Canadian ATLAS Tier-1 Report

Di Qing

On behalf of TRIUMF ATLAS Tier-1 team

HEPiX Autumn 2023 Workshop

October 16-20, 2023

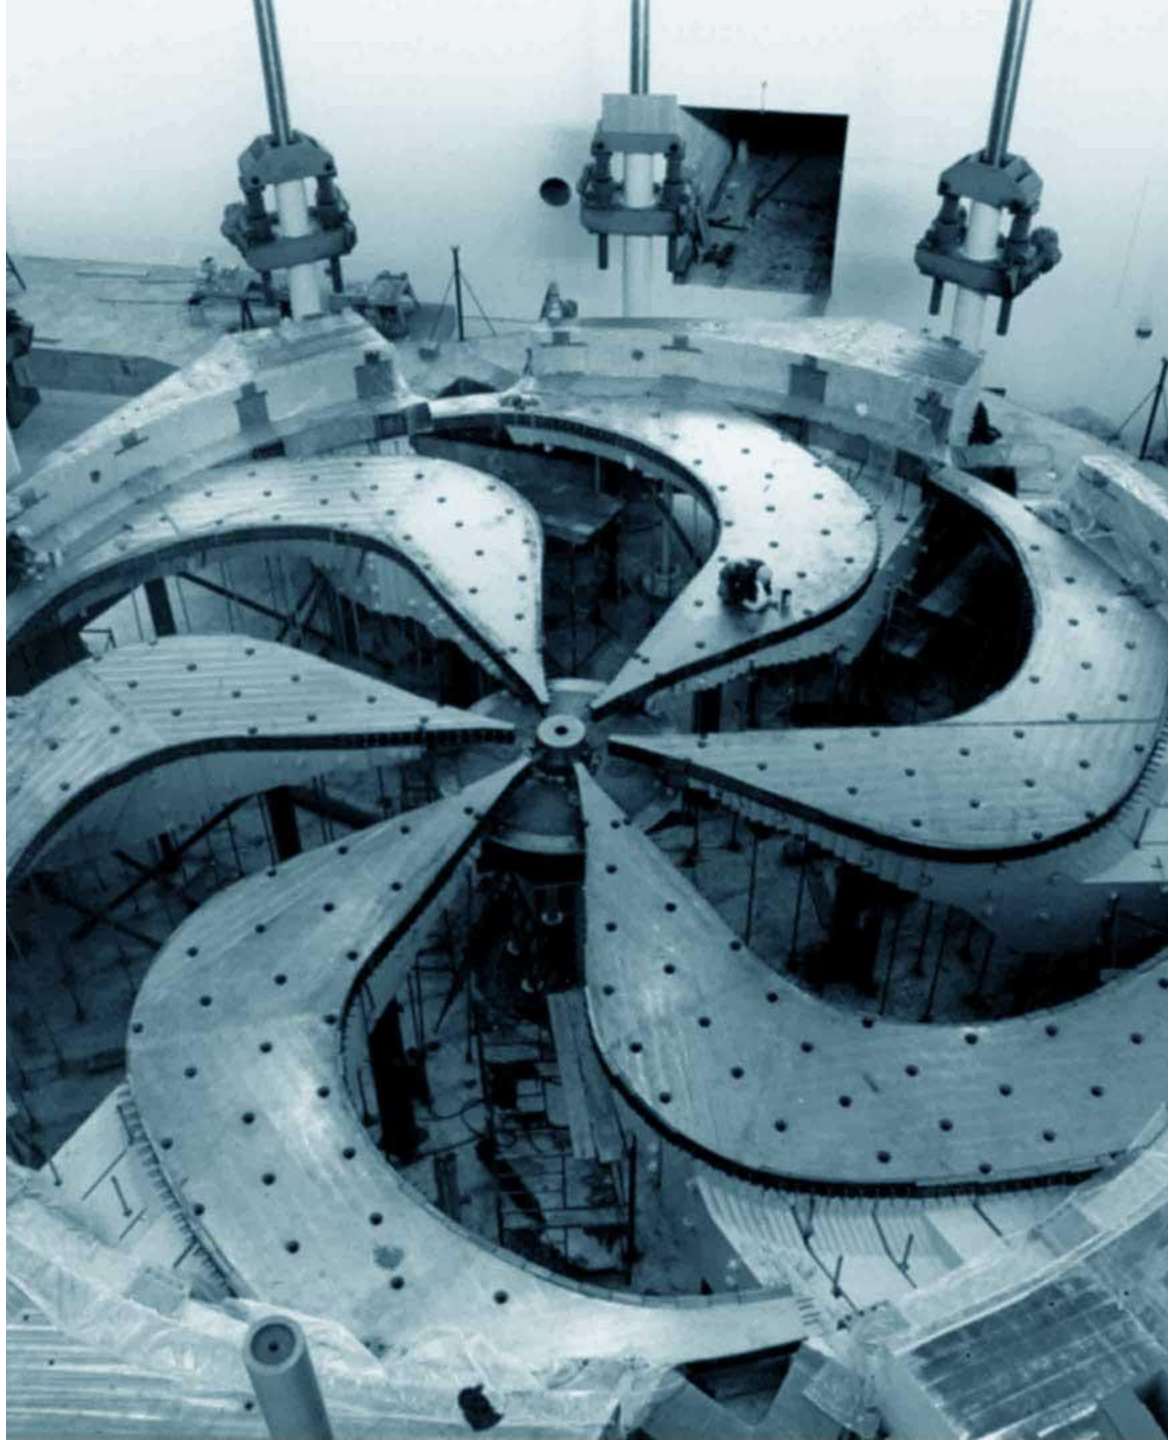

2023-10-16

Discovery, accelerated

Canadian ATLAS Tier-1 Overview

- Providing 10% of worldwide Tier-1 resources for ATLAS
- A ‘Federated’ Tier-1
  - Primary Tier-1 services and resources relocated from TRIUMF to a new data centre at Simon Fraser University since 2018
  - Some old equipment at TRIUMF are still providing computing resources
- Current capacity:
  - 7824 cores (SFU) + 1128 cores (TRIUMF, simulation only)
  - 17.1 PB disk (SFU)
  - 36 PB tape (SFU)
- An additional 11.5 PB of tape capacity expected to be delivered soon

Legacy TRIUMF Data center still running but …

- Limited space
- All equipment out of warranty
- Major cooling devices stopped working
- Another project planning to use part of the space
- Available CPU cores dropped from ~4800 to ~1100 in three years

SFU water tower building data center

- Share the same data center infrastructure with Compute Canada/Digital Research Alliance of Canada Cedar shared HPC facility including SFU ATLAS Tier-2
- New data center infrastructure aspects are the responsibility of SFU data center team
  - MoU and SLA between TRIUMF and SFU in place
- TRIUMF Tier-1 personnel continue to be responsible for the Tier-1 operations
- Core services are in HA area with UPS and diesel generator

Tier-1 Networking

- Collaborative effort between TRIUMF, BCNET, CANARIE and SFU
- Juniper QFX10008 as core switch for both internal and external network
- The connectivity of external network goes through another shared Cedar core switch (Arista)
- Tunnel between SFU and TRIUMF for the old compute nodes at TRIUMF
- IPv6 fully implemented at Tier-1 at SFU location
  - Dual stack on all storage servers

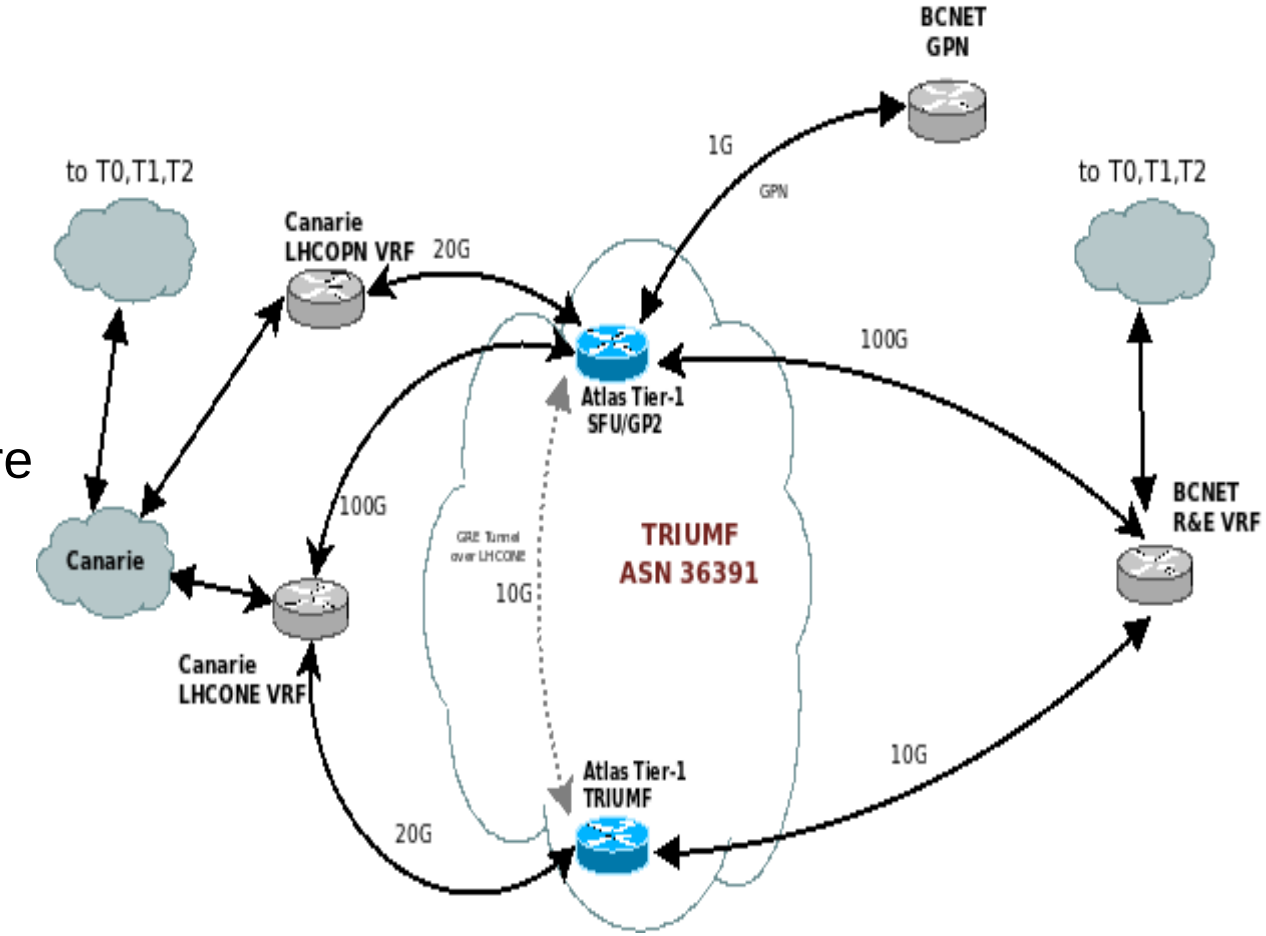

ATLAS Tier-1 cluster at SFU

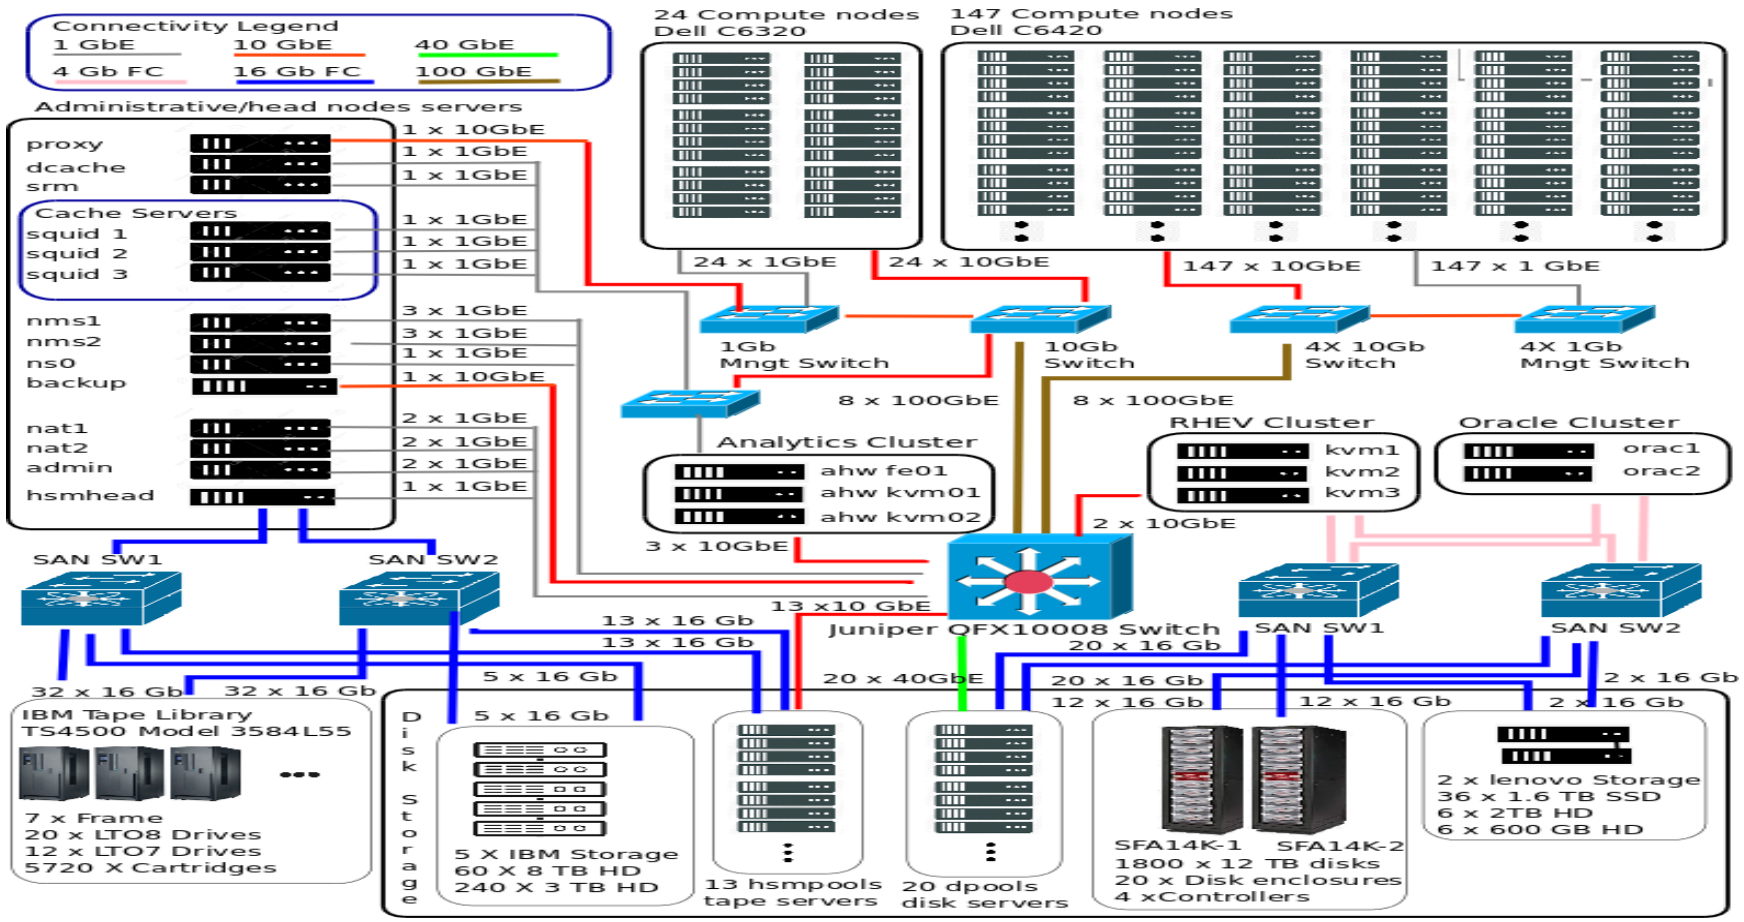

Storage systems

- DCache 7.2.40
- 16 PB disk capacity, 13.6 PB populated for production
  - 2 X DDN SFA14kx, 2 X 24 X 16Gb Fibre connections
  - 20 Lenovo servers which has dual port 16Gb HBA, 16 cores,196GB ram, 40Gb network
- 36 PB tape
  - IBM TS4500 tape library
  - 20 LTO8 drivers, 12 LTO7 drivers
  - 3319 LTO6, 1000 LTO7, 1820 LTO8 tape cartridges
  - Another 11.5 PB LTO8 tape cartridges coming within 1-2 weeks
  - 1PB disk buffer
  - one HSM head node,13 HSM pool nodes, 10GbE, dual HBA 16Gb/port
- Webdav and xrootd protocols for internal data access, and webdav protocols for external data access
- In 2022, ~105 millions transfers, ~145PB(WAN + LAN) data transferred, 10x times more than total disk capacity

Compute nodes and batch systems

- 2 HTCondor batch systems, one for compute nodes at SFU and one for the old compute nodes at TRIUMF
  - Each cluster has two HTCondor head servers with HA configurations plus two ARC CEs
- 7824 cores at SFU
  - 24 DELL C6320 with 32 cores
  - 147 DELL C6420 with 48 cores
  - 4GB RAM/core and ~43GB/core SSD disk space
- 1128 cores at TRIUMF
  - Out of warranty

Updates on other services/resources

- Virtualization system
  - 240 logical CPUs, 564.5 GB RAM and 14.7 TB storage
  - 41 virtual machines, majority of non storage Tier-1 servers
    - 123% of 240 logical CPUs committed, and used 239GB RAM and 10.2 TB storage
  - RedHat Virtualization (RHV) on three hypervisors
    - Very stable and good support
  - Current license only valid till the end of 2023
    - Need to find a replacement
- Frontier servers and Oracle Database
  - Decommissioned as all ATLAS frontier servers only hosted by CERN now
- Elastic search system
  - Migrated to new and more powerful hardware
  - See Fernando’s talk for details
- CVMFS stratum 1 server for EGI VOs
  - Now 32 small VOs

Cyber Security (I)

- Joined TRIUMF cyber security team
  - Closely cooperate with TRIUMF on cyber security aspects
  - Receiving feeds from the Canadian Shared Security Operators Centre
  - Access to CANSSOC MISP portal and SLACK community channels for cyber security information exchange and discussions
- Asked BCNET to scan our subnets and services regularly with Nessus for any potential security risks
- Reviewed Tier-1 web pages and added different level of protections based on the contents
  - Three levels: public, research community, internal only
  - Currently controlling the access through X509 user certificates

Cyber Security (II)

- A tool developed to monitor the centralized syslog and detect any suspicious SSH connections
- Syslog injected to Elastic Stash with Filebeat for visualization of all connection attempts
- Plan to do
  - Implement Multifactor Authentication for SSH connections
  - Implement filter on the edge router to block attempts from suspicious IPs
  - Also plan to check what we can do for WLCG community, for example, test pDNS

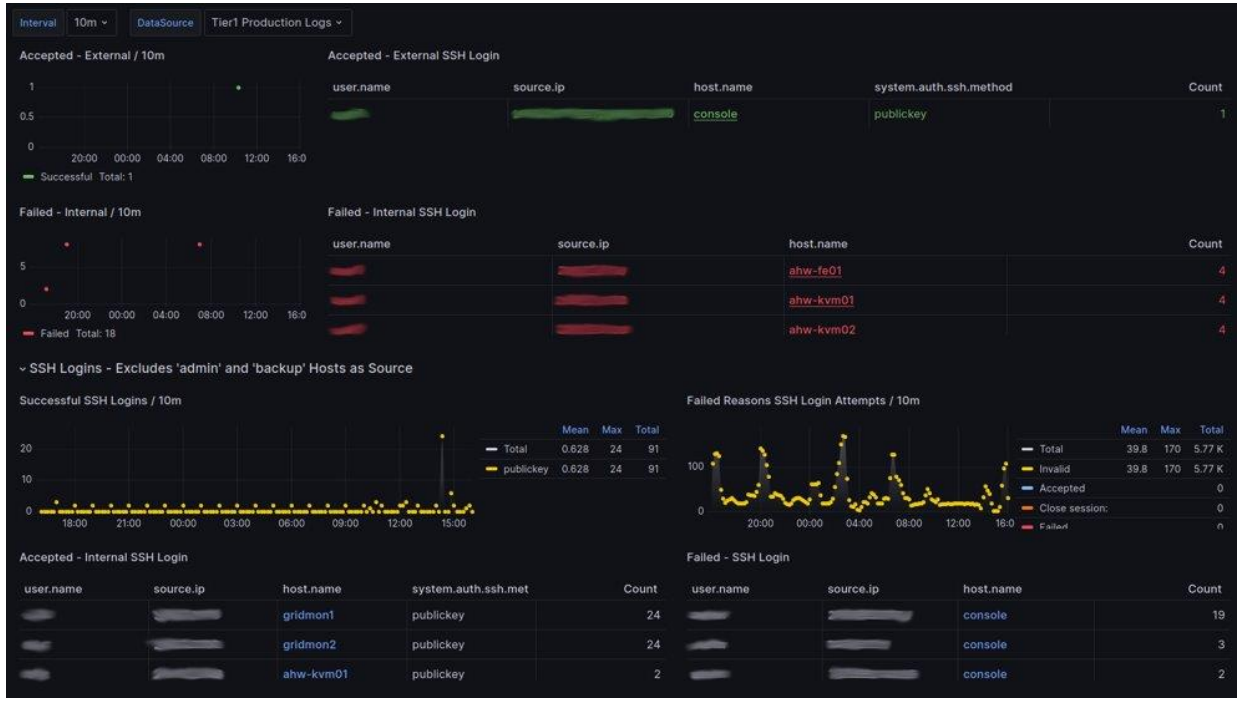

Operating systems

- Got an Academic license package with RedHat through CERN extended research network
  - Sufficient for our future operations
- Majority of nodes are still with SL7
- Started to install new nodes with RHEL9
- Plan to migrate majority of Tier-1 nodes to RHEL9 next year
- Was planning to use Alma Linux in parallel
  - Ubuntu or other Linux distributions instead?

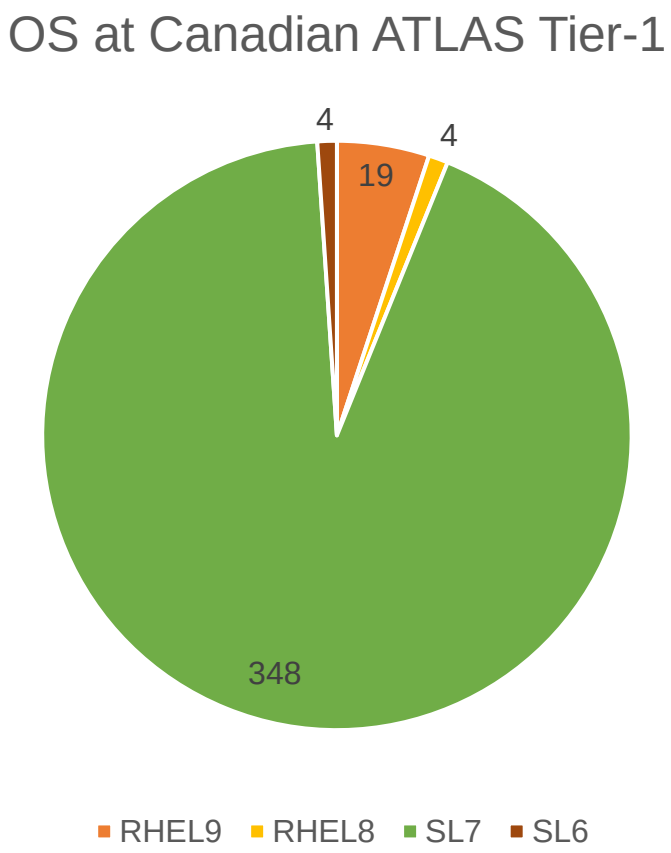

Thank you
Merci

Thanks to the TRIUMF ATLAS Tier-1 team for contributing to this presentation:

Asoka De Silva, Rajan Devbhandari, Fernando Fernandez Galindo, Vitaliy Kondratenko, Simon Liu, Yun-Ha Shin, Reda Tafirout, Andrew Wong

www.triumf.ca

Follow us **@TRIUMFLab**

Backup slides

‘Federated’ Tier-1

- Our computing resources are crossing two remote sites with distance ~22.8km
- All services at both locations are in the same triumf.ca network domain
- All storages are at SFU location
- 7824 cores at SFU location
- 1128 cores at TRIUMF location
  - Out of Warranty
  - Stage in/out the data through a 10Gb GRE tunnel
  - Created a dedicated ATLAS Panda Queue
  - Run simulation jobs only to reduce data access

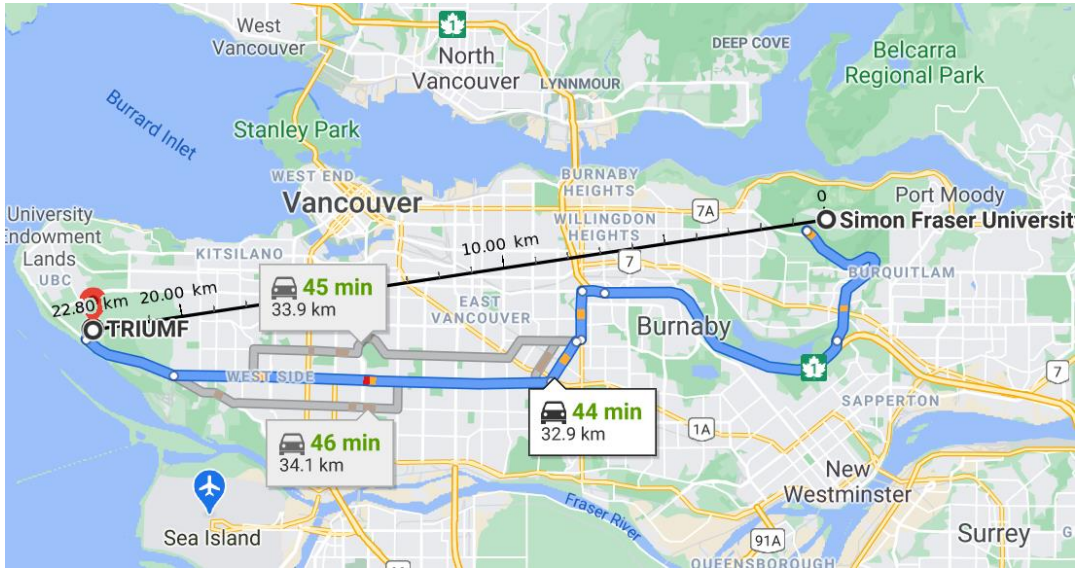

Hardware concerns

- Current DDN systems
  - Too many drive failures – might be caused by a batch of problematic disks
    - Firmware update will be scheduled soon to mitigate this issue
  - And other component issues, controller, IOM, cable issue
  - Support response time too long, no local parts center, no local engineer
- Current Library
  - Too many media related issue, replaced many drives, discarded a lot cartridges, but issue still there
  - IBM is taking actions and sending 4 drives for further analysis
- DELL C6420 compute nodes
  - Too many CMOS battery failures
  - There were also some SSD, CPU, Motherboard and RAM failures