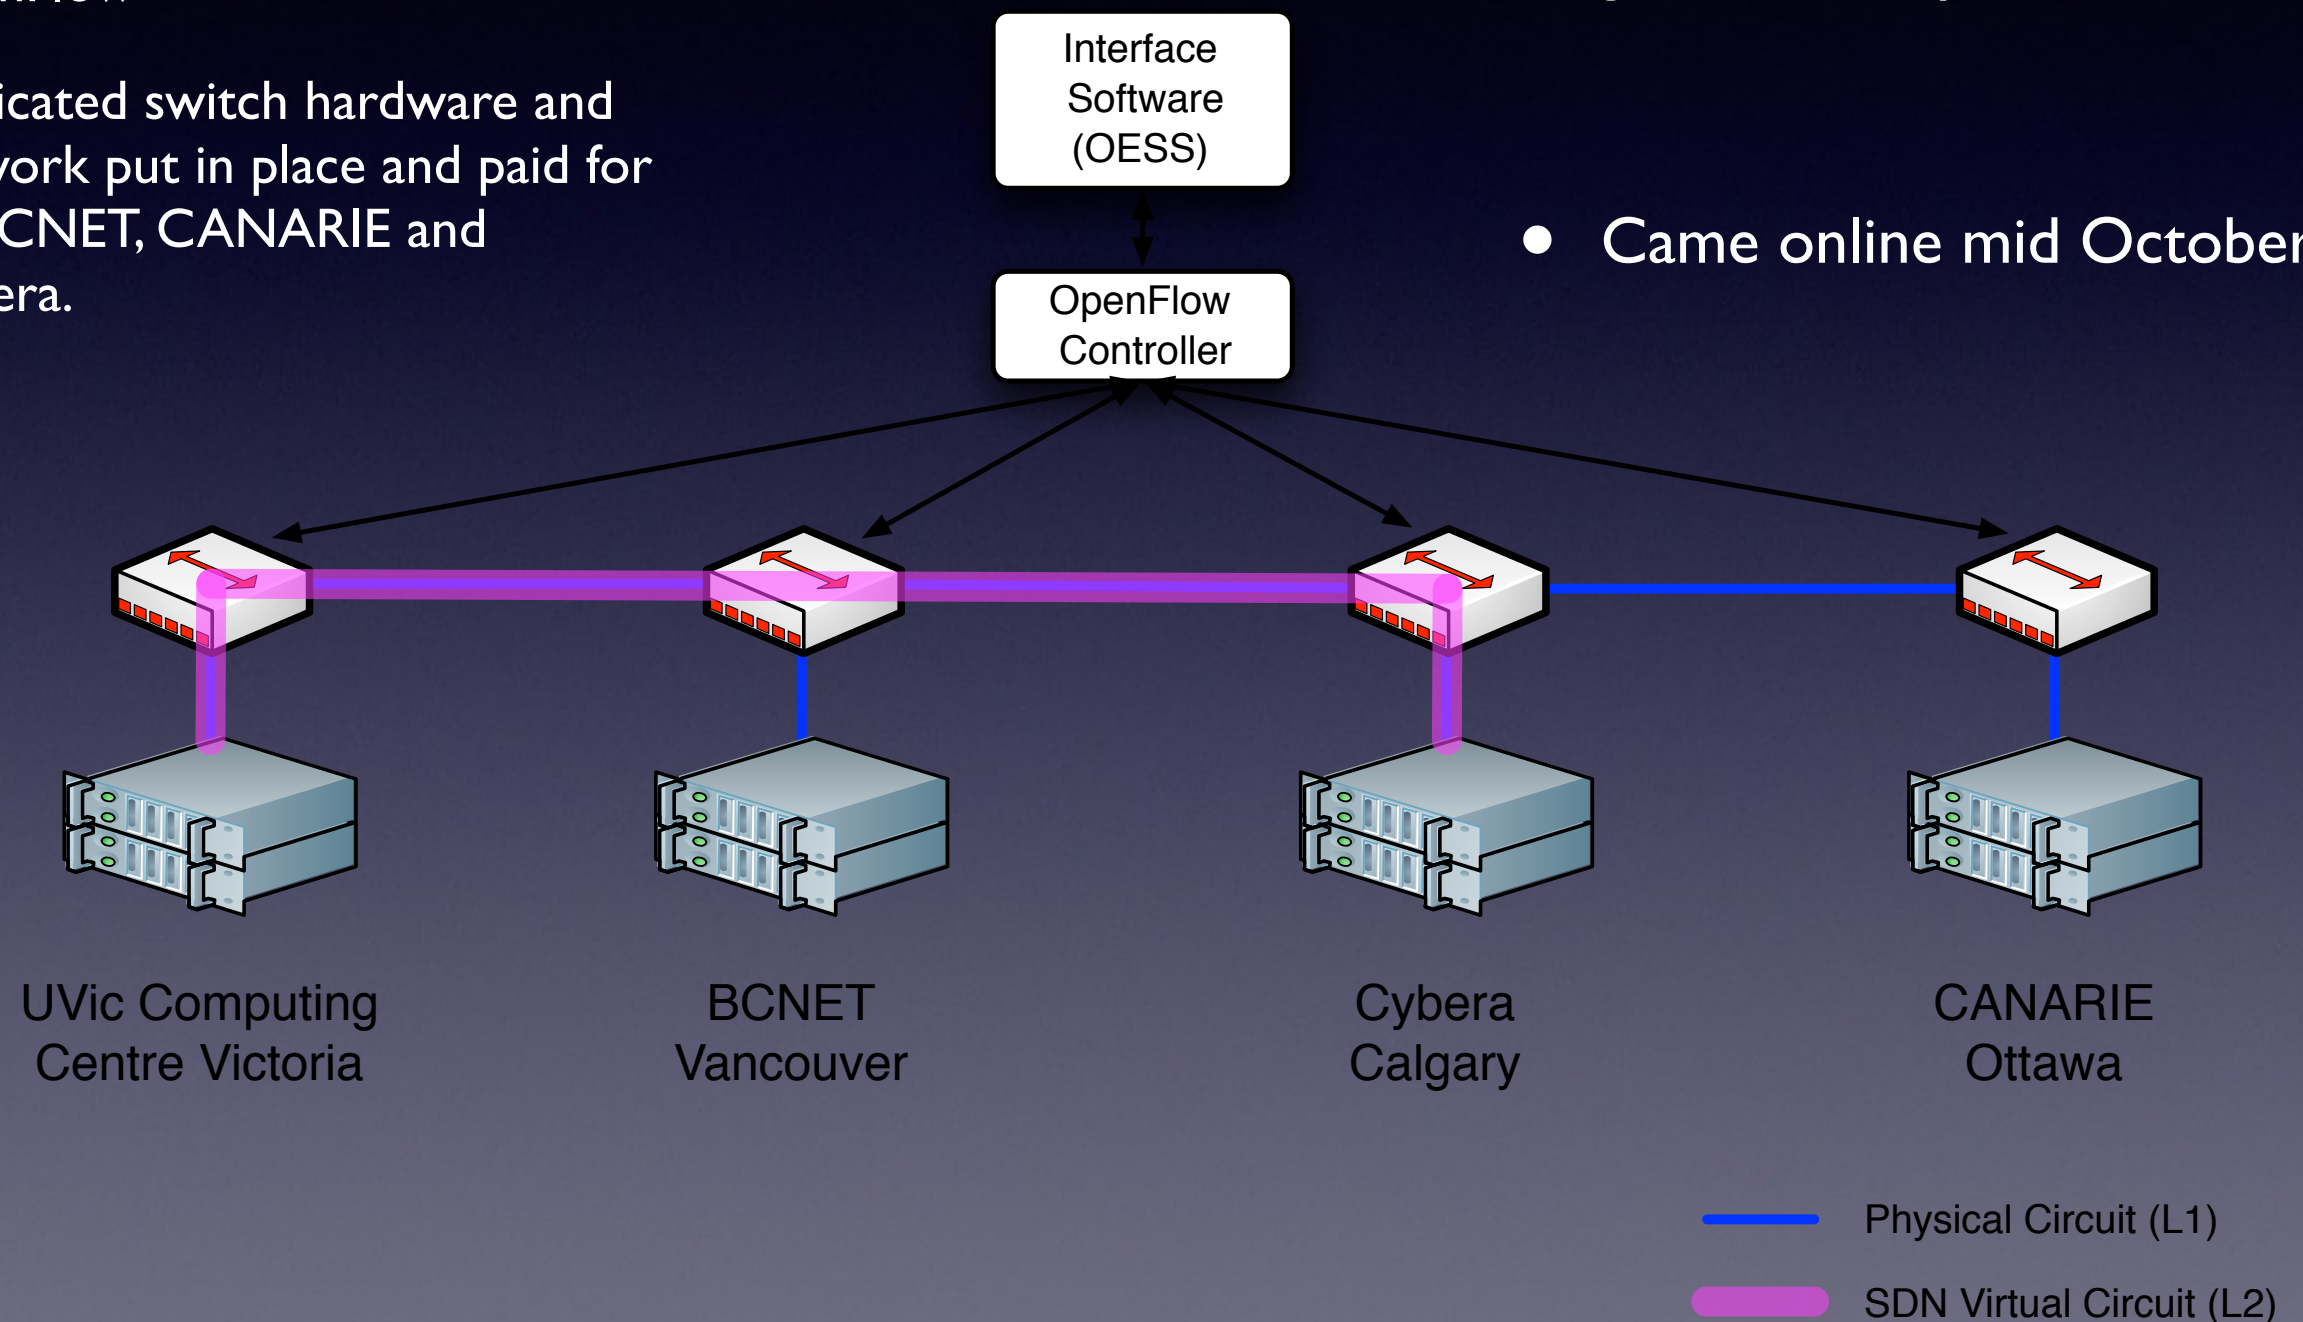


Software Defined Network Testbed

- HEPnet has established a SDN testbed to explore capabilities offered by SDN in Particular OpenFlow
- Dedicated switch hardware and network put in place and paid for by BCNET, CANARIE and Cybera.

- Goal is also to help Canadian R&E networking community gain familiarity with SDN


- Came online mid October



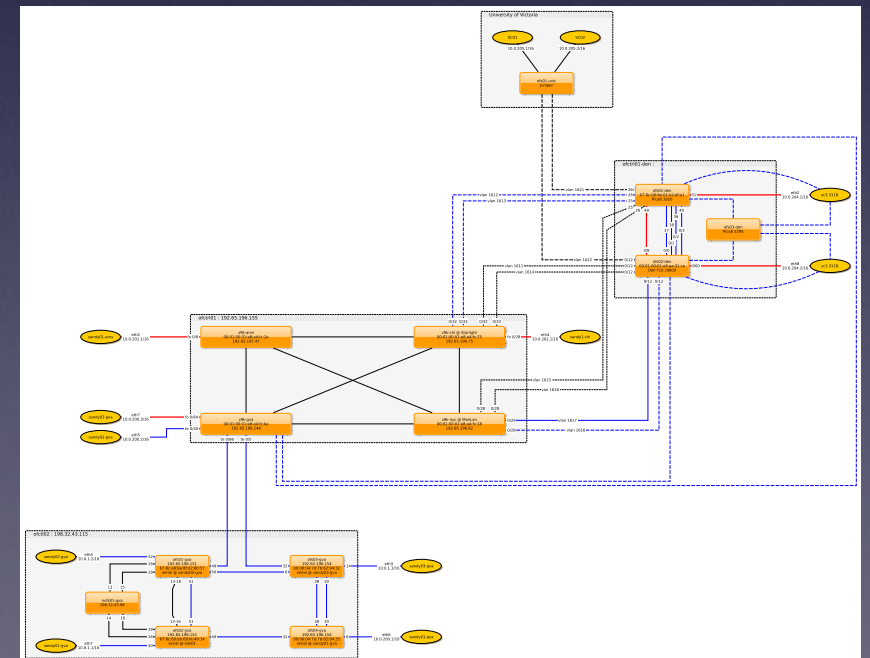
SCI3 SDN Demo



- Intercontinental OpenFlow Testbed formed for SCI3. Links Victoria – Denver- São Paulo – Chicago – New York – Amsterdam – CERN



- Link Layer Multipath Switching



- <http://supercomputing.uvic.ca>

