

# TRIUMF Site Report



Denice Deatrach

HEPiX

17 October 2016 Berkeley, CA

Accelerating Science for Canada  
 Un accélérateur de la démarche scientifique canadienne

Owned and operated as a joint venture by a consortium of Canadian universities via a contribution through the National Research Council Canada  
 Propriété d'un consortium d'universités canadiennes, géré en co-entreprise à partir d'une contribution administrée par le Conseil national de recherches Canada



- **TRIUMF Site-wide** (slides from Steve McDonald)
  - CIO, ERP business process management
  - Network status
- **TRIUMF Tier-1**
  - Notable projects in 2016:
    - Ansible for configuration management
    - Docker experience
  - Compute Canada transition
    - Next hardware acquisition
    - Transitional federated site operations

- **New CIO – Davis Swan (Feb 2016)**
  - First TRIUMF CIO appointment
- **Deployed new ERP system** (Finance, Procurement, Logistics)
  - Legacy system was built in-house over 28 yrs
  - Replaced with Agresso Business World – Unit4
  - 5yrs - deployment, no Steve, no HEPiX ☹️
  - Finally launched in Jan 2016
- **Deployed new WiFi** (802.11ac) (indoor & outdoor)
  - Authentication based on 802.11x
  - ~100 Aruba IAP-225's
  - Eduroam on the horizon

- **Deployed new Network Core**
  - Routing Juniper cluster 2x EX-9208 (10G, 100G capable)
  - Firewall Juniper cluster 2x SRX-3400 (10Gbps)
    - all traffic except (LHCOPN, LHCONE)
  - ATLAS Tier-1 network core Juniper EX-9214
- **Network Edge**
  - Migrating from Avaya 5520's to Juniper Ex-3300's
- **External connectivity**
  - 1x 10G LHCOPN
  - 2x 10G LHCONE
  - 1x 10G General research
- **Initiated IPv6 with ATLAS Tier-1**

# Tier-1 Ansible and Git - 1

- Transitioned to Git from CVS in 2015
- Late in 2015 we did an Ansible (from EPEL) evaluation using Git version control and proof-of-concept tests
- Ansible deployment in 2016:
  - 1) Early Ansible roles and Yum playbooks
  - 2) Generic roles to replace old Kickstart scripts
  - 3) Documentation / training
  - 4) Some server and Storage roles added
- Approximately 2 man-months of work

- Current set of 45 Ansible roles:

```

File Edit View Search Terminal Help
Last login: Thu Oct 13 15:24:07 2016 from shrugged.triumf.ca
Agent pid 3484
[root@gridadm ~]# . dcd
[root@gridadm ~]# cd /ks/ansible
[root@gridadm ansible]# ls
Archive/  inventory/  library/  playbooks/  readme.md  roles/  scripts/
[root@gridadm ansible]# ls roles/
anaconda/          docker-registry/  iptables/         network/          ssh/
autofs/            docker-worker/    ks/               ntpd/            sudoers/
base/              ds5020/           kvm-guest/        oracle-java/      syslog/
batch-users/       fetch-crl/        logrotate/        pakiti-mwr/      tcpwrap/
cvmfs/             ganglia/          logs/             pdsh-hosts/      unpriv/
dcache-config/     generic-node/     mail/             postgres/         worker/
dirvish/           grid-certs/       mhvtl/           root/            xrootd/
dns-client/        httpd-2.4/        multipath/        serial-console/  yaim/
docker-distribution/ ipmi/             mysql/            smallhsm/        zol/
[root@gridadm ansible]#
[root@gridadm ansible]# █

```

- Initial Ansible presentation at TRIUMF:  
[http://grid.triumf.ca/share/documents/Ansible\\_presentation.pdf](http://grid.triumf.ca/share/documents/Ansible_presentation.pdf)
- Currently using Ansible 2.1 modified slightly and rebuilt at TRIUMF to preserve file timestamps by default
- Next steps:
  - continue adding new roles, up to now done by 3 people in our group
  - involve 2 more staff
  - finish certifying all roles
  - consider making public our Ansible repository

# Docker for worker nodes - 1

- This Spring we began evaluating Docker as a solution for worker node deployment
- URL describing this work:  
<https://twiki.atlas-canada.ca/bin/view/AtlasCanada/TRIUMFDocker>
- Goal: Run ATLAS SL-6 based Docker containers on SL-7 Docker engine nodes
- Current tests are done running containers with condor-master as the container application



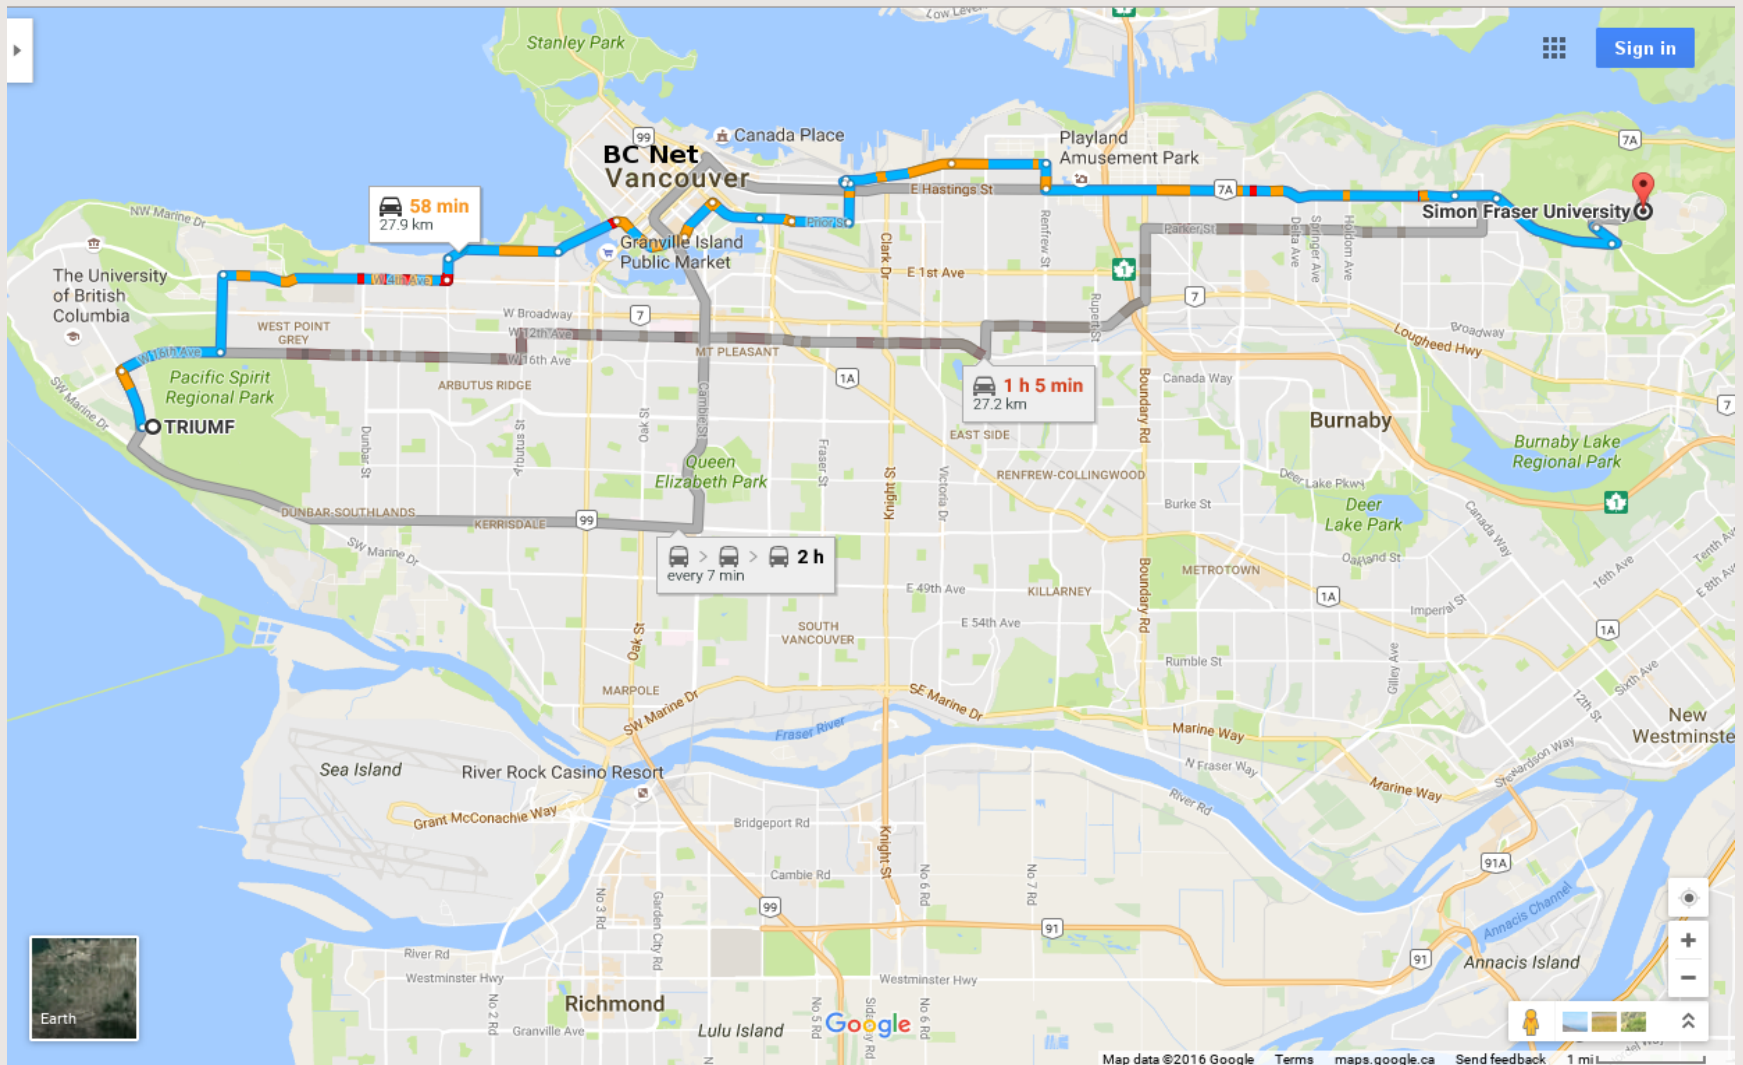
# Docker for worker nodes - 2

- Good results comparing container to host performance with ATLAS software validation
- Good results comparing container to host performance for generic benchmarks:
  - HEPSPEC '06
  - iozone local I/O
  - data transfer protocol tests (gsiftp, dcap, root, http)
- Work to do: Distribution server and image signing, deployment strategies, container logging, image maintenance

# Tier-1 Compute Canada Transition

- **Compute Canada (CC)**
  - a non-profit corporation of dozens of member universities and research institutes
  - is the national advanced research computing (ARC) facility of Canada
  - funded by the Canadian Foundation for Innovation (CFI), which also provides funds for the Tier-1
  - CFI wants the Tier-1 integrated under the Compute Canada umbrella to reduce infrastructure and operating costs
  - the Tier-1 needs space to grow in the next hardware refresh, and will thus transition to the new Compute Canada data centre at Simon Fraser University (SFU) starting in early 2017

# TRIUMF and SFU Locations



# Tier-1 Hardware refresh plans

- Tier-1 data centre at TRIUMF at capacity, the last major refresh was in 2012.
- Tape capacity increased in 2016 by replacing LTO-4 with LTO-6 and adding last frame.
- Current capacity: 4830 cores, 7.8 PB disk (usable), 12 PB tape, and ~85 servers
- Core router replaced with Juniper EX9214 in September 2015
- We are extending warranties for critical hardware into 2018
- New equipment going to SFU, starting with tape services next year, then disk and CPU. Some PO's are imminent – new HSM pools and tape disk SAN buffer space.

# TRIUMF/SFU Tier-1 Federated Site

- We need to operate both sites as one federated site (mostly dcache/storage side) while we make the transition to SFU
  - Phase 1
    - Set up pre-production site at SFU - network, basic administrative functions, test setup of grid and storage services on new servers
    - commissioning of new tape library and federation testing
  - Phase 2
    - full production mode with new tape-destined data going only to SFU; start tape migration
    - gradually bring online new disk and CPU resources with disk migration
  - Phase 3
    - gradual diminution of TRIUMF-side services until final decommissioning in 2018 before the end of warranties

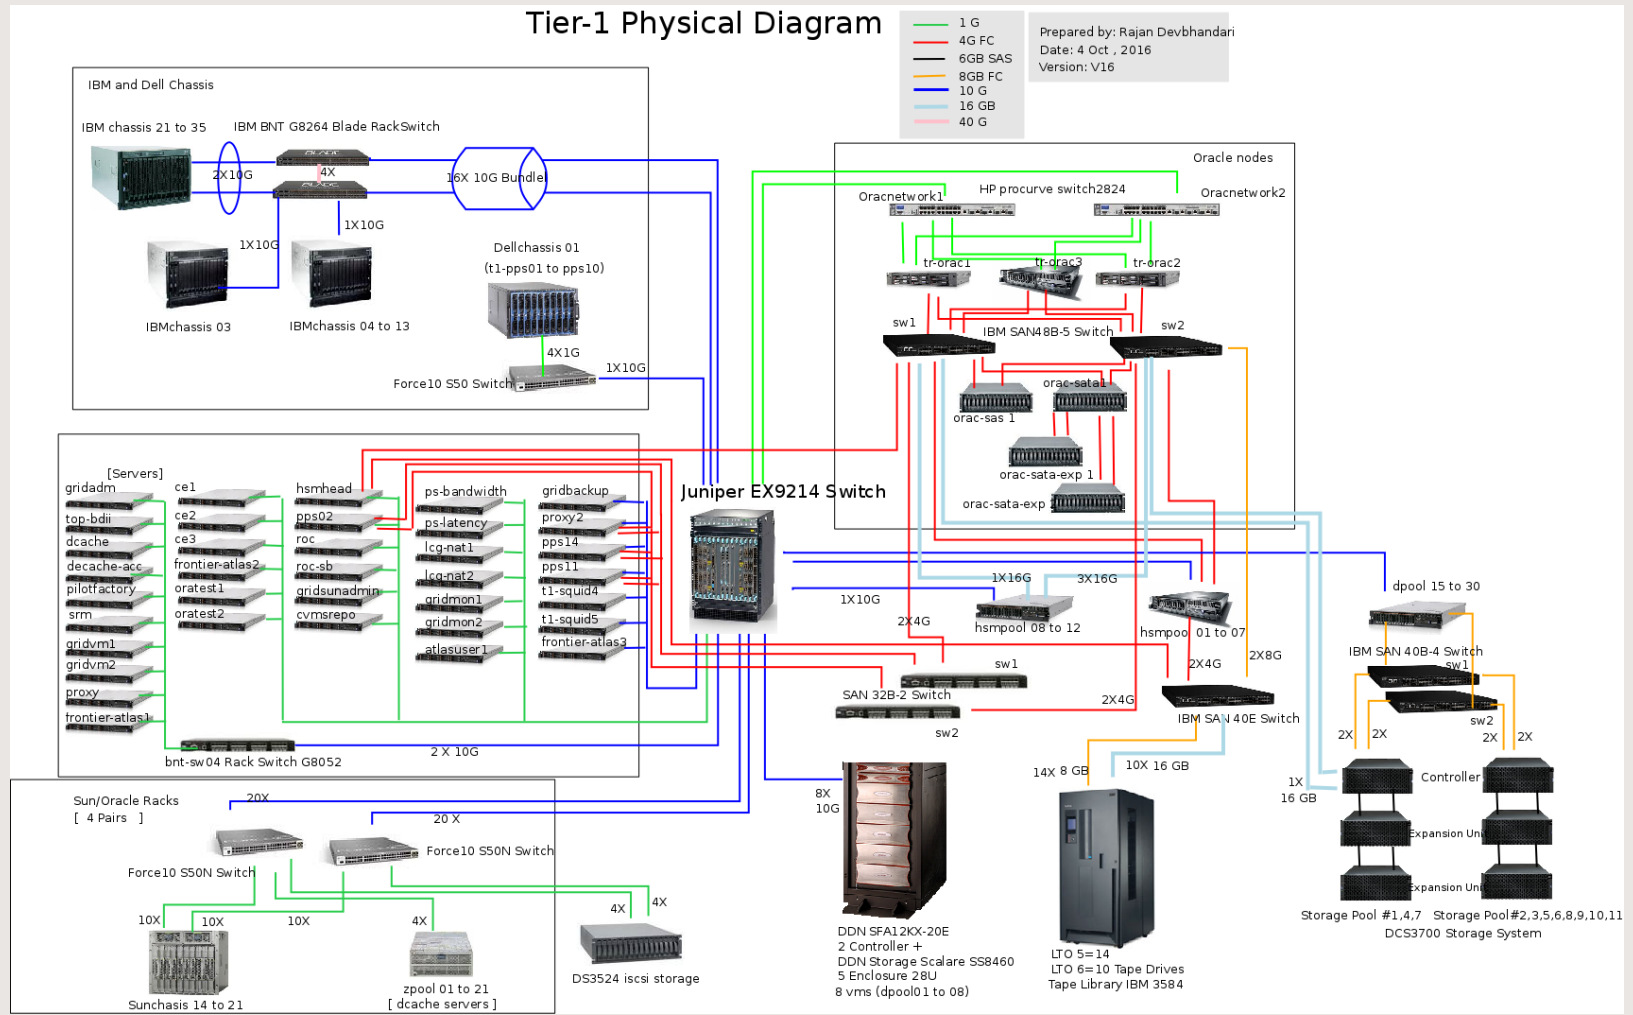
# Thank you!

# Merci

TRIUMF: Alberta | British Columbia |  
 Calgary | Carleton | Guelph | Manitoba |  
 McGill | McMaster | Montréal | Northern  
 British Columbia | Queen's | Regina |  
 Saint Mary's | Simon Fraser | Toronto |  
 Victoria | Western | Winnipeg | York



# Tier-1 Physical Diagram



# Tier-1 Network Topology

