# Canadian ATLAS Tier-1

Di Qing

On behalf of TRIUMF ATLAS Tier-1 Team

HEPiX Spring 2021 online Workshop

March 15-19, 2021

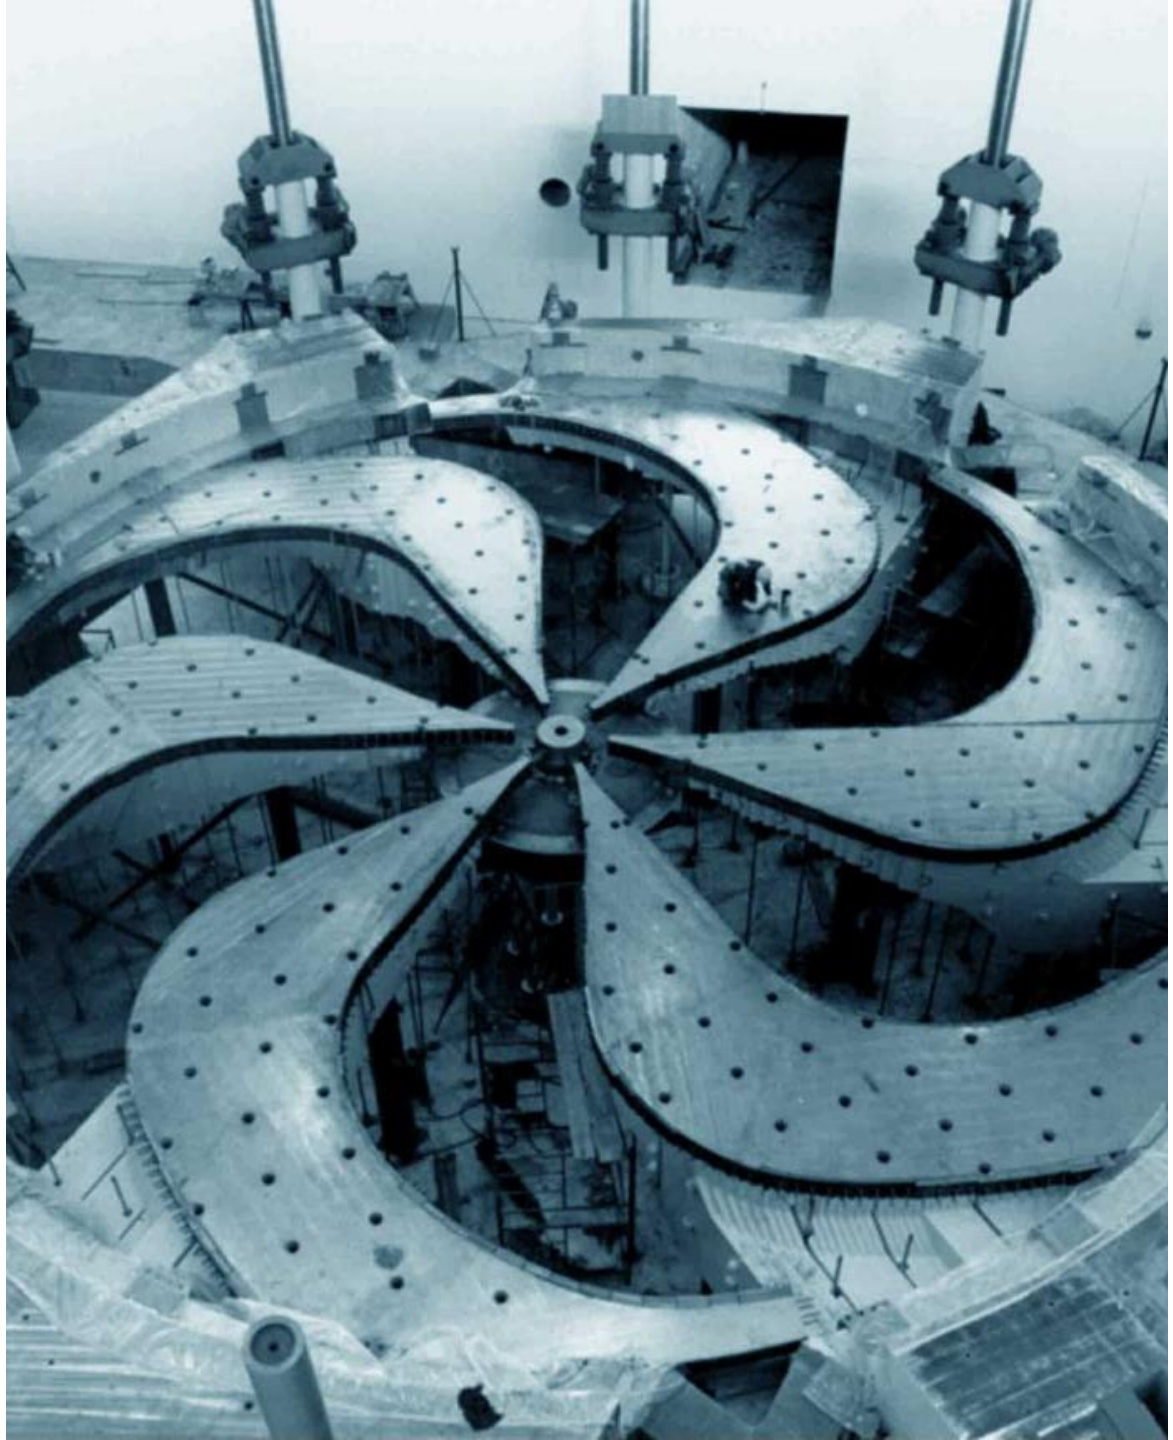

2021-03-15

Discovery, accelerated

# Canadian ATLAS Tier-1 Overview



- Key player in large-scale distributed computing (ATLAS experiment only)
- Providing 10% of worldwide Tier-1 resources
- Primary Tier-1 services and resources relocated from TRIUMF to a new data centre at Simon Fraser University since 2018
- Current capacity in production:
  - 7680 cores (SFU) + 4744 cores (TRIUMF, simulation only)
  - 11 PB disk (SFU)
  - 31 PB tape (SFU)
- Recently funding secured for future expansion
- Stable 24x7 operations since 2007

# New data centre in SFU Water Tower Building



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

SFU Water Tower Building

# ‘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
- 7680 cores at SFU location
- 4744 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

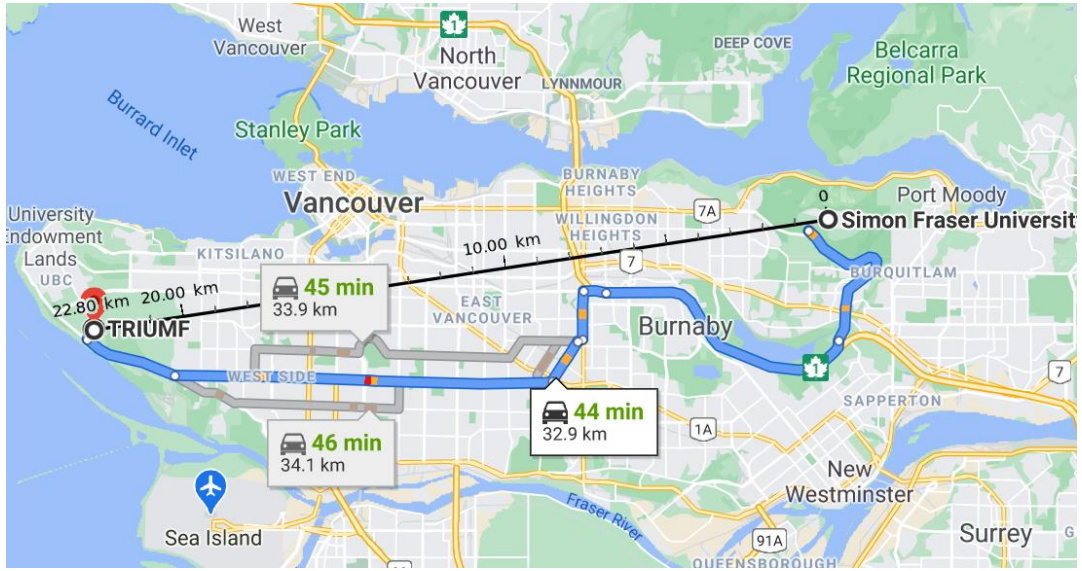


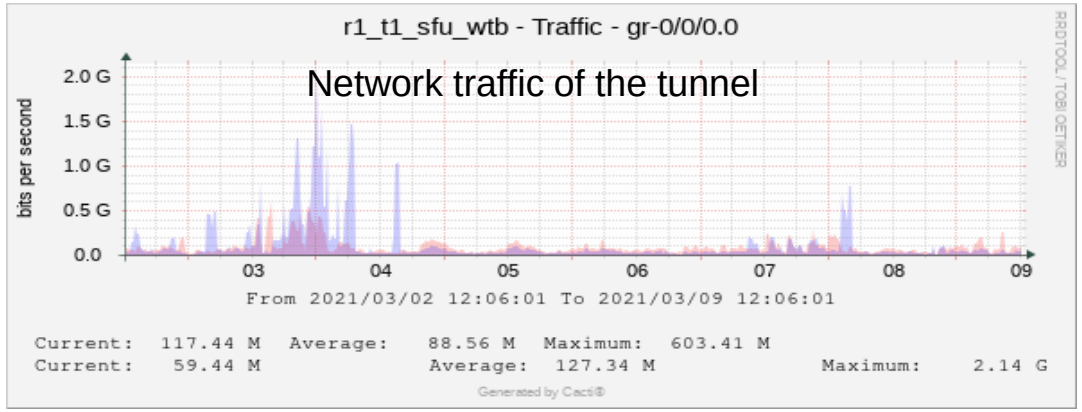

# Tier-1 Networking



- Collaborative effort between TRIUMF, BCNET, CANARIE and SFU
- Juniper QFX10008 as core switch for both internal and external network
- 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

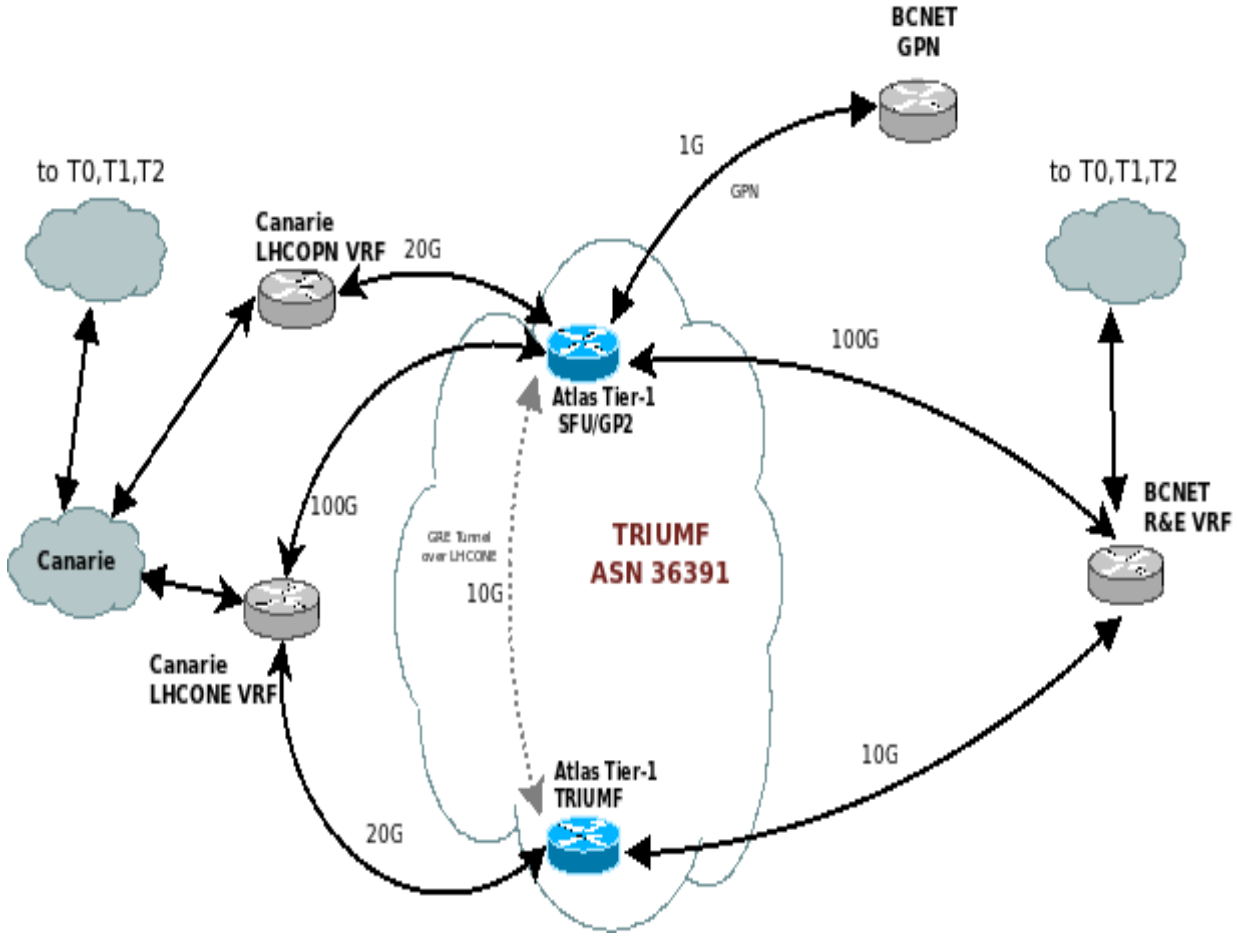

# Tier-1 cluster at SFU



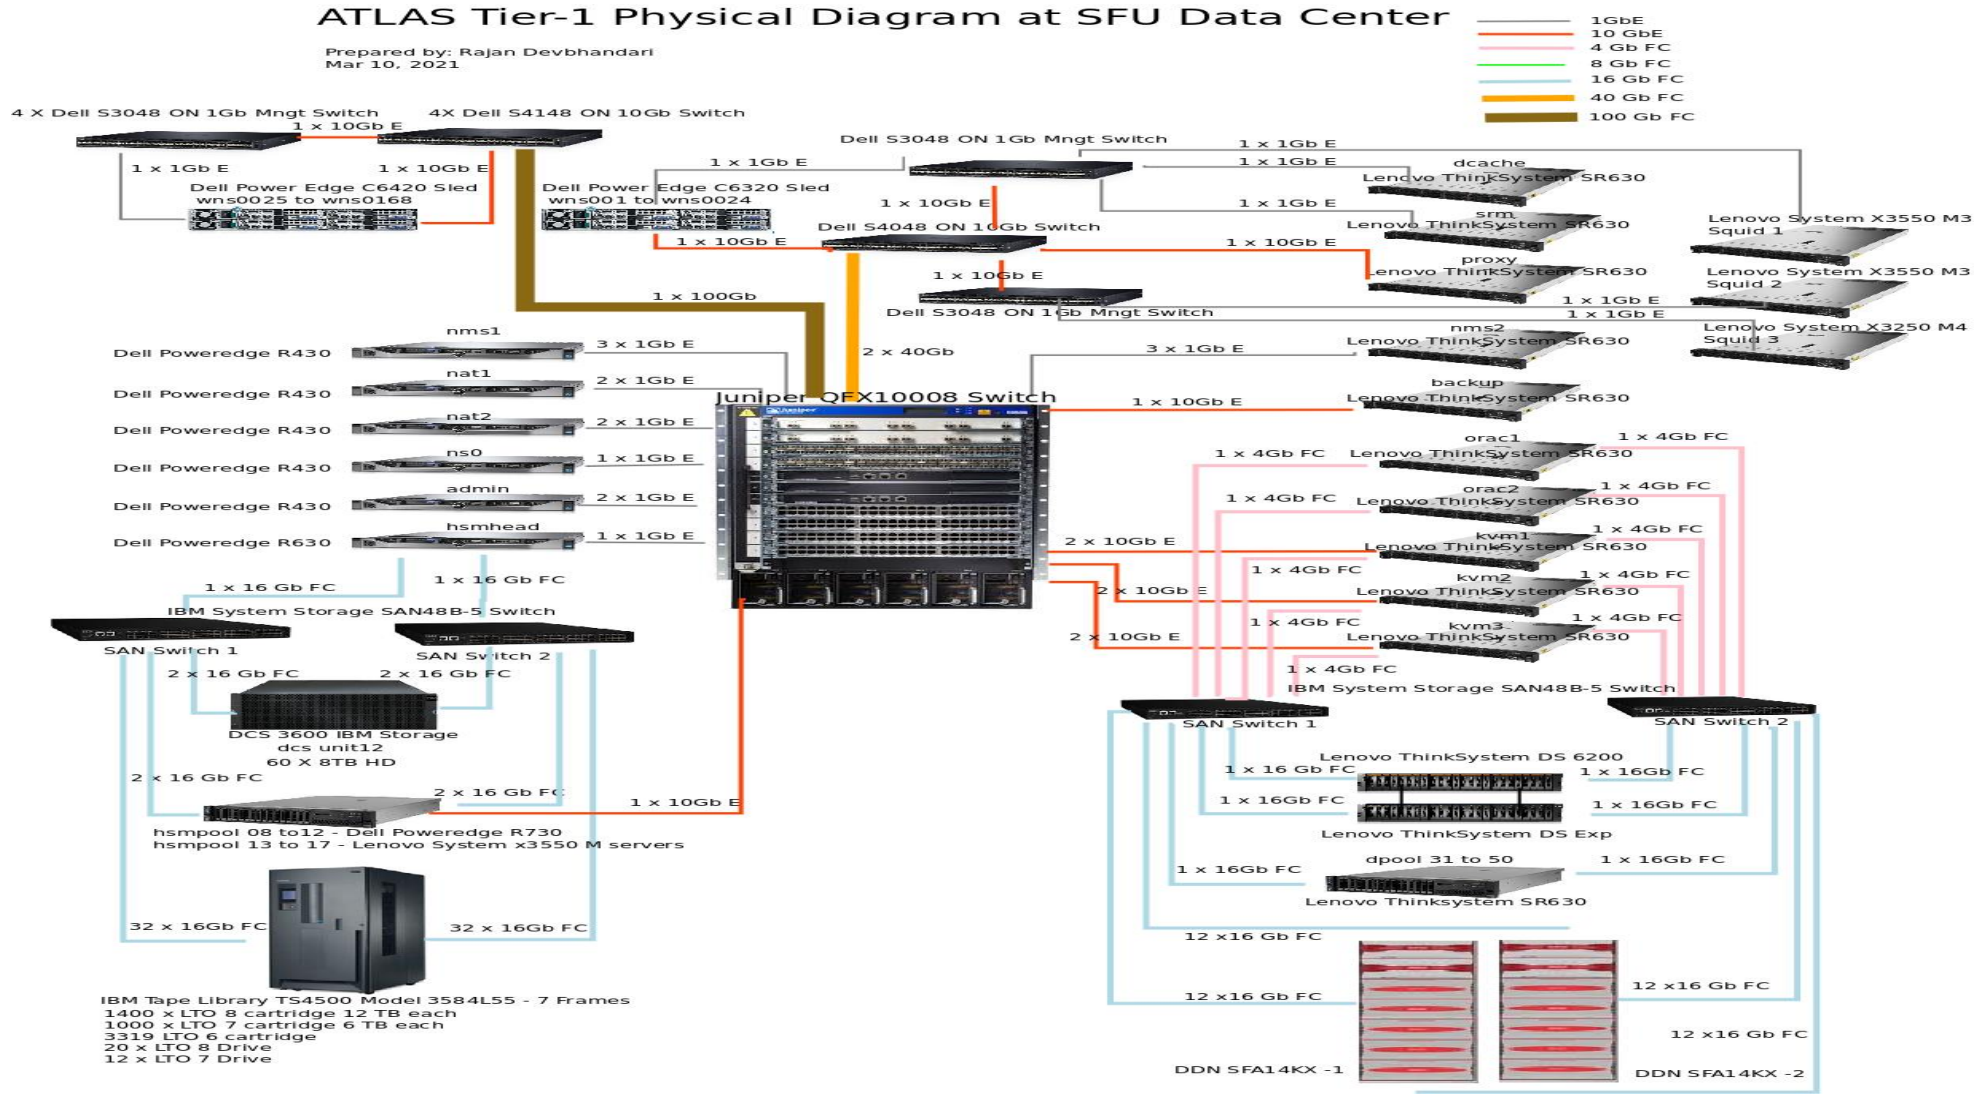

# Storage



- DCache 5.2.35
- 11PB disk
  - 2 X DDN SFA14kx, 2 X 24 X 16Gb Fibre connections
  - 20 Lenovo servers, each server has dual port 16Gb HBA, 16 cores,196GB ram, 40GbE, 10 servers per DDN
- 31PB tape
  - IBM TS4500 tape library
  - 20 LTO8 drivers, 12 LTO7 drivers
  - 3319 LTO6, 1000 LTO7, 1400 LTO8 tape cartridges
  - 800TB disk buffer
  - one hsmhead node,10 hsm pool nodes, 10GbE, dual HBA 16Gb/port
- Webdav and xrootd protocols for internal data access, and gridftp and webdav protocols for external data access

Tape library

# Computing system



- 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
- 7680 cores at SFU
  - 24 DELL C6320 with 32 cores
  - 144 DELL C6420 with 48 cores
  - 4GB RAM/core and ~43GB/core SSD disk space
- 4744 cores at TRIUMF
  - Out of warranty

# Virtualization



- RedHat Enterprise Virtualization (RHEV), a licensed implementation
- 3 hypervisor servers with redundant power supplies and OS disks
  - 2 x Intel Xeon Gold 6138 (20 core) CPUs
  - 192GB RAM
  - 2 x 10 Gb network cards and 1 dual-port QLogic 16 Gb FC HBA
- Fibre-channel based shared storage
  - DS6200 array and expansion unit
  - Mix of SSD and spinning disks
    - ‘Hot’ data stays on high performance disk groups with SSD disk and ‘cold’ data sinks to low performance disk groups
- Majority of non-storage Tier-1 servers are hosted by the virtualization system
  - 37 virtual machines as now
  - Critical services such as three ATLAS frontier server, 4 ARC CEs, CVMFS stratum 1 server for EGI VOs, our web server, GIT and software repo

# RHEV Dashboard



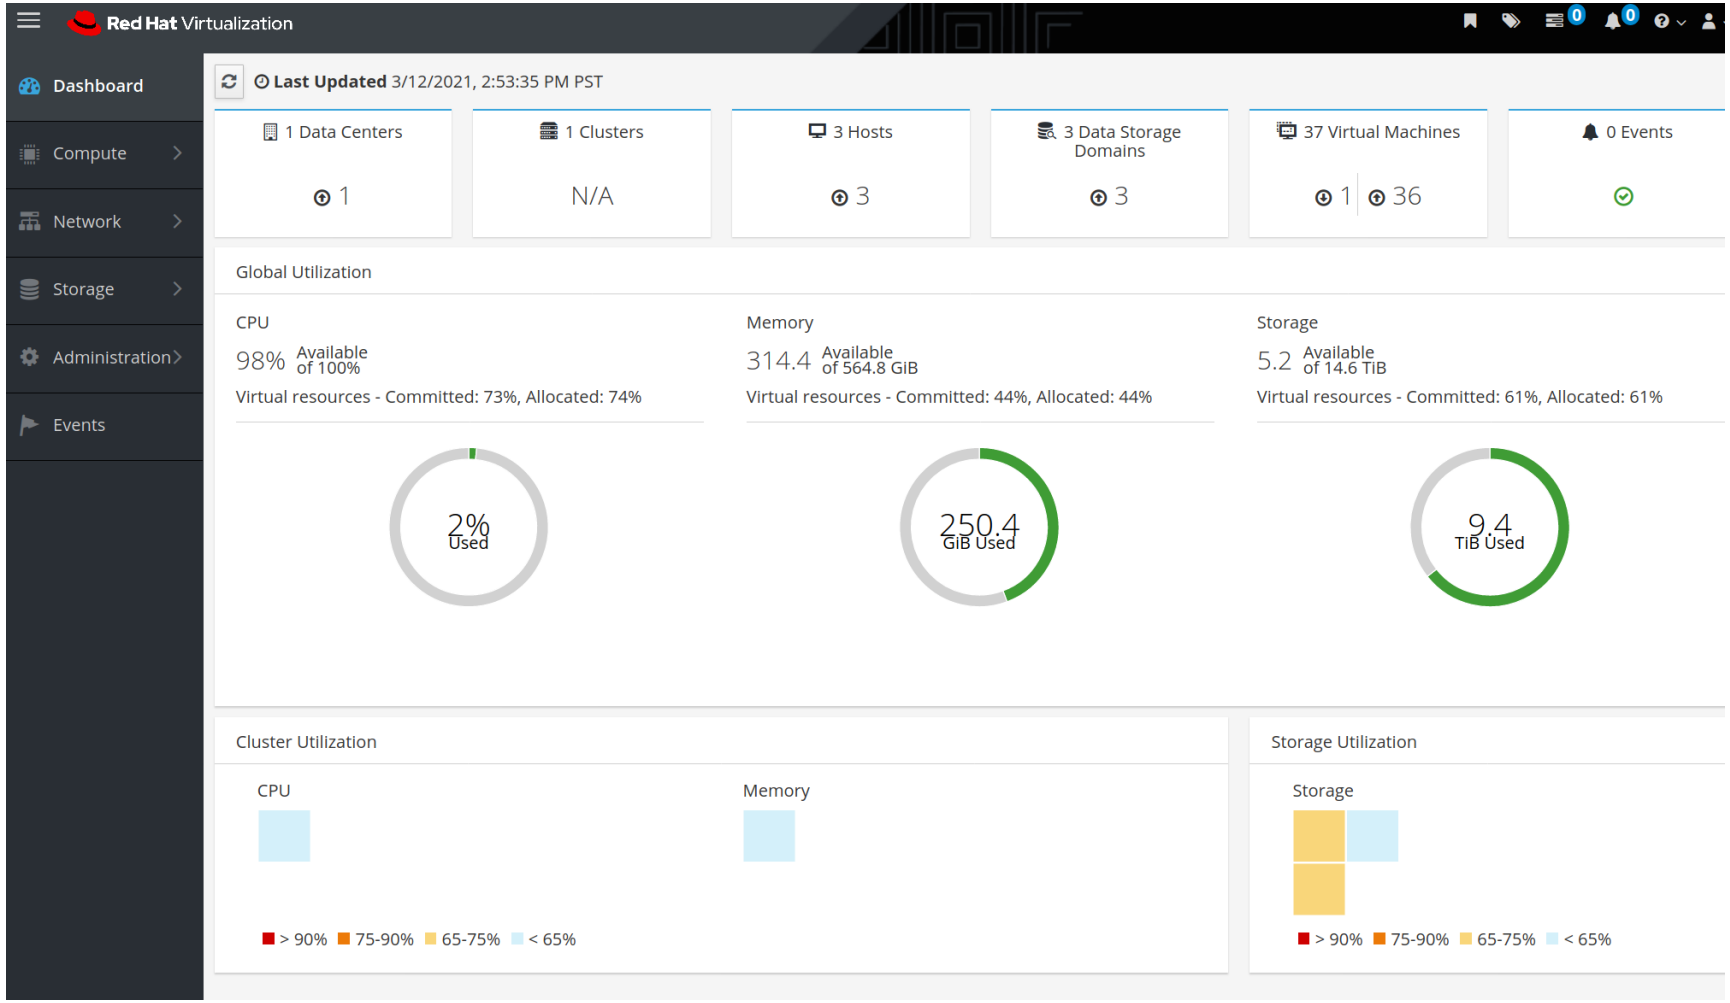

# Monitoring system



- Nagios as the core monitoring system
  - Customized probes created for different Tier-1 system aspects
- Ganglia as a live and recorded statistics of system metrics
- SNMP for device monitoring
- A lot of scripts created to scan and monitor the health of Tier-1 system with CRON and send alerts
- Elastic Search at Tier-1
  - Fully implemented for our system last year
  - Collecting metrics for data transfers, computing resources and syslog etc
  - Visualization with Kibana and Grafana
  - Started doing analysis based on these data, for example, CPU efficiency

# Metrics shipped to Elastic Search

- Ship with Filebeat and process with Logstash all the ".access" logs from dCache doors
  - This gives us insight and information about the connections and their status
  - Combined with the "billing" logs we also get insight into the activity between the doors and the pools for each transfer
- Packetbeat to capture protocol metrics on the doors
- Metrics from dCache PostgreSQL databases and MariaDB/MySQL
- Metrics from compute nodes' SSD disks to give us an idea about their usage and remaining life
- Inlet temperature of Tier-1 devices
- All the syslogs including the logs from router
- Usages of the batch systems

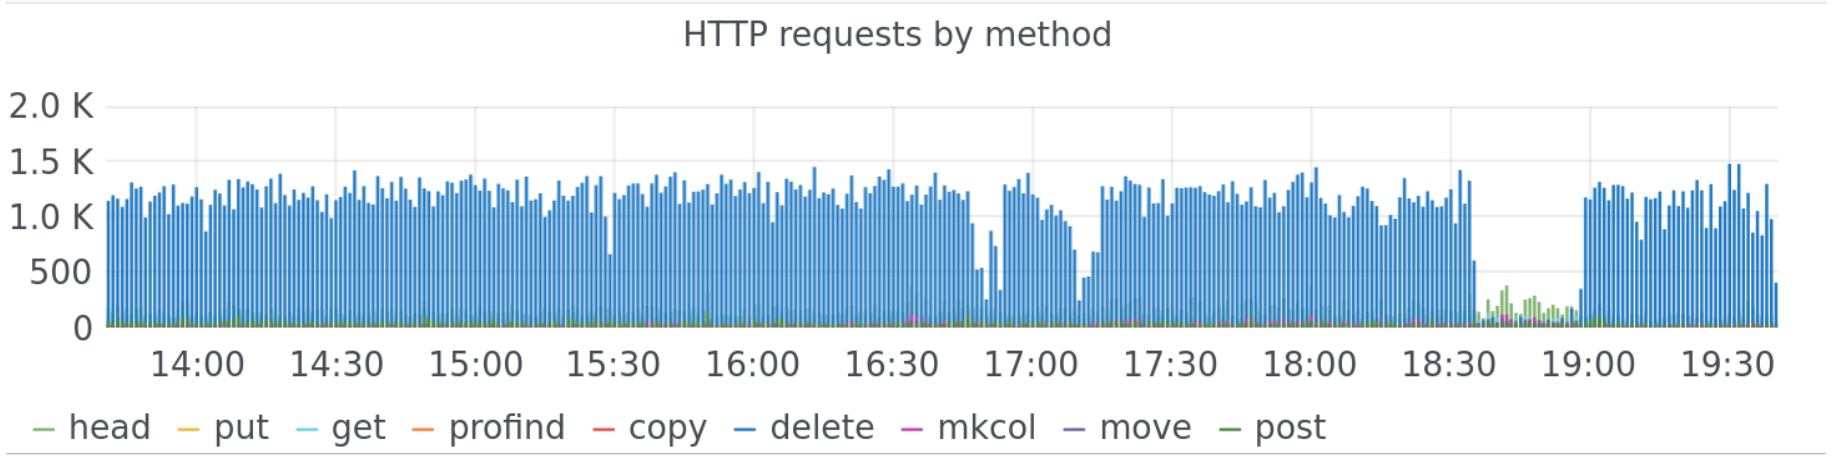

# Database



- Oracle database for ATLAS condition data
  - 2 Oracle RAC servers: 2 x Intel Xeon Gold 6134, 96GB RAM
  - The same shared Fibre-channel based storage as RHEV system
  - ~4TB database
- 3 ATLAS Frontier servers under DNS alias providing ATLAS frontier services to North American sites
- PostgreSQL databases for dCache storage system
- MySQL/MariaDB for other services such as tape system and APEL accounting data for BOINC
- MongoDB for the meta data on disk pools
  - Replaced Berkeley DB

# ATLAS Frontier Server Usage

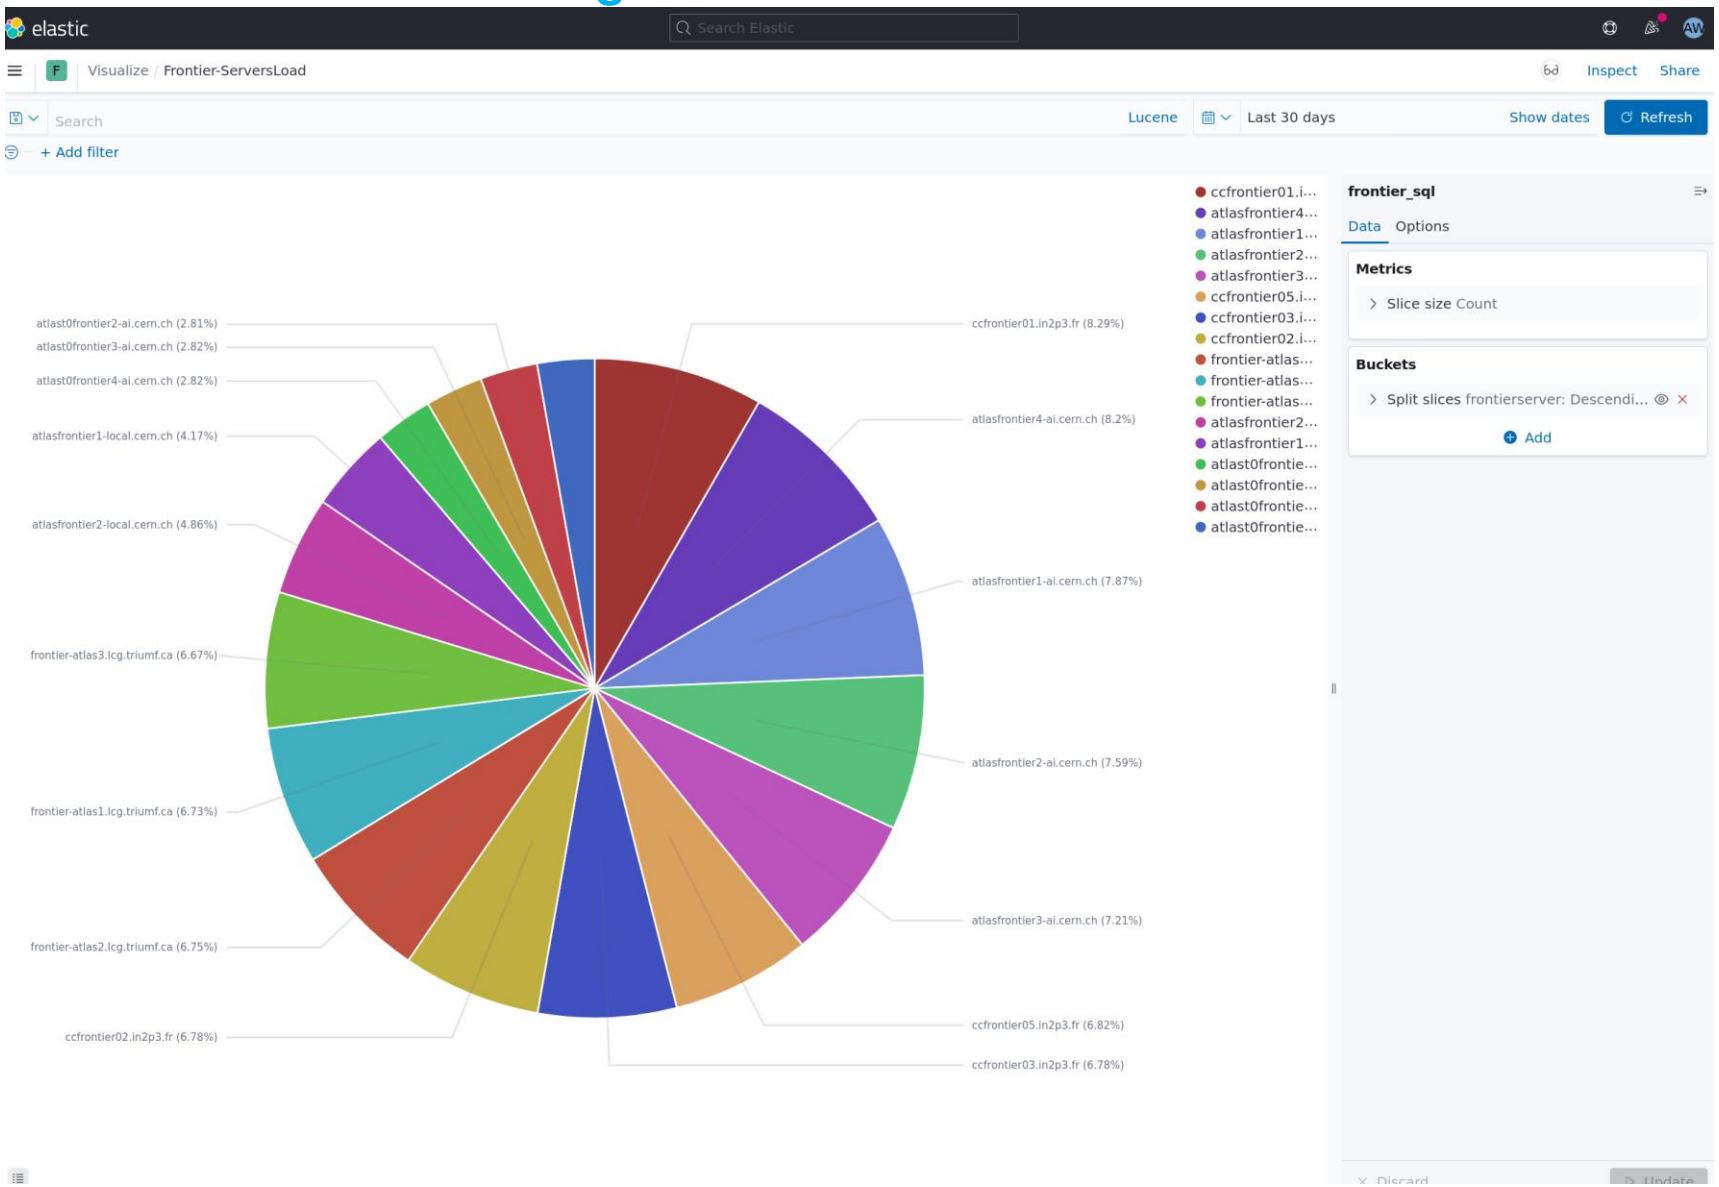

# System administration tools



- Automated Installation Infrastructure of QUATTOR is being used for automated installation
- Ansible with a GIT back-end as configuration management tool
  - All Tier-1 systems are being configured and managed by Ansible
  - ~120 roles created for Tier-1 tasks and configuration since 2016
  - System admins create and maintain the roles for their own areas
- Backup systems
  - Dirvish is our rotated backup system
  - In addition, home-made scripts are doing backups and archiving copy to storage and tape weekly
  - Amanda is still being used on few systems
- System safety
  - 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

# CPU resources are not fully utilized



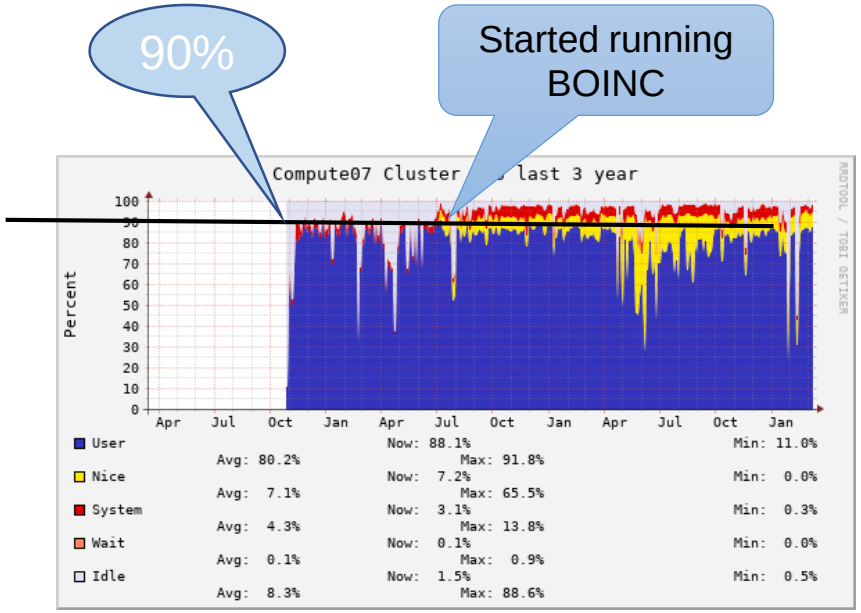


- Average CPU efficiency ~91.3% in 2020
  - CPU utilization rate even lower, ~88.9%
- Reasons
  - Staging in/out data
  - Sequential step of multi-core job
  - No payloads from ATLAS
  - Draining jobs
  - Switching between single core/multi-core productions

# Running ATLAS@home on TRIUMF ATLAS Tier-1 clusters



- ATLAS at home project with BOINC
  - A project to use the internet-connected volunteer resources
  - Run simulation of ATLAS experiment - CPU intensive
  - Now also used by some grid sites and HPC to backfill nodes
- Started in March 2018 at TRIUMF location
- Started to run BOINC jobs at SFU location in June 2019
- Implementations
  - Create account on LHC@home
  - Select ATLAS project
  - Install and configure BOINC client
  - Tune the configurations
    - Avoid the impacts on normal productions
    - Use CGroup to control the CPU share of BOINC jobs
    - Limit the number of cores (25%) which BOINC can use
  - Accounting data is being published to APEL

# What are gained with BOINC jobs

- Backfill nodes

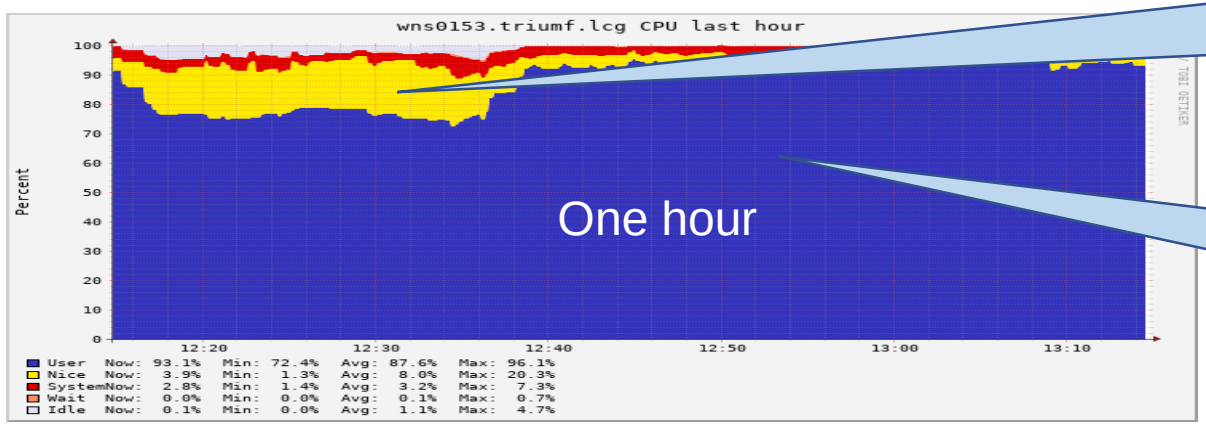


CPU usage of BOINC processes with very low priority

Processes of ATLAS jobs with normal priority

- Increase the CPU utilization of the cluster

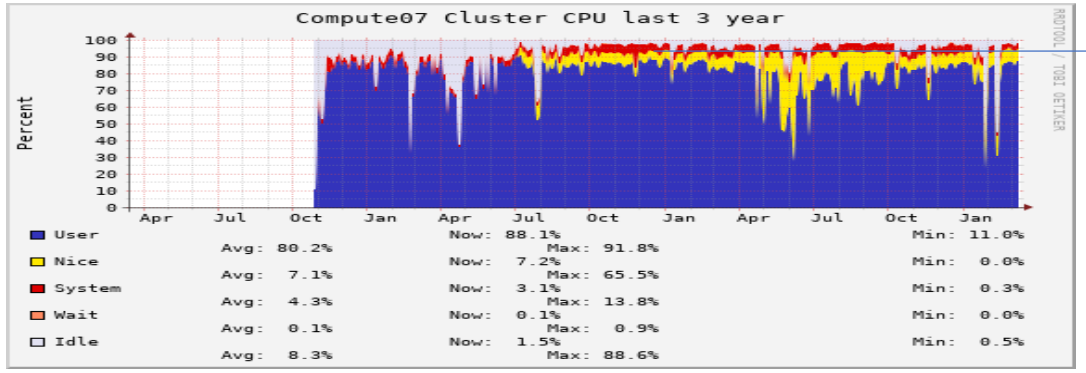


Total CPU usage

# Contribution to ATLAS@home project



- Provided a lot of ‘additional’ resources to ATLAS, ~6.6% CPU time of the cluster at SFU location in 2020


LHC@home Project Computing Community Jobs Site Join Login

| Rank | Name | Recent average credit | Total credit | Country | Participant since |
|---|---|---|---|---|---|
| 1 | Agile Boincers | 9,240,671 | 4,594,076,474 | Switzerland | 20 Sep 2012, 13:19:40 UTC |
| 2 | AGLT2 | 1,464,976 | 2,685,700,886 | United States | 23 Jun 2014, 2:32:15 UTC |
| 3 | TRIUMF-LCG2 | 773,479 | 847,046,975 | Canada | 15 Mar 2018, 21:05:31 UTC |
| 4 | NDGF-T1 | 1,181,065 | 783,100,439 | Norway | 26 Feb 2019, 12:43:24 UTC |
| 5 | Greger | 262,486 | 426,802,288 | Sweden | 9 Jan 2015, 13:08:57 UTC |

TRIUMF Contributions as March 9, 2021

# 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**

Discovery, accelerated